Monitoring Microsoft Applications

eG Enterprise v6

Restricted Rights Legend

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

Trademarks

Microsoft Windows, Windows NT, Windows 2003, and Windows 2000 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

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

Copyright

Table of Contents

Table of Figures

Chapter 1

Introduction

Microsoft applications are common-place in IT infrastructures today. From web interfaces to domain controllers to authentication servers to Microsoft RDS servers to simple browsers, a wide range of Windows-based applications are being increasingly utilized by infrastructure operators to keep the IT environment afloat and easily accessible to end-users.

This means that even a slight slowdown in the performance of one of these applications, if not resolved soon, can prove to be fatal to the critical end-user service riding on it. This is reason enough for bringing Microsoft applications under the purview of '24x7 monitoring'.

eG Enterprise provides 100% web-based monitoring models to continuously monitor and report on the status of critical Microsoft applications such as Active Directory servers, Microsoft RDS servers, Windows Domain Controllers, etc.

This document describes the monitoring model that eG Enterprise prescribes for every Microsoft application, and the performance metrics each model collects.

Chapter 2

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 2.1), which defines the key performance metrics that need to be tracked to determine the service level achieved by the server/server farm.

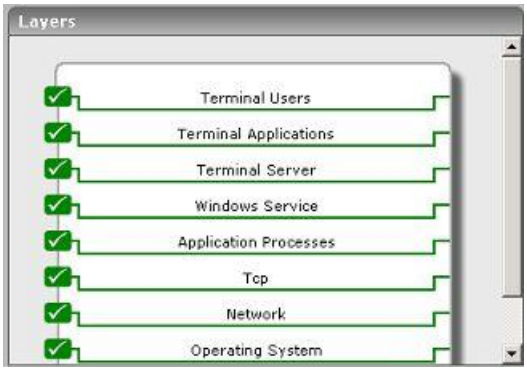

Figure 2.1: Layer model of a Microsoft RDS server

Using the metrics reported by each of the layers depicted by Figure 2.1, 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 2.1 have already been discussed in the *Monitoring Unix and Windows Servers* document, the sections to come will discuss the top 4 layers only.

2.1 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.

2.1.1 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.

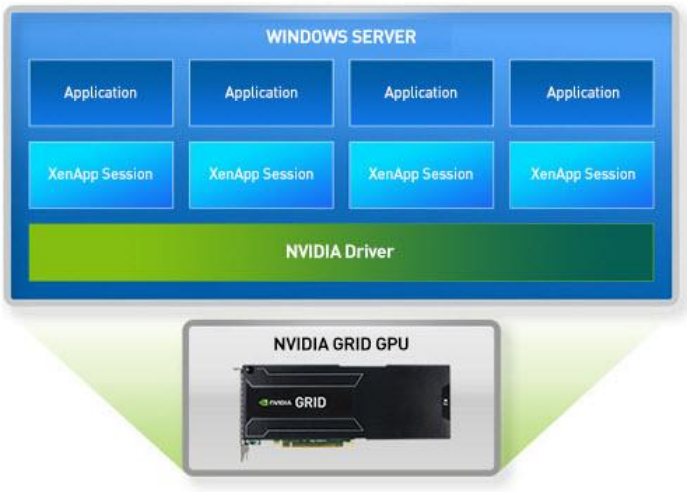

Figure 2.2: Architectural diagram for NVIDIA GRID with XenApp

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.

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.

Purpose	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 VMs/virtual dekstops/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.
Target of the test	A Citrix XenApp server / Microsoft RDS
Agent deploying the test	An internal/remote agent

<table>
<tr><td>Configuration Parameters</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Citrix server
4. 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/NVSMI directory. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each GRID physical GPU assigned to the host being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>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.</td><td>Percent</td><td>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 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.</td></tr>
<tr><td>Used frame buffer memory:
Indicates the amount of frame buffer memory on-board this GPU that is being used by the host.</td><td>MB</td><td>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.</td></tr>
</table>

	Free frame buffer memory: Indicates the amount of frame buffer memory on-board this GPU that is yet to be used by the host.	MB	
	Frame buffer memory utilization: Indicates the percentage of frame buffer memory on-board 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 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 GPU that is currently being used by the host.	MB	For better user experience with graphic applications, enough BAR1 memory should be available to the host.
	Free BAR1 memory: Indicates the amount of BAR1 memory of this GPU that is still to be used by the host.	MB	
	BAR1 memory 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.

<table>
<tr>
<td></td>
<td>Power management:
Indicates whether/not power management is enabled for this GPU.</td>
<td></td>
<td>Many NVIDIA graphics cards support multiple performance levels so that the server can save power when full graphics performance is not required.

The default Power Management Mode of the graphics card is Adaptive. In this mode, the graphics card monitors GPU usage and seamlessly switches between modes based on the performance demands of the application. This allows the GPU to always use the minimum amount of power required to run a given application. This mode is recommended by NVIDIA for best overall balance of power and performance. If the power management mode is set to Adaptive, the value of this measure will be Supported.

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

The numeric values that correspond to these measure values are discussed in the table below:
<table>
<tr><th>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>
</table>

	Power limit: Indicates the power limit configured for this GPU.	Watts	**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 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.
	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
	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	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-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.

	Fan speed: Indicates the percent of maximum speed that this GPU's fan is currently intended to run at.	Percent	The value of this measure could range from 0 to 100%. An abnormally high value for this measure could indicate a problem condition – eg., a sudden surge in the temperature of the GPU that could cause the fan to spin faster. Note that the reported speed is only the intended fan speed. If the fan is physically blocked and unable to spin, this output will not match the actual fan speed. Many parts do not report fan speeds because they rely on cooling via fans in the surrounding enclosure. By default the fan speed is increased or decreased automatically in response to changes in temperature.
	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 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.

2.2 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 2.1: The tests mapped to the Windows Service layer

2.2.1 App-V Client Admin Log Test

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

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

Purpose	Reports the statistical information about the admin events generated by the target system
Target of the test	An App-V Client on the target Microsoft RDS Server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT** – Specify the port at which the specified **HOST** listens to. By default, this is 8080. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. The default value is *Microsoft-AppV-Client/Admin*. 5. **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. 6. **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 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 (;).

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 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.

	7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs 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'.** 8. **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. 9. **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**. 10. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.

<table>
<tr><td></td><td colspan="3">11. 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 DDFREQ.
12. 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:
o The eG manager license should allow the detailed diagnosis capability
o Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the App-V Client that is to be monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Information messages:
Indicates the number of App-V Client admin information events generated when the test was last executed.</td><td>Number</td><td>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 admin logs in the Event Log Viewer for more details.</td></tr>
<tr><td>Warnings:
Indicates the number of App-V Client admin warnings that were generated when the test was last executed.</td><td>Number</td><td>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.</td></tr>
<tr><td></td><td>Error messages:
Indicates the number of App-V Client admin error events that were generated during the last measurement period.</td><td>Number</td><td>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.</td></tr>
</table>

	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 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.

2.2.2 App-V Client Operational Log Test

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

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

Purpose	Reports the statistical information about the operation events generated by the target system
Target of the test	An App-V Client on the target Microsoft RDS Server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT** – Specify the port at which the specified **HOST** listens to. By default, this is 8080. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. The default value is *Microsoft-AppV-Client/Operational*. 5. **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. 6. **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 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 (;).

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 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.

7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs 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'.**

8. **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.

9. **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**.

10. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.

11. **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 **DDFREQ**.

12. **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.

Outputs of the test	One set of results for the App-V Client that is to be monitored		
Measurements made by the test	**Measurement**	**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 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 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.

2.2.3 App-V Client Virtual Application Log Test

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

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

Purpose	Reports the statistical information about the virtual application events generated by the target system
Target of the test	An App-V Client on the target Microsoft RDS Server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT** – Specify the port at which the specified **HOST** listens to. By default, this is 8080. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. The default value is *Microsoft-AppV-Client/Virtual Applications*. 5. **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. 6. **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 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 (;).

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 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.

7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs 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'.**

8. **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.

9. **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**.

10. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.

11. **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**.

12. **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.

Outputs of the test	One set of results for the App-V Client that is to be monitored		
Measurements made by the test	**Measurement**	**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.
	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.

2.3 The Remote Desktop Services Layer

The tests associated with this layer (see Figure 2.3) 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

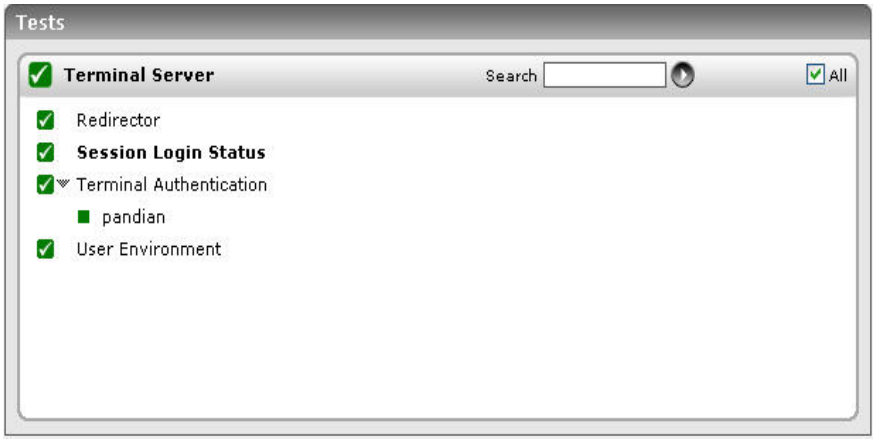

Figure 2.3: Tests associated with the Remote Desktop Services layer

2.3.1 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.

Purpose	Periodically reports the status of logons from client sessions to the Citrix / Microsoft RDS server
Target of the test	A Microsoft RDS server
Agent	An internal agent

deploying the test			
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - Host name of the server for which the test is to be configured 3. **PORT** - Enter the port to which the **HOST** listens		
Outputs of the test	One set of results the Microsoft RDS server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Session login status: Indicates whether the client sessions to the server are currently enabled or not.	Percent	If the value for this measure is *100*, it indicates that client logons are 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 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. 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*.

2.3.2 Terminal Connection Test

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

Purpose	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

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - Host name of the server for which the test is to be configured 3. **PORT** - Enter the port to which the specified **TARGETHOST** listens 4. **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.		
Outputs of the test	One set of results for every port being monitored		
Measurements made by the test	**Measurement**	**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.

2.3.3 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.

Purpose	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

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">
1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. PASSWORD - Enter the password that corresponds to the specified USERNAME.
6. CONFIRM PASSWORD – Confirm the password by retyping it here.
7. 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/or multiple 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 Environments document.

8. 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 configured for this test. 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.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every user account being checked</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Authentication status:
Indicates whether the login was successful or not</td><td>Percent</td><td>A value of 100 % indicates that the login has succeeded. The value 0 is indicative of a failed login.</td></tr>
<tr><td>Authentication time:
Indicates the time it took to login</td><td>Secs</td><td>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.</td></tr>
</table>

2.3.4 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).

Purpose	Monitors the activity of redirector component of the Microsoft windows operating system and tracks the status of the file serving as seen by a file server's client.		
Target of the test	Any Microsoft RDS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Microsoft RDS server		
Outputs of the test	One set of results for the Microsoft RDS server being monitored		
Measurements made by the test	**Measurement**	**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 service.	Number	The **Current Commands** measure indicates the number of pending commands from the local computer to all destination servers. This 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 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.

2.3.5 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 UserProfile test monitors the size of the Microsoft RDS server user profiles and raises an alarm if the profile size exceeds the profile quota size.

Purpose	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	Any Microsoft RDS server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Microsoft RDS server 4. **PROFILESIZELIMIT** - Specify the profile quota size (in MB). The default value is 50 MB. 5. **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. 6. **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. 7. **FILESIZELIMIT** - Takes the file quota size (in KB). The default size is 10000 KB. 8. **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**. 9. **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. 10. **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 strin *Favo*, and with names that end with the string *Destop*, will be excluded from size computation. Moreover, all sub-folders within these folders will also be ignored during size computation.

	11. **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.		
Outputs of the test	One set of results for every user profile on the Microsoft RDS server monitored		
Measurements made by the test	**Measurement**	**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 profile size.	MB	
	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**.

2.3.6 User Environment 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.

The UserEnvironment 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 *Microsoft Terminal* as the **Component type**, *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.

Purpose	Periodically checks the userenv log file to monitor the user login and profile loading process
Target of the test	Any Microsoft RDS server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Microsoft RDS server monitored</td></tr>
<tr><td rowspan="5">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Profile load starts:
Indicates the number of profile loads in the last measurement period.</td><td>Number</td><td>This metric gives an idea of the rate at which users are logging in to the server.</td></tr>
<tr><td>Profile load successes:
Indicates the number of successful profile loads in the last measurement period.</td><td>Number</td><td></td></tr>
<tr><td>Profile loading failures:
Indicates the number of profile load failures in the last measurement period.</td><td>Number</td><td>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.</td></tr>
<tr><td>Profile load failures percent:
Indicates the percentage of profile loads that failed in the last measurement period.</td><td>Percent</td><td></td></tr>
<tr><td></td><td>Avg user profile load time:
Indicates the average time it took to load a profile successfully in the last measurement period.</td><td>Secs</td><td></td></tr>
</table>

	Max profile load time: Indicates the maximum time it took to load a profile during the last measurement period.	Secs	
	System policy starts: Indicates the number of system policy applications started in the last measurement period.	Number	
	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.	Secs	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.	Secs	
	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.	Secs	
	Max group policy time: Indicates the average time taken for applying group policies.	Secs	

	Profile unload starts: Indicates the number of profile unloads started during the last measurement period.	Number	
	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.	Secs	
	Max profile unload time: Indicates the maximum time for unloading a profile during the last measurement period.	Secs	

2.3.7 Terminal Server 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 *Microsoft Terminal* as the **Component type**, *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.

<table>
<tr><td>Purpose</td><td colspan="3">Reports the usage statistics pertaining to a Microsoft RDS server's client access licenses</td></tr>
<tr><td>Target of the test</td><td colspan="3">A Microsoft RDS server</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Microsoft RDS server being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Active licenses:
Represents number of active client access licenses that were currently consumed from the Microsoft RDS server license server.</td><td>Number</td><td>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 target server, the client from which the session was instantiated, etc.</td></tr>
<tr><td>Temporary licenses:
Indicates the number of temporary client access licenses that were currently consumed from the Microsoft RDS server license.</td><td>Number</td><td>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.</td></tr>
</table>

2.3.8 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).

<table>
<tr><td>Purpose</td><td colspan="3">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</td></tr>
<tr><td>Target of the test</td><td colspan="3">A Microsoft RDS server</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each user to the Microsoft RDS server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total GDI objects:
Indicates the total number of GDI handles that this user has created.</td><td>Number</td><td>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.</td></tr>
</table>

	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 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.

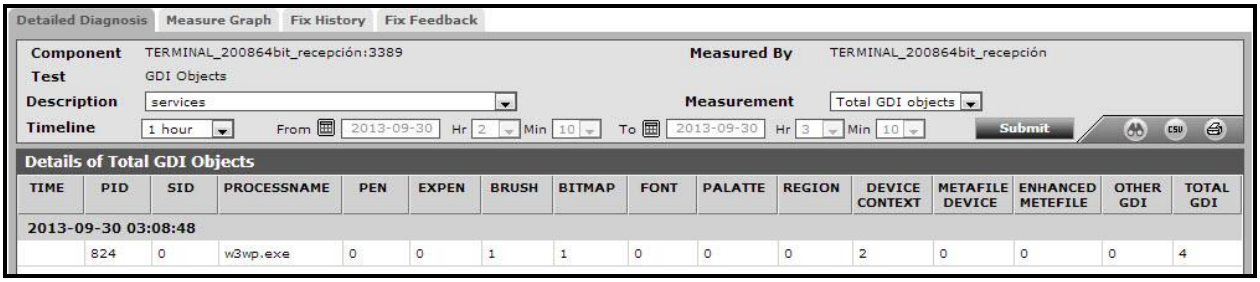

Figure 2.4: The detailed diagnosis of the *Total GDI objects* measure

2.3.9 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.

Purpose	Tracks the status of the default listener ports and reports whether any of the ports is down
Target	A Microsoft RDS server
Agent deploying this test	Internal agent
Configurable	1. **TEST PERIOD** - How often should the test be executed

<table>
<tr><td>parameters for this test</td><td colspan="3">2. HOST - The host for which the test is to be configured.
3. PORT - The port at which the HOST listens
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of outputs for every listener port configured</td></tr>
<tr><td>Measurements of the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Is listener down?:
Indicates whether/not this listener port is down.</td><td></td><td>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:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>0</td></tr><tr><td>No</td><td>1</td></tr></table>
Note:
By default, this measure reports the above-mentioned Measure Values to indicate the status of a listener port. However, the graph of this measure will represent the same using the numeric equivalents only.</td></tr>
</table>

2.3.10 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 Citrix, 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?

By default, this test is disabled. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *Microsoft RDS* as the **Component type**, *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.

<table>
<tr><td>Purpose</td><td colspan="3">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?</td></tr>
<tr><td>Target</td><td colspan="3">A Microsoft RDS server</td></tr>
<tr><td>Agent deploying this test</td><td colspan="3">Internal agent</td></tr>
<tr><td>Configurable parameters for this test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured.
3. PORT - The port at which the HOST listens
5. 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.
6. 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.
7. 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</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every user to the Microsoft RDS server monitored</td></tr>
<tr><td>Measurements of the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

	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 *User account discovery, LDAP bind time to active directory, Client side extension processed time, DC discovery time, Total group policy object file accessed time* and *User profile duration* measures to know exactly where that user's login process experienced a bottleneck - is it when loading the profile? is it when processing system policies? is it when processing group policies? is it when interacting with AD for authenticating the user login?
	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 the value of this measure is *0* for a unusually longer duration, then administrators should verify if the pre-requisites are fulfilled. Once the pre-requisites are fulfilled, this measure will report the values using which administrators can accurately identify where exactly has the logon process stalled.
	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	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. 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.
	Group policy time: Indicates the time taken for applying group policies for this user in the last measurement period.	Secs	This measure enforces the domain policy and settings for the user session and in case of synchronous processing, the user will not see their desktop at logon until user GP processing completes. If the value of this measure is *0* for a unusually longer duration, then administrators should verify if the pre-requisites are fulfilled. Once the pre-requisites are fulfilled, this measure will report the values using which administrators can accurately identify where exactly has the logon process stalled.

	GP scripts duration: Indicates the time taken for executing group policy scripts for this user in the last measurement period.	Secs	The output of the script will showow if the Group policy scripts were executed in synchrnous mode or asyynchronous mode. When running the logon scripts configured in the relevant GPO's, in a case of synchronous logon scripts, Windows Explorer does not start until the logon scripts have finished running. If the value of this measure is *0* for a unusually longer duration, then administrators should verify if the pre-requisites are fulfilled. Once the pre-requisites are fulfilled, this measure will report the values using which administrators can accurately identify where exactly has the logon process stalled.
	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*. A low value is desired for this measure. A sudden/gradual increase in the value of this measure is a cause of concern. If the value of this measure is 0 for a unusually longer duration, then administrators should verify if the pre-requisites are fulfilled. Once the pre-requisites are fulfilled, this measure will report the values using which administrators can accurately identify where exactly has the logon process stalled.
	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 during the last measurement period.	Secs	A low value is desired for this measure. If the value of this measure is 0 for a unusually longer duration, then administrators should verify if the pre-requisites are fulfilled. Once the pre-requisites are fulfilled, this measure will report the values using which administrators can accurately identify where exactly has the logon process stalled.

<table>
<tr>
<td></td>
<td>Group policy:
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:
By default, this measure reports the above-mentioned Measure Values 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., 1to 3.</td>
</tr>
<tr>
<td></td>
<td>User account discovery:
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.</td>
<td>Secs</td>
<td>Compare the value of this measure across users to know which user’s logon process spent maximum time in retrieving account information.</td>
</tr>
<tr>
<td></td>
<td>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.</td>
<td>Secs</td>
<td>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.</td>
</tr>
<tr>
<td></td>
<td>DC 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.</td>
<td>Secs</td>
<td>Compare the value of this measure across users to know which user’s logon process spent maximum time in domain controller discovery.</td>
</tr>
<tr>
<td></td>
<td>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.</td>
<td>Secs</td>
<td>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.</td>
</tr>
</table>

	Number of CSE applied: Indicates the total number of client side extensions used for processing group policies for this user during the last measurement period.	Number	
	Number of CSE success: Indicates the number of client side extensions that were successfully used for processing group policies for this user during the last measurement period.	Number	
	Number of CSE warning: 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	
	Number of CSE error: 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.

	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.	Secs	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.

2.4 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.

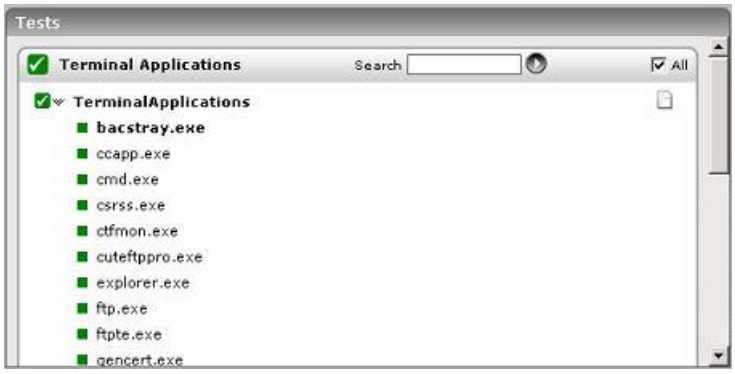

Figure 2.5: Tests associated with the Terminal Applications layer

2.4.1 Terminal Applications Test

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

This test will report metrics only if the Microsoft RDS server being monitored uses the .Net framework v3.0 (or above).

Purpose	Returns the performance measures pertaining to the applications published on the Microsoft RDS server
Target of the test	A Microsoft RDS server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Microsoft RDS server 4. **APPS** - By default, **all** is displayed here, which will auto-discover and monitor all the 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. 5. **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 accessed an application on the server. This way, administrators will be able to quickly determine which user logged into the server from which domain. If you want the detailed diagnosis to display only the *username* of these users, set this flag to **No**. 6. **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. 7. **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.
Outputs of the test	One set of results is reported for each application

Measurements made by the test	Measurement	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.

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.

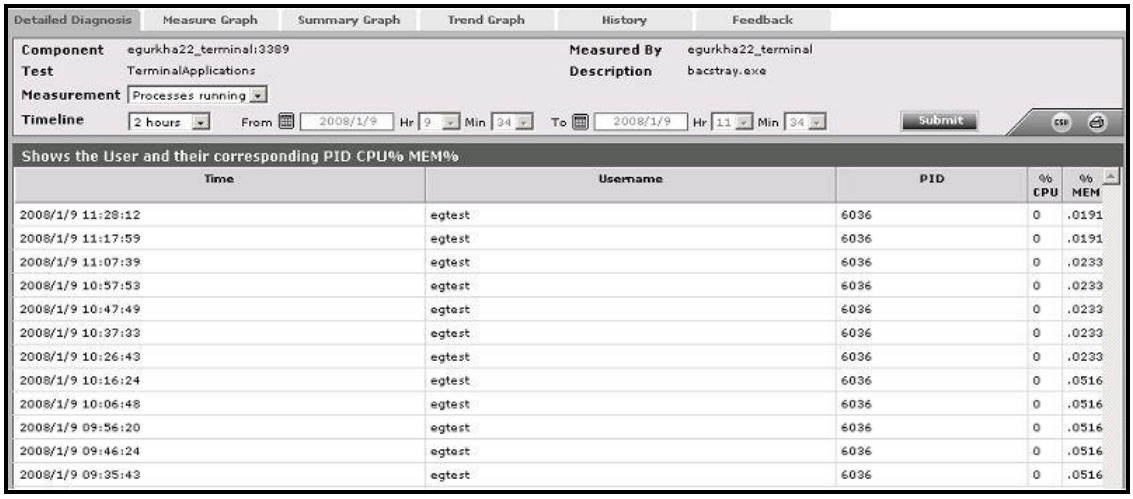

Time	Username	PID	% CPU	% MEM
2008/1/9 11:28:12	egtest	6036	0	.0191
2008/1/9 11:17:59	egtest	6036	0	.0191
2008/1/9 11:07:39	egtest	6036	0	.0233
2008/1/9 10:57:53	egtest	6036	0	.0233
2008/1/9 10:47:49	egtest	6036	0	.0233
2008/1/9 10:37:33	egtest	6036	0	.0233
2008/1/9 10:26:43	egtest	6036	0	.0233
2008/1/9 10:16:24	egtest	6036	0	.0516
2008/1/9 10:06:48	egtest	6036	0	.0516
2008/1/9 09:56:20	egtest	6036	0	.0516
2008/1/9 09:46:24	egtest	6036	0	.0516
2008/1/9 09:35:43	egtest	6036	0	.0516

Figure 2.6: 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.

- 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**.

2.4.2 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.

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

Purpose	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.
Target of the test	An App-V Client on the target Mincorsoft RDS server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – The port at which the specified HOST listens. By default, this is NULL.
4. 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.
5. 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 the EXTENDED STATISTICS flag to No.
6. 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:
o The eG manager license should allow the detailed diagnosis capability
o Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each application of the target App-V Client that is to be monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total size:
Indicates the total size of this virtual application package.</td><td>MB</td><td>The detailed diagnosis of this measure lists the Version of the application, Application ID, Version ID of the applicaition and the application path.</td></tr>
</table>

<table>
<tr>
<td></td>
<td>Is loading?:
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></td>
<td>Loaded percentage:
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></td>
<td>In use?:
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 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>
</table>

<table>
<tr>
<td></td>
<td>Any user based pending tasks available?
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 using the numeric equivalents - 0 or 1.</td>
</tr>
<tr>
<td></td>
<td>Any global based pending tasks available:
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></td>
<td>Processes running:
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>
</table>

	CPU utilization: Indicates the percentage of CPU used by this application.	Percent	A very high value could indicate 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	
	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 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.

2.4.3 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, select the **Enable/Disable** option from the **Tests** menu of the **Agents** tile, select **Component type** as *Microsoft RDS*, pick this test from the **DISABLED TESTS** list, click the **<** button, and click **Update** to save the changes.

Purpose	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
Target of the	Microsoft RDS

<table>
<tr><td>test</td><td colspan="3"></td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. 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
6. 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.
7. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every ‘published application’ on the Microsoft RDS server that is currently launched</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Launch count:
Represents the number of instances of this published application that have been launched currently.</td><td>Number</td><td>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.</td></tr>
</table>

	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.

2.5 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 2.7) facilitate such user-related analysis.

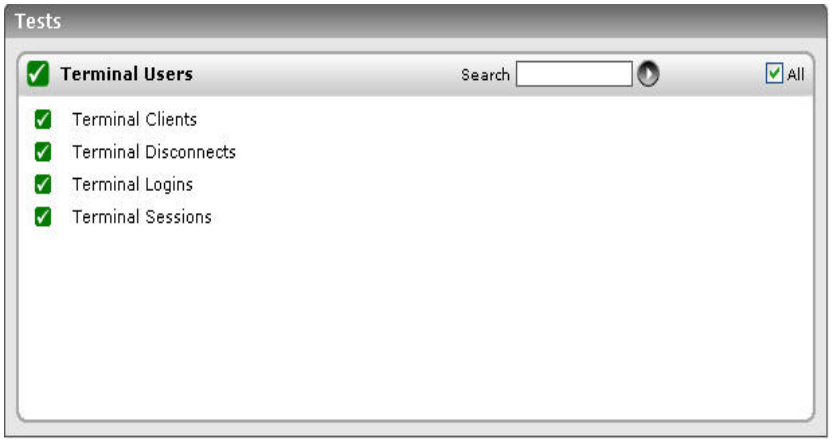

Figure 2.7: Tests associated with the Terminal Users layer

2.5.1 Terminal Sessions Test

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

Purpose	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

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. 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.
6. 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.
7. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Active sessions:
Indicates the number of active terminal services sessions currently on the server.</td><td>Number</td><td>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 the Microsoft RDS server.</td></tr>
</table>

	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 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.

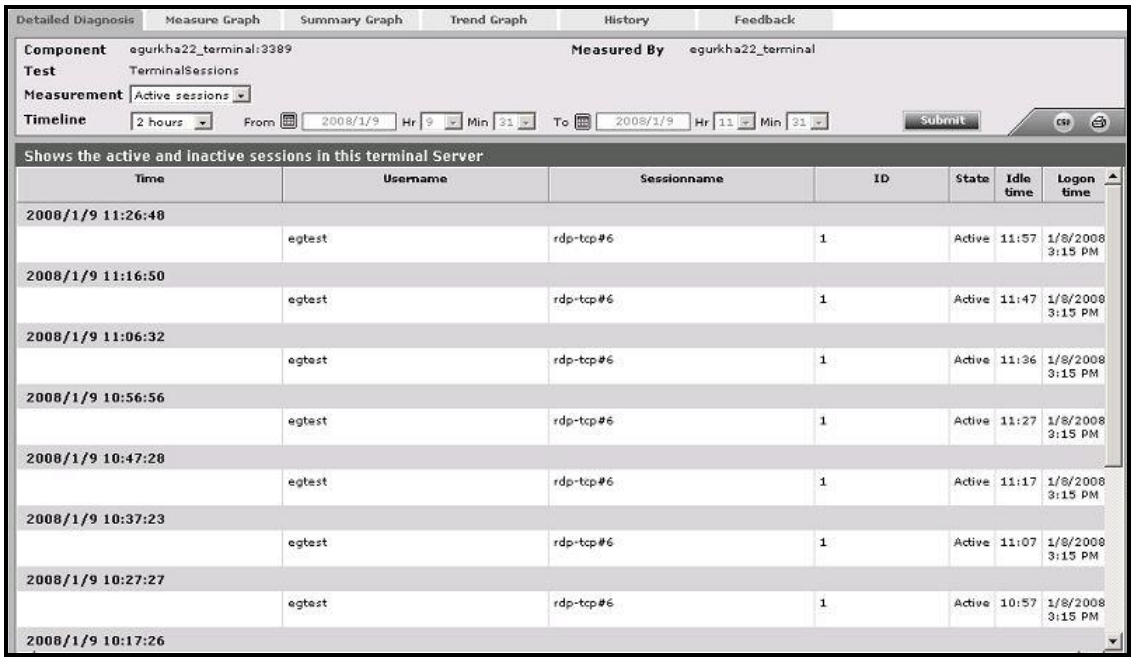

Figure 2.8: The detailed diagnosis of the Active sessions measure

2.5.2 Terminal Logins Test

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

Purpose	Monitors the new logins to the Microsoft RDS server
Target of the test	Any Microsoft RDS server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. 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.
6. 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.
7. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results is reported for each Microsoft RDS server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>New logins:
Indicates the number of new logins to the Microsoft RDS server in the last measurement period.</td><td>Number</td><td>A consistent zero value could indicate a connection issue.</td></tr>
</table>

	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.

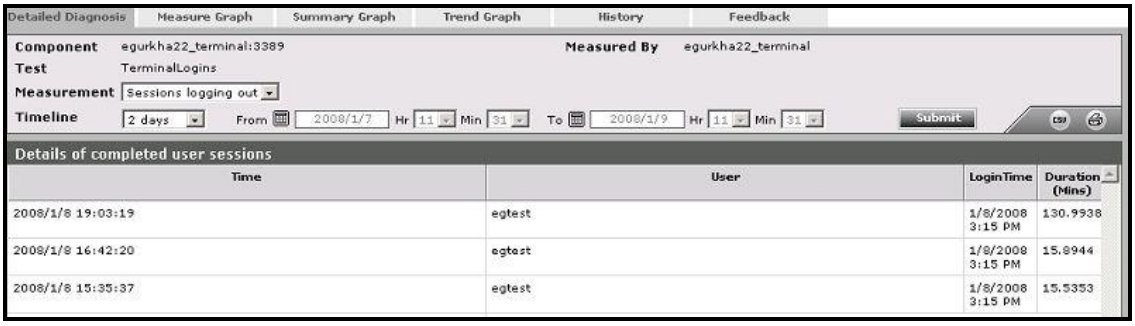

Figure 2.9: The detailed diagnosis of the Sessions logging out measure

2.5.3 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 *Microsoft Terminal* as the **Component type**, *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.

Purpose	To monitor the client connections to and from a Microsoft RDS server
Target of the test	A Microsoft RDS server
Agent deploying the test	Internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. SERVERIP - By default, the SERVERIP field will display the IP address of the Microsoft RDS server.
5. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every server being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Current connections:
The number of TCP connections currently established by clients to the Microsoft RDS server</td><td>Number</td><td>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.</td></tr>
<tr><td>New connections:
The number of new TCP connections initiated by clients to the Microsoft RDS server during the last measurement period</td><td>Number</td><td>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 server.</td></tr>
<tr><td></td><td>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</td><td>Number</td><td>A large number of sudden connection drops may be early warning indicators of problems with the Microsoft RDS server.</td></tr>
</table>

	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).

2.5.4 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 TerminalUsers 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.

This test will report metrics only if the Microsoft RDS server being monitored uses the .Net framework v3.0 (or above).

Purpose	Tracks every user connection from the Microsoft RDS client to the server, and monitors the resource utilization of every user on the Microsoft RDS server
Target of the test	A Microsoft RDS server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. 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.
6. 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.
7. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every user logged into the Microsoft RDS server</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>User sessions:
Represents the current number of sessions for a particular user</td><td>Number</td><td>A value of 0 indicates that the user is not currently connected to the Microsoft RDS server.</td></tr>
</table>

	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 reported is the sum of all memory utilizations across all the sessions.	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.
	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. 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.

	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 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.

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 2.10).

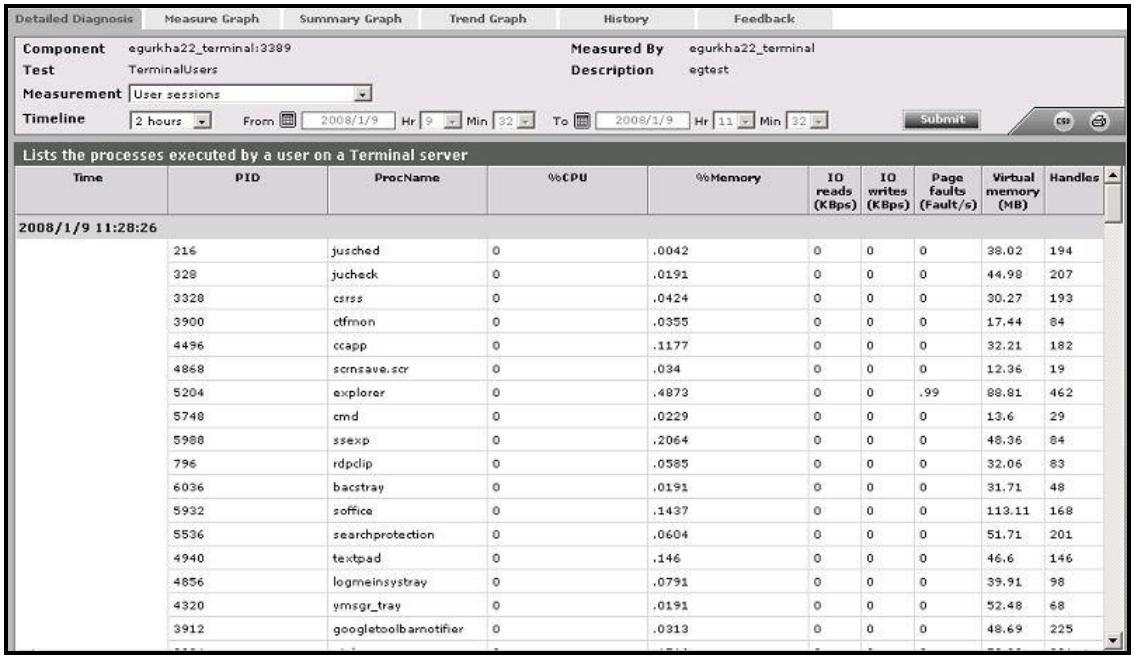

Time	PID	ProcName	%CPU	%Memory	IO reads (KBps)	IO writes (KBps)	Page faults (Fault/s)	Virtual memory (MB)	Handles
2008/1/9 11:28:26									
	216	jusched	0	.0042	0	0	0	38.02	194
	328	jucheck	0	.0191	0	0	0	44.98	207
	3328	csrss	0	.0424	0	0	0	30.27	193
	3900	ctfmon	0	.0355	0	0	0	17.44	84
	4496	ccapp	0	.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 2.10: 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**.

2.5.5 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.

Purpose	Measures the number of disconnected Microsoft RDS server sessions
Target of the test	Any Microsoft RDS server
Agent deploying the	An internal agent

<table>
<tr><td>test</td><td colspan="3"></td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.
5. 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 disconnected from the server recently. This way, administrators will be able to quickly determine which user belongs to which domain. If you want the detailed diagnosis to display the username alone, then set this flag to No.
6. 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.
7. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results is reported for each Microsoft RDS server being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total disconnected sessions:
Indicates the total number of sessions that are in the disconnected state.</td><td>Number</td><td></td></tr>
<tr><td>New disconnects:
Indicates the number of sessions that were disconnected in the last measurement period</td><td>Number</td><td>The detailed diagnosis of this measure, if enabled lists the users who have recently disconnected.</td></tr>
</table>

	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 2.11).

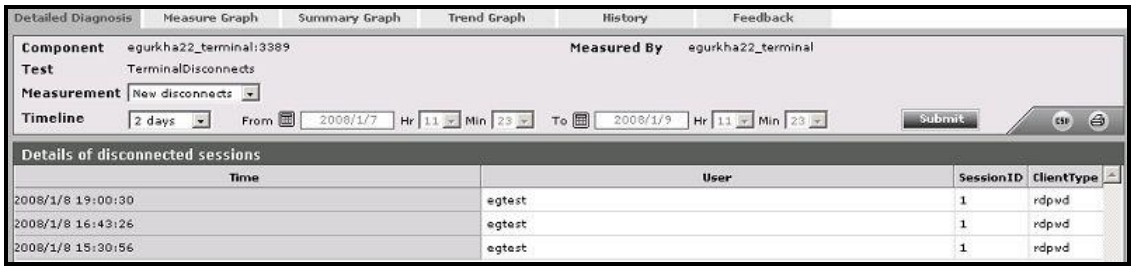

Figure 2.11: The detailed diagnosis of the New disconnects measure

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 2.12).

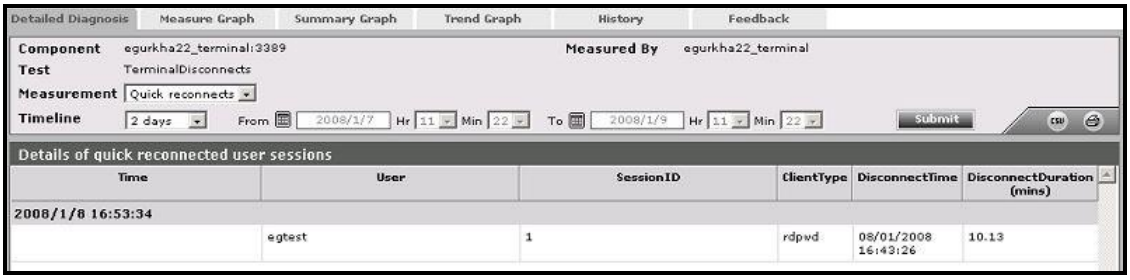

Figure 2.12: The detailed diagnosis of the Quick reconnects measure

2.5.6 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 *Microsoft Terminal* as the **Component type**, *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.

Purpose	Reports on the latencies seen by users connecting to a Microsoft RDS server		
Target	A Microsoft RDS server		
Agent deploying this test	Internal agent		
Configurable parameters for this test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured. 3. **PORT** - The port at which the **HOST** listens 4. **DISPLAYDOMAIN** - By default, the **DISPLAYDOMAIN** flag is set to **Yes**; this indicates that this test, by default, will report metrics for every *domainname*|*username* who is currently connected to the server. This way, administrators can quickly figure out which 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. 5. **PACKETSIZE** - The size of packets used for the test (in bytes) 6. **PACKETCOUNT** – The number of packets exchanged between the Microsoft RDS server and the user terminal during the test 7. **TIMEOUT** - How long after transmission should a packet be deemed lost (in seconds) 8. **PACKETINTERVAL** - Represents the interval (in milliseconds) between successive packet transmissions during the execution of this test. 9. **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.		
Outputs of the test	One set of outputs for every user currently connected to the Microsoft RDS server		
Measurements of the test	**Measurement**	**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 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	A significant increase in the 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 ehnance 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.

2.5.7 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.

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

Purpose	For each remote user connecting to a RemoteFX-enabled Microsoft RDS server, this test measures user experience and reports abnormalities (if any)
Target of the test	A Microsoft RDS server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Microsoft RDS server
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every user logged into the Microsoft RDS server</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>User sessions:
Represents the current number of sessions for a particular user.</td><td>Number</td><td>A value of 0 indicates that the user is not currently connected to the Microsoft RDS server.</td></tr>
</table>

	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 input for this user.	Percent	The compression ratio typically affects the quality of the picture. Generally, the higher the compression ratio, the poorer the quality of the resulting picture. Ideally therefore, the value of this measure should be 0. You can compare the value of this measure across users to identify that user whose picture output was very poor owing to high compression.

	Input frames: Indicates the number of source frames provided per second as input to the RemoteFx graphics for this user.	Frames/Sec	
	Output Frames: Indicates the number of source frames sent per second to this user as output of RemoteFx graphics.	Frames/Sec	
	Source frames: Indicates number of frames per second composed at the source for this user.	Frames/Sec	
	Base TCP round trip time: Indicates the time between initiating a network request and receiving a response over TCP for this user.	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 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 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, open the **ENABLE / DISABLE TESTS** page using the Agents -> Tests -> Enable/Disable menu sequence, select **Microsoft Terminal** as the component-type, **Performance** as the *Test type*, select the test from the **DISABLED TESTS** list, and click on **>>** to move it to the **ENABLED TESTS** list. Finally, click on the **Update** button. This test is mapped to the **Windows Service** layer of the Microsoft RDS server component.

Chapter 3

Monitoring Active Directory Servers

A directory service consists of both a directory storage system called the "directory store" and a mechanism that is used to locate and retrieve information from the system. The primary functions of the directory service are managed by the Directory System Agent (DSA), which is a process that runs on each domain controller (abbreviated as DC). Active Directory is the directory service that is included with Microsoft Windows. It stores objects that provide information about the real entities that exist in an organization's network like printers, applications, databases, users etc. Active Directory is a part of the domain controller. It is associated with one or more domains. It stores information about users, specific groups of users like the Administrator, computers, applications, services, files, and distribution lists etc. Active Directory then makes this information available to the users and applications throughout the organization.

Active Directory is an important component of the Windows environment. Like any other Windows applications, its performance can affect the rest of the target environment. Active Directory consumes resources and the administrator needs to be aware of how much of the system's resources are being consumed over a long term. This helps the administrators to plan for future upgrades. Gathering performance data gives the administrators a good way to see the effects of any optimization efforts that he/she might attempt, and provides a great way for diagnosing problems when they occur. Most of the Windows servers and components are dependent on Active Directory either directly or indirectly. So monitoring the Active Directory server's performance regularly is necessary to make sure that the target environment is meeting your business and networking goals.

The eG Enterprise suite provides extensive monitoring support to the Active Directory (AD) server operating on Windows 2000, 2003, and 2008/2012. The specialized monitoring model that the eG Enterprise offers (see Figure Figure 3.1) periodically executes a number of tests on the AD server to extract a wide gamut of metrics indicating the availability, responsiveness, and overall health of the AD server and its underlying operating system. Using this model, Active Directory servers can be monitored in an agent-based or an agentless manner.

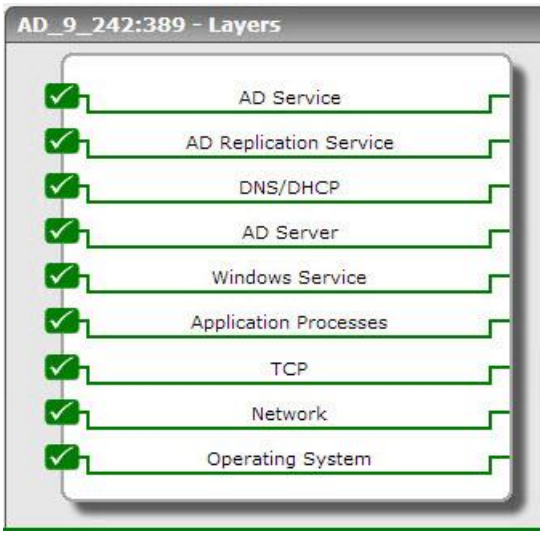

Figure 3.1: Layer model for Active Directory

Using these metrics, an administrator can find quick answers to the following performance queries:

- Is the AD server available?
- How quickly is the server responding to user requests?
- Are there adequate work items to service blocking requests, or are too many requests getting rejected?
- Have any internal server errors been reported recently?
- Have too many login attempts failed?
- Did session timeouts occur too frequently?
- Is the schema cache effectively utilized, or is disk read/write activity high?
- Is the server currently overloaded? Are sufficient domain controllers available in the environment to handle the load?
- Are all changes to the AD server getting replicated across and within sites?
- How many directory synchronizations are in queue? Is the number high enough to force a replication?

The last 5 layers of Figure 3.1 have been discussed in the *Monitoring Unix and Windows Servers* document, and will hence not be discussed again. However, for the *Active Directory* server alone, the **Operating System** layer is mapped to an additional **Net Logon** test. The section that follows will discuss this test in detail. All other sections in this chapter will focus only on the top 3 layers of Figure 3.1.

3.1 The Operating System Layer

The **Operating System** layer of a monitored *Active Directory* server typically runs all the tests that are mapped to the same layer for a *Windows* server or a *Windows Generic* server. The only difference however is that for the Active Directory server, an additional **Net Logon** test is mapped to this layer. This section provides details of the **Net Logon** test.

3.1.1 Net Logon Test

The Netlogon service is responsible for communication between systems in response to a logon request, a domain synchronization request, and a request to promote a Backup Domain Controller (BDC) to a Primary Domain Controller (PDC). The Netlogon service performs several tasks when servicing network logon requests. They are as follows:

- Selects the target domain for logon authentication
- Identifies a domain controller in the target domain to perform authentication
- Creates a secure channel for communication between Netlogon services on the originating and target systems
- Passes an authentication request to the appropriate domain controller
- Returns authentication results to Netlogon on the originating system

Delays in the Netlogon authentication process can often scar a user's overall experience with not just the domain controller, but also with the application that requests for the authentication. In order to avoid undue authentication delays, you can use the **Net Logon** test. This test monitors the Netlogon authentication feature, proactively detects potential authentication bottlenecks, and promptly alerts administrators to what is causing the bottleneck, so that remedial actions can be initiated in good time.

Purpose	Monitors the Netlogon authentication feature, proactively detects potential authentication bottlenecks, and promptly alerts administrators to what is causing the bottleneck, so that remedial actions can be initiated in good time		
Target of the test	An Active Directory server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every AD server being monitored		
Measurements	**Measurement**	**Measurement Unit**	**Interpretation**

made by the test	**Semaphore waiters:** Indicates the number of threads currently waiting to acquire the semaphore.	Number	A consistent increase in the value of this measure is a cause forconcern, as it indicates that the count of 'busy' semaphores is steadily increasing. This in turn could cause many threads/logon requests to be enqueued, due to the lack of adequate semaphores. Consequently, authentication will be delayed.
	Semaphore acquires: Indicates the number of times the semaphore has been acquired over this secure channel during the last measure period.	Number	
	Semaphore holders: Indicates the number of threads currently holding the semaphore.	Number	This is a good indicator of the current authentication workload over the secure channel. If the value of this measure is equal to the *MaxConcurrentApi* registry setting or is fast approaching that value, it indicates that the server is getting overloaded. Authentication delays and timeouts may occur as a result. The typical way to resolve the problem is to raise the maximum allowed worker threads that service that authentication. You can do this by altering the *MaxConcurrentApi* registry value and then restarting the Net Logon service on the servers.
	Semaphore timeouts: Indicates the number of times a thread has timed out waiting for the semaphore over the secure communication channel during the last measure period.	Number	Ideally, this measure has to be 0. A non-zero value for the measure indicates that one/more authentication threads have hit the time-out for the waiting and the logon was denied. This is a sign of a very bad user experience, and typically occurs when the secure channel is overloaded, hung or broken. The typical way to resolve the *overload* problem is to raise the maximum allowed worker threads that service that authentication. You can do this by altering the *MaxConcurrentApi* registry value and then restarting the Net Logon service on the servers.

3.2 The AD Server Layer

The **AD Server** layer verifies the availability and responsiveness of the Active Directory (AD) service from an external location. This layer also monitors the user accesses to the AD server and reports how well the server handles access requests. In the process, the layer also reports useful session-related metrics pertaining to the user sessions on the AD server. Besides, the layer also reports the overall health of the AD database (see Figure 3.2).

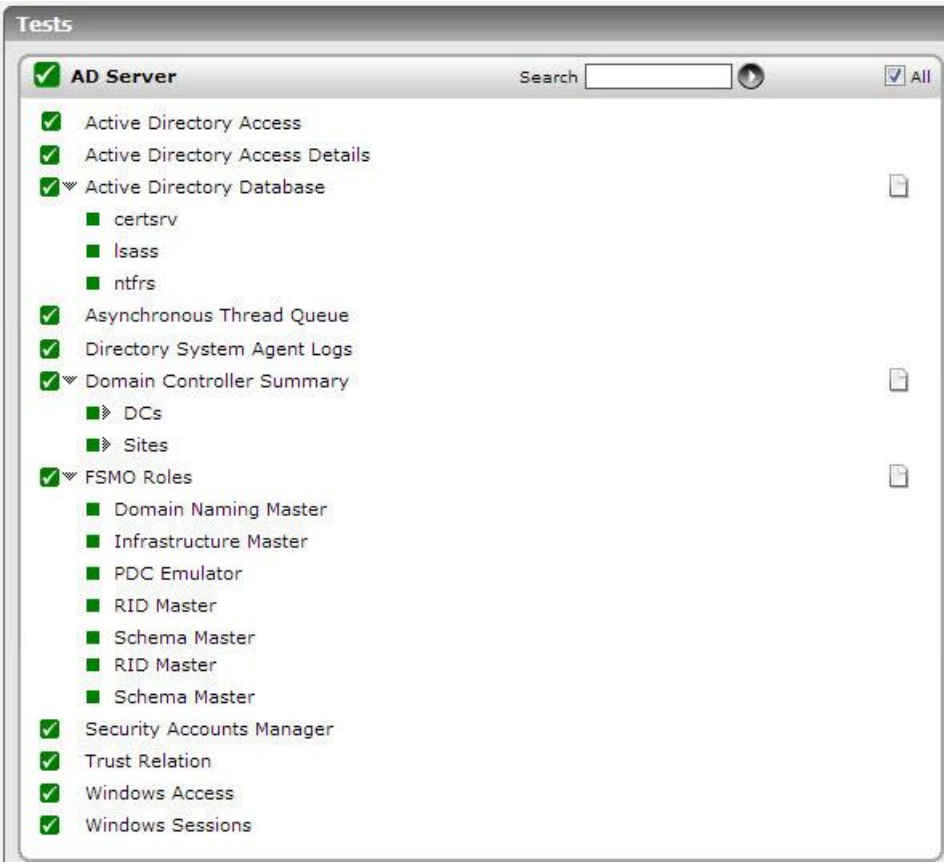

Figure 3.2: The tests associated with the AD Server layer

3.2.1 Asynchronous Thread Queue Test

Monitoring the asynchronous thread queue (ATQ) on an AD server will provide useful pointers to the request processing ability of the server. This test monitors the ATQ, reports the number and nature of requests queued in the ATQ, captures a steady growth (if any) in the length of the queue over time, and thus reveals potential processing bottlenecks on the AD server.

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	Monitors the ATQ, reports the number and nature of requests queued in the ATQ, captures a

<table>
<tr><td></td><td colspan="3">steady growth (if any) in the length of the queue over time, and thus reveals potential processing bottlenecks on the AD server</td></tr>
<tr><td>Target of the test</td><td colspan="3">An Active Directory or Domain Controller on Windows 2008</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The IP address of the machine where the Active Directory is installed.
3. PORT – The port number through which the Active Directory communicates. The default port number is 389.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Active Directory being monitored</td></tr>
<tr><td rowspan="4">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>ATQ estimated queue delay:
Indicates the estimated time the next request will spend in the queue prior to being serviced by the directory service.</td><td>Secs</td><td></td></tr>
<tr><td>ATQ outstanding queued requests:
Indicate how many requests are queued at the domain controller.</td><td>Number</td><td>A high level of queuing indicates that requests are arriving at the domain controller faster than they can be processed. This can also lead to a high latency in responding to requests. Delay is the estimated time the next request will spend in the queue prior to being serviced by the directory service, 1.265 seconds.</td></tr>
<tr><td>ATQ request latency:
Indicates the average length of time to process a request, not including time spent on the queue.</td><td>Secs</td><td>A high value of this measure is a cause for concern, as it indicates a processing bottleneck on the AD server.</td></tr>
<tr><td></td><td>ATQ threads ldap:
Indicates the number of threads that ATQ has currently allocated to servicing LDAP requests.</td><td>Number</td><td></td></tr>
</table>

	ATQ thread others: Indicates the number of threads that ATQ has currently allocate to DS services other than LDAP.	Number	
	ATQ threads total: Indicates the total number of threads that are either waiting to service an incoming request or are already servicing a request.	Number	If values for this counter and **ATQ Threads ldap** counter are equal, a queue is likely building on the LDAP port, which will result in long response times. If the two counters are always equal, use Server Performance Advisor to troubleshoot the problem.

3.2.2 ADAM Access Details Test

This test measures the load on the AD server in terms of the level of read-write activity on the server and the count of search operations performed by the server. In the process, the test reveals the following:

- Which AD services initiated the read-write operations? Which of these services generated the maximum I/O load on the server - is it the LSA? the NSPI? the NTDS? SAM? or the replication service? - this information is useful when administrators are faced with an AD overload, as it accurately points them to the probable sources of the load;
- Which AD service performed the maximum searches on the server? - in the event of an overload, this metric will help you identify that service which could be contributing to the overload;
- Is the server sized with adequate threads to handle the I/O load?

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	Measures the load on the AD server in terms of the level of read-write activity on the server and the count of search operations performed by the server. In the process, the test reveals the following: • Which AD services initiated the read-write operations? Which of these services generated the maximum I/O load on the server - is it the LSA? the NSPI? the NTDS? SAM? or the replication service? • Which AD service performed the maximum searches on the AD server? • Is the server sized with adequate threads to handle the I/O load?
Target of the test	An Active Directory or Domain Controller on Windows 2008
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Schema cache hit ratio: Indicates the percentage of object name lookups serviced by the Schema Cache.	Percent	All changes made to Active Directory are validated first against the schema. For performance reasons, this validation takes place against a version of the schema that is held in memory on the domain controllers. This "in-memory version," called the schema cache, is updated automatically after the on-disk version has been updated. The schema cache provides mapping between attribute identifiers such as a database column identifier or a MAPI identifier and the in-memory structures that describe those attributes. The schema cache also provides lookups for class identifiers to get in-memory structures describing those classes. A low value of this measure indicates that the Directory Service needs high disk read/write activity to perform its job. This results in poor response time of the components available in the Active Directory.
	Notify queue size: Indicates the number of pending update notification requests that have been queued and not transmitted.	Number	When any change in the Active Directory occurs, the originating domain controller sends an update notification requests to the other domain controllers. A high value of this measure indicates that the Active Directory is changing frequently but the update notification requests have not been transmitted to the other domain controllers. This results in a loss of data integrity in the directory store. This problem can be corrected by forcing the replication.

	Current threads in use: Indicates the current number of threads in use by the directory service (which is different from the number of threads in the directory service process).	Number	This is the number of threads currently servicing client API calls; it can be used to indicate whether additional processors should be used. A fluctuating value for this measure indicates a change in the load. A low value could point to network problems that are preventing client requests from succeeding.
	Server binds: Indicates the number of domain controller–to–domain controller binds per second that are serviced by this domain controller.	Binds/Sec	
	Directory reads: Indicates the rate of directory reads.	Reads/Sec	These measures serve as effective indicators of the ability of the AD server to process read, write, and search requests.
	Directory writes: Indicates the rate of directory writes.	Writes/Sec	
	Directory searches: Indicates the number of directory searches per second.	Searches/Sec	
	DS reads from DRA: Indicates the percentage of reads on the directory by replication.	Percent	If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the *DS reads from KCC*, *DS reads from LSA*, *DS reads from NSPI*, *DS reads from NTDS*, and *DS reads from SAM* measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?

	DS reads from KCC: Indicates the percentage of reads performed by the Knowledge Consistency Checker on the directory.	Percent	The **Knowledge Consistency Checker (KCC)** generates the replication topology by specifying what domain controllers will replicate to which other domain controllers in the site. The KCC maintains a list of connections, called a **replication topology**, to other domain controllers in the site. The KCC ensures that changes to any object are replicated to all site domain controllers and updates go through no more than three connections. If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the *DS reads from DRA*, *DS reads from LSA*, *DS reads from NSPI*, *DS reads from NTDS*, and *DS reads from SAM* measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS reads from LSA: Indicates the percentage of reads performed by the Local Security Authority on the directory.	Percent	The Local Security Authority (LSA) is the security subsystem responsible for all interactive user authentication and authorization services on a local computer. The LSA is also used to process authentication requests made through the Kerberos V5 protocol or NTLM protocol in Active Directory. If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the *DS reads from DRA*, *DS reads from KCC*, *DS reads from NSPI*, *DS reads from NTDS*, and *DS reads from SAM* measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS?

	DS reads from NSPI: Indicates the percentage of reads performed by the Name Service Provider Interface (NSPI) on the directory.	Percent	The **Name Service Provider Interface (NSPI)** is the protocol by which Messaging API (MAPI) clients access the AD DS. Exchange Address Book clients use the client MAPI provider Emsabp32.dll to look up e-mail addresses in the global catalog. The client-side MAPI provider communicates with the server through the proprietary Name Service Provider Interface (NSPI) RPC interface. If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the *DS reads from KCC*, *DS reads from LSA*, *DS reads from DRA*, *DS reads from NTDS*, and *DS reads from SAM* measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? or the NSPI?
	DS reads from NTDS: Indicates the percentage of reads performed by the name service directory APIs on the directory.	Percent	If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the *DS reads from KCC*, *DS reads from LSA*, and *DS reads from DRA*, *DS reads from NSPI*, and *DS reads from SAM* measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? or the SAM?
	DS reads from SAM: Indicates the percentage of reads performed by the Security Account Manager (SAM) on the directory.	Percent	The Security Accounts Manager (SAM) is used for verifying passwords and for checking passwords against any existing password policies that are in effect on a domain controller. If the AD server is experiencing abnormally high read activity, then, you can compare the value of this measure with the values reported by the *DS reads from KCC*, *DS reads from LSA*, and *DS reads from DRA*, *DS reads from NSPI*, and *DS reads from NTDS* measures to know which AD service is performing the maximum reads on the AD server - is it the replication service? the LSA? the KCC? the NSPI? or the NTDS?

	DS writes from DRA: Indicates the percentage of writes on the AD server by replication.	Percent	If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the *DS writes from KCC*, *DS writes from LSA*, *DS writes from NSPI*, *DS writes from NTDS*, and *DS writes from SAM* measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS writes from KCC: Indicates the percentage of writes performed by the Knowledge Consistency Checker on the directory.	Percent	If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the *DS writes from DRA*, *DS writes from LSA*, *DS writes from NSPI*, *DS writes from NTDS*, and *DS writes from SAM* measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the KCC? the LSA? the NSPI? the NTDS? or the SAM?
	DS writes from LSA: Indicates the percentage of writes performed by the Local Security Authority on the directory.	Percent	If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the *DS writes from DRA*, DS *writes from KCC*, *DS writes from NSPI*, *DS writes from NTDS*, and *DS writes from SAM* measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS writes from NSPI: Indicates the percentage of writes performed by the Name Service Provider Interface (NSPI) on the directory.	Percent	If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the *DS writes from DRA*, *DS writes from KCC*, *DS writes from LSA*, *DS writes from NTDS*, and *DS writes from SAM* measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS writes from NTDS: Indicates the percentage of writes performed by the name service directory APIs on the directory.	Percent	If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the *DS writes from DRA*, *DS writes from KCC*, *DS writes from LSA*, *DS writes from NSPI*, and *DS writes from SAM* measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?

	DS writes from SAM: Indicates the percentage of writes performed by the Security Accounts Manager (SAM) on the directory.	Percent	If the AD server is experiencing abnormally high write activity, then, you can compare the value of this measure with the values reported by the *DS writes from DRA*, *DS writes from KCC*, *DS writes from LSA*, *DS writes from NSPI*, and *DS writes from NTDS* measures to know which AD service is performing the maximum writes on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS searches from DRA: Indicates the percentage of searches performed by the replication service on the AD server.	Percent	If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the *DS searches from KCC*, *DS searches from LSA*, *DS searches from NSPI*, *DS searches from NTDS*, and *DS searches from SAM* measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS searches from KCC: Indicates the percentage of searches performed by the Knowledge Consistency Checker on the directory.	Percent	If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the *DS searches from DRA*, *DS searches from LSA*, *DS searches from NSPI*, *DS searches from NTDS*, and *DS searches from SAM* measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS searches from LSA: Indicates the percentage of searches performed by the Local Security Authority on the directory.	Percent	If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the *DS searches from DRA*, *DS searches from KCC*, *DS searches from NSPI*, *DS searches from NTDS*, and *DS searches from SAM* measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?

	DS searches from NSPI: Indicates the percentage of searches performed by the Name Service Provider Interface (NSPI) on the directory.	Percent	If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the *DS searches from DRA*, *DS searches from KCC*, *DS searches from LSA*, *DS searches from NTDS*, and *DS searches from SAM* measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS searches from NTDS: Indicates the percentage of searches performed by the name service directory APIs on the directory.	Percent	If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the *DS searches from DRA*, *DS searches from KCC*, *DS searches from LSA*, *DS searches from NSPI*, and *DS searches from SAM* measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?
	DS searches from SAM: Indicates the percentage of searches performed by the Security Accounts Manager (SAM) on the directory.	Percent	If the AD server is processing an abnormally large number of search requests, then, you can compare the value of this measure with the values reported by the *DS searches from DSA*, DS *searches from KCC*, *DS searches from LSA*, *DS searches from NSPI*, and *DS searches from NTDS* measures to know which AD service is performing the maximum number of searches on the AD server - is it the replication service? the LSA? the KCC? the NSPI? the NTDS? or the SAM?

3.2.3 ADAM Database Test

This test reports critical statistics pertaining to the usage of the database caches, and the overall health of the AD database.

Purpose	Reports critical statistics pertaining to the usage of the database caches, and the overall health of the AD database
Target of the test	An Active Directory server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every AD server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Database cache hits : Indicates the percentage of page requests of the database file that were occupied in a cache before responding to the request.	Percent	Ideally, the value of this measure should be moderate. A high value of this measure indicates the high utilization of physical memory. In such a case, you can add the required memory to the database.
	Database table cache hits: Indicates the percentage of database tables that were opened using cached schema information.	Percent	Ideally, the value of this measure should be high.
	Log records waiting: Indicates the rate of log record stalls, per second.	Records/Sec	
	Log threads waiting: Indicates the current number of threads waiting for data to be written to the log so that database updation will be executed.	Number	

3.2.4 Active Directory Access Test

This test monitors the availability and response time from clients of an Active Directory server from an external perspective.

Purpose	Monitors the availability and response time from clients of an Active Directory server from an internal perspective
Target of the test	An Active Directory or Domain Controller
Agent deploying the test	An external agent

Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389. 4. **DOMAIN** - The default value of the **DOMAIN** parameter will be *none*. In Windows 2003 environments however, the ADServerTest will function effectively only if a "fully qualified domain name" is provided in the **DOMAIN** text box. 5. **USER** - Provide the name of a domain user in the **USER** text box. This can be *none* for Windows 2000 environments. 6. **PASSWORD** - Provide the password for the domain user specified above, in the **PASSWORD** text box. This can be *none* for Windows 2000 environments. 7. **CONFIRM PASSWORD** – Confirm the **PASSWORD** by retyping it here. 8. **CONNECTTIMEOUT** - By default, this is set to *30* seconds. This implies that by default, the test will wait for 30 seconds to establish a connecton with the target Active Directory server. If a connection is established within the default 30 second period, then the test will report that the server is available; if the test is unable to connect to the server within the default period, then it will report that the server is unavailable. If it generally takes a longer time for clients to connect to the AD server in your environment, then, you may want to change the **CONNECTTIMEOUT** period so that, the test does not time out before the connection is established, and consequently present an "untrue" picture of the availability of the server.
Outputs of the test	One set of results for every Active Directory being monitored

Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active directory availability: Indicates the availability of the server.	Percent	The availability is 100% when the server is responding to a request and 0% when it is not. Availability problems may be caused by a misconfiguration / malfunctioning of the server, or if the server has not been started.
	Active directory response time: Indicates the time taken by the server to respond to a user query	Secs	A sudden increase in response time is indicative of a bottleneck at the server.

3.2.5 Windows Access Test

This test monitors the accesses to an AD server.

Purpose	Monitors the accesses to the Windows server
Target of the test	An Active Directory server or a Domain Controller
Agent	An internal agent

deploying the test			
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every AD server or domain controller being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Blocking request rejects: The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items	Reqs/sec	If the number of blocking request rejects is high, you may need to adjust the MaxWorkItem or MinFreeWorkItems server parameters
	Permission errors: The number of times opens on behalf of clients have failed with STATUS_ACCESS_DENIED in the last measurement period	Number	Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected.
	File access denied errors: The number of times accesses to files opened successfully were denied in the last measurement period	Number	This number indicates attempts to access files without proper access authorization.
	Internal server errors: This value indicates the number of times an internal server error was detected in the last measurement period.	Number	Unexpected errors usually indicate a problem with the server.
	Data received: The rate at which the server has received data from the network	Kbytes/sec	This metric indicates how busy the server is.
	Data transmitted: The rate at which the server has sent data over the network	Kbytes/sec	This metric indicates how busy the server is.
	Resource shortage errors: The number of times STATUS_DATA_NOT_ACCEPTED was returned to clients in the last measurement period	Number	A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the InitWorkItems or MaxWorkItems server parameters might need to be adjusted.

	Avg response time: Average time taken by the server to respond to client requests	Secs	This is a critical measure of server health.

3.2.6 Windows Sessions Test

This test reports various session-related statistics for an AD server.

Purpose	Reports various session-related statistics for a Windows server		
Target of the test	An AD server or a Windows Domain Controller		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every AD server or domain controller being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Logons: Rate of logons to the server	Reqs/sec	This measure reports the rate of all interactive, network, and service logons to a windows server. The measure includes both successful and failed logons.
	Logon errors: Number of logons in the last measurement period that had errors	Number	This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server.
	Current sessions: The number of sessions currently active in a server	Number	This measure is one of the indicators of current server activity.
	Sessions with errors: The number of sessions in the last measurement period that were closed to unexpected error conditions	Number	Sessions can be closed with errors if the session duration reaches the autodisconnect timeout.

	Sessions forced off: The number of sessions in the last measurement period that have been forced to logoff	Number	This value indicates how many sessions were forced to logoff due to logon time constraints.
	Sessions logged off: The number of sessions in the last measurement period that were terminated normally	Number	Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low.
	Sessions timed out: The number of sessions that have been closed in the last measurement period due to their idle time exceeding the AutoDisconnect parameter for the server	Number	The number of session timed out gives an indication of whether the AutoDisconnect setting is helping to conserve server resources

3.2.7 FSMO Roles Test

FSMO stands for Flexible Single Master Operations, and FSMO roles (also known as operations master roles) help you prevent conflicts in your Active Directory.

For most Active Directory objects, the task of updating can be performed by any Domain Controller except those Domain Controllers that are read-only. Updates such as computer object properties, renamed organizational units, and user account password resets can be handled by any writable domain controller.

After an object is changed on one domain controller, those changes are propagated to the other domain controllers through replication. During replication all of the Domain Controllers share their updates. So a user that has their password reset in one part of the domain may have to wait until those changes are replicated to the Domain Controller that they are signing in from.

This model works very well for most objects. In the case of any conflicts, such as a user's password being reset by both the central helpdesk as well as an administrator working at the user's site, then conflicts are resolved by whichever made the last change. However, there are some changes that are too important, and are not well suited to this model.

There are 5 specific types of updates to Active Directory that are very specific, and conflicts should be avoided. To help alleviate any potential conflicts, those updates are all performed on a single Domain Controller. And though each type of update must be performed on a single Domain Controller, they do not all have to be handled by the same Domain Controller.

These types of updates are handled by Domain Controllers Flexible Single Master Operations roles, or FSMO roles. Each of the five roles is assigned to only one domain controller.

There are five FSMO roles in every Active Directory forest. They are:

- Schema Master
- Domain Naming Master

- Infrastructure Master
- Relative ID (RID) Master
- Primary Domain Controller (PDC) Emulator

Among these five FSMO roles, the following three FSMO roles are needed only once in every domain in the forest:

- Infrastructure Master
- Relative ID (RID) Master
- Primary Domain Controller (PDC) Emulator

If a domain controller configured with a specific FSMO role is suddenly rendered unavailable or is unreachable, then that particular function cannot be performed. This in turn implies that the types of updates that will otherwise be handled by that domain controller can no longer be processed, thus creating a climate of conflict in the AD environment. With the held of the **FSMO Roles** test however, you can rapidly detect the unavailability of an FSMO domain controller over the network, isolate potential network connectivity issues and latencies, and spot real/probable delays in LDAP binding, so that such issues can be promptly remedied and conflicts prevented.

Purpose	Helps rapidly detect the unavailability of an FSMO domain controller over the network, isolate potential network connectivity issues and latencies, and spot real/probable delays in LDAP binding, so that such issues can be promptly remedied and conflicts prevented.		
Target of the test	An AD server or a Windows Domain Controller		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for each FSMO role		
Measurements	**Measurement**	**Measurement Unit**	**Interpretation**

made by the test	**LDAP bind time:** Indicates the time taken for the last successful LDAP bind.	Secs	In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as binding. When a programmatic object, such as an IADs (Interface Adapter Device) or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be bound to the directory object. The method for programmatically binding to an Active Directory object will depend on the programming technology that is used. All bind functions and methods require a binding string. The form of the binding string depends on the provider. Active Directory Domain Services are supported by two providers, *LDAP* and *WinNT*. Beginning with Windows 2000, the LDAP provider is used to access Active Directory Domain Services. The LDAP binding string can take one of the following forms: "LDAP://<host name>/<object name>" "GC://<host name>/<object name>" Ideally, the value of this measure should be low. A high value for this measure could be a possible indication of network-related problems or of the hardware that needs to be upgraded immediately. This measure will not be reported if the value of the *Availability* measure is *0*.
	Avg network delay: Indicates the average delay between transmission of packet to a target and receipt of the response to the packet at the source.	Secs	An increase in network latency could result from misconfiguration of the router(s) along the path, network congestion, retransmissions at the network, etc. The detailed diagnosis capability, if enabled, lists the hop-by-hop connectivity and delay. This measure will not be reported if the value of the *Availability* measure is *0*.

	Minimum network delay: Indicates the minimum time between transmission of a packet and receipt of the response back.	Secs	A significant increase in the minimum round-trip time is often a sure sign of network congestion. This measure will not be reported if the value of the *Availability* measure is *0*.
	Packet loss: Indicates the percentage of packets lost during transmission from source to target and back.	Percent	Packet loss is often caused by network buffer overflows at a network router or by packet corruptions over the network. The detailed diagnosis for this measure provides a listing of routers that are on the path from the external agent to target server, and the delays on each hop. This information can be used to diagnose the hop(s) that could be causing excessive packet loss/delays. This measure will not be reported if the value of the *Availability* measure is *0*.
	Availability: Indicates whether/not this FSMO role is available over the network.	Percent	A value of *100* indicates that the FSMO role is available. The value *0* indicates that the FSMO role is not available. Typically, the value 100 corresponds to a *Pkt_loss_pct* of 0. If the FSMO role is not available over the network i.e., if this measure reports a value 0, all other measures applicable for this test will not be reported.

3.2.8 Directory System Agent Logs Test

This test monitors the Active Directory database files and log files for file size, and also monitors free disk space on the hosting volumes.

Purpose	Monitors the Active Directory database files and log files for file size, and also monitors free disk space on the hosting volumes
Target of the test	An AD server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every AD server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Directory system agent DB size: Indicates the size of the database files on the AD server.	MB	
	System volume size: Indicates the size of the SYSVOL folder - SYSVOL is the shared directory on domain controllers that contains Group Policy and logon script information.	MB	
	Directory system agent log file size: Indicates the size of the log files on the AD server.	MB	
	Directory system agent free log space: Indicates the amount of free space on the volume hosting log files.	MB	Ideally, this value should be high.
	Directory system agent free DB space: Indicates the amount of free space on the volume hosting database files.	MB	Ideally, this value should be high. If the free space for database files is very low, then the AD server might be rendered unable to update objects.
	System volume share availability: Indicates whether the SYSVOL folder is available or not.	Percent	If the value of this measure is 100, it indicates the SYSVOL folder is available. The value 0 on the other hand, indicates that the folder is not available.

3.2.9 Domain Controller Summary

Use this test to know the number and names of all domain controllers that manage the servers and users in the domains of interest to you.

This test runs only on Active Directory servers operating on Windows 2008.

Purpose	Reports the number and names of all domain controllers that manage the servers and users in the domains of interest to you		
Target of the test	An AD server on Windows 2008		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the AD server 4. **DNS NAME** - Provide a comma-separated list of the fully qualified domain names of all the domains that you want the test to scan for domain controllers. For instance, your specification can be, *chn.eginnovations.com,mas.eginnovations.com*.		
Outputs of the test	One set of results for every domain name configured against **DNS NAME**		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Domain Controllers: Indicates the number of domain controllers in this domain.	Number	The detailed diagnosis of this measure lists the names of all domain controllers in a chosen domain.

3.2.10 Security Accounts Manager Test

Every Windows computer has a local Security Accounts Manager (SAM). The SAM is responsible for a few functions. First, it is responsible for storing the local users and groups for that computer. Second, the local SAM is responsible for authenticating logons. When a computer is not joined to a domain, the only option is to use the local SAM to perform the authentication.

If too many computer/user creations in SAM fail or if SAM takes too long to enumerate, evaluate, and authenticate users/user groups, the user experience with the computer is bound to be impacted adversely. By periodically monitoring the operations of SAM, administrators can proactively detect potential problem conditions and plug the holes, so that the user experience remains unaffected. The **Security Accounts Manager** test does just that. At configured intervals, this test checks how well SAM performs its core functions, and promptly reports real/probable failures and latencies to the administrator.

This test applies only to Active Directory Servers installed on Windows 2008 and above.

Purpose	At configured intervals, this test checks how well SAM performs its core functions, and promptly reports real/probable failures and latencies to the administrator.
Target of the test	An Active Directory or Domain Controller on Windows 2008 and above
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Machine creation attempts: Indicates the number of attempts per second to create computer accounts.	Number	
	User creation attempts: Indicates the number of attempts per second to create user accounts.	Number	
	Successful user creations: Indicates the number of user accounts successfully created per second.	Number	Ideally, the value of this measure should be equal to the value of the *User creation attempts* measure. A low value is a cause for concern, as it indicates that many user creation attempts are failing; the reasons for the same have to be ascertained and addressed soon.
	Successful computer creations: Indicates the number of computers successfully created per second.	Number	Ideally, the value of this measure should be equal to the value of the *Machine creation attempts* measure. A low value is a cause for concern, as it indicates that many machine creation attempts are failing; the reasons for the same have to be ascertained and addressed soon.
	GC evaluations: Indicates the number of SAM global catalog evaluations per second.	Number	
	Enumerations: Indicates the number of net user, net group, and net local function enumerations per second.	Connections/Sec	

	Display information queries: Indicates the number of queries per second to obtain display information.	Connections/Sec	
	Account group evaluation latency: Indicates the time taken by SAM to evaluate an account group.	Secs	This indicates the mean latency of the last 100 account and universal group evaluations performed for authentication. A high value could indicate a bottleneck.
	Resource group evaluation latency: Indicates the time taken by SAM to evaluate a resource group.	Secs	This indicates the mean latency of the last 100 resource group evaluations performed for authentication. A high value could indicate a bottleneck.

3.2.11 Trust Relation Test

Trusts are relationships that are established between domains or forests that enable users in one domain or forest to be authenticated by a domain controller in another domain or forest. Trusts allow users in one domain or forest to access resources in a different domain or forest.

This test automatically discovers the trust relationship that the configured domain shares with other domains, and brings to light problems (if any).

Note:

This test will not work on an Active Directory server running on Windows 2000.

Purpose	Automatically discovers the trust relationship that the configured domain shares with other domains, and brings to light problems (if any)
Target of the test	An AD server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Windows server
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every AD server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Trust errors:
Indicates the number of errors in the trust relationship between the configured domain and other domains.</td><td>Number</td><td>Ideally, this value should be 0. In the event of the occurrence of one/more errors, you can use the detailed diagnosis capability of this measure to view elaborate error descriptions, and accordingly investigate the problem further.</td></tr>
</table>

3.3 The DNS/DHCP Layer

The tests mapped to this layer perform periodic health checks on the DNS and DHCP services that AD relies on.

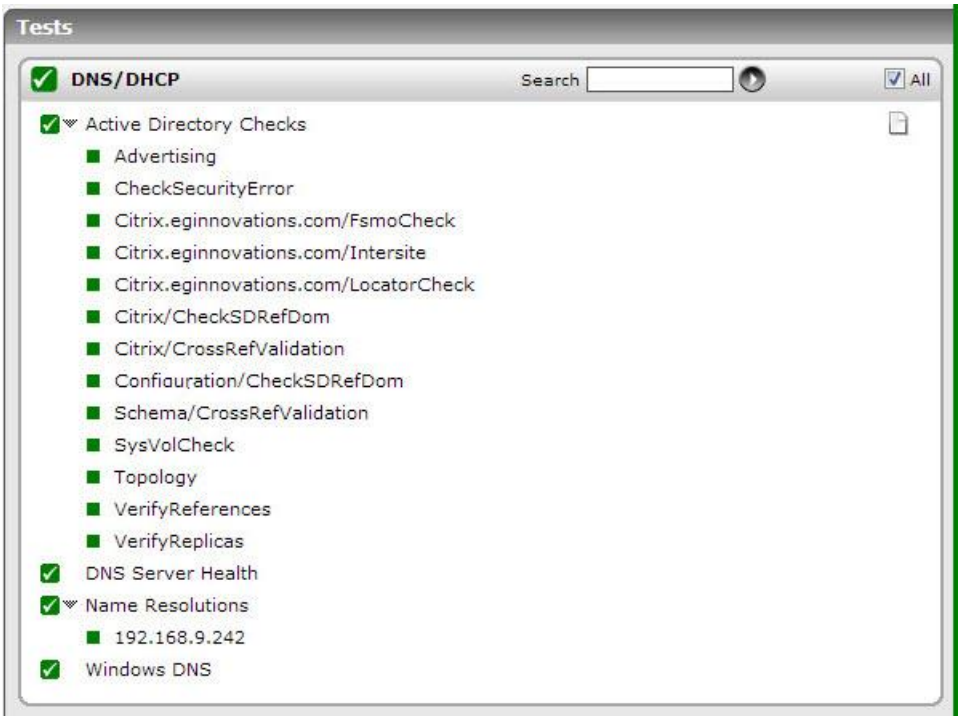

Figure 3.3: The tests mapped to the DNS/DHCP layer

3.3.1 Active Directory Checks Test

Domain controllers are the backbone of a Windows network. If your domain controllers are not working then the Active Directory does not work either. If the Active Directory does not work, then users cannot log on, group policies cannot be enforced, and a whole slew of other features become unavailable. To enable administrators to quickly detect and troubleshoot issues with the domain controller before they affect the operations of the AD server, Windows ships with a specialized tool called the Domain Controller Diagnostic (**DCDIAG**) Utility. **DCDIAG** is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in the system. The tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in troubleshooting. It consists of a framework for executing tests and a series of tests to verify different functional areas of the system - eg., replication errors, domain controller connectivity, permissions, proper roles, etc.

Using the **Active Directory Checks** test, the eG Enterprise Suite leverages the **DCDIAG** utility's ability to report on a wide variety of health parameters related to the domain controller. This ensures that even the smallest of aberrations in the performance of the domain controller is captured and promptly brought to the attention of the administrators. The **Active Directory Checks** test executes the **DCDIAG** command at configured intervals, and based on the output of the command, discovers the **DCDIAG** health checks that were performed, and the current status of each check - whether it reported a success or an error. In case the check resulted in an error/failure, you can use the detailed diagnosis of the test to understand the reason for the same, so that troubleshooting is easier!

Note:

For this test to run, the **DCDIAG.exe** should be available in the **<WINDOWS_INSTALL_DIR>\windows\system32** directory of the AD server to be monitored. The **DCDIAG** utility ships with the Windows Server 2003 Support Tools and is built into Windows 2008 R2 and Windows Server 2008. This utility may hence not be available in older versions of the Windows operating sytem. When monitoring the AD server on such Windows hosts, this test will run only if the **DCDIAG.exe** is copied from the **<WINDOWS_INSTALL_DIR>\windows\system32** directory on any Windows 2003 (or higher) host in the environment to the same directory on the target host.

Purpose	Executes the **DCDIAG** command at configured intervals, and based on the output of the command, discovers the **DCDIAG** health checks that were performed, and the current status of each check - whether it reported a success or an error
Target of the test	An Active Directory or Domain Controller on Windows 2003 or above
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389. 4. **DOMAIN**, **USERNAME**, **PASSWORD**, and **CONFIRM PASSWORD** - In order to execute the **DCDIAG** command, the eG agent has to be configured with a *domain administrator's* privileges. Therefore, specify the domain name and login credentials of the *domain administrator* in the **DOMAIN**, **USERNAME** and **PASSWORD** text boxes. Confirm the **PASSWORD** you provide by retyping it in the **CONFIRM PASSWORD** text box. 5. **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.
Outputs of the test	One set of results for every **DCDIAG** health check that was performed

<table>
<tr><th>Measurements made by the test</th><th>Measurement</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td>Status:

Indicates the status of this DCDIAG health check.</td><td></td><td>If the health check returns a positive result, the value of this measure will be Pass. If not, the value of this measure will be Fail. The numeric values that correspond to these measure values have been discussed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Pass</td><td>1</td></tr><tr><td>Fail</td><td>0</td></tr></table>
Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail, you can use the detailed diagnosis of this measure to know the reason for the failure and the domain controller where the failure occurred. This eases the pain involved in troubleshooting problem conditions.</td></tr>
</table>

3.3.2 AD Checks Test

Domain controllers are the backbone of a Windows network. If your domain controllers are not working then the Active Directory does not work either. If the Active Directory does not work, then users cannot log on, group policies cannot be enforced, and a whole slew of other features become unavailable. To enable administrators to quickly detect and troubleshoot issues with the domain controller before they affect the operations of the AD server, Windows ships with a specialized tool called the Domain Controller Diagnostic (**DCDIAG**) Utility. **DCDIAG** is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in the system. The tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in troubleshooting. It consists of a framework for executing tests and a series of tests to verify different functional areas of the system - eg., replication errors, domain controller connectivity, permissions, proper roles, etc.

Using the **AD Checks** test, the eG Enterprise Suite leverages the **DCDIAG** utility's ability to report on a wide variety of health parameters related to the domain controller. This ensures that even the smallest of aberrations in the performance of the domain controller is captured and promptly brought to the attention of the administrators. The **AD Checks** test executes the **DCDIAG** command at configured intervals, and based on the output of the command, reports the count of **DCDIAG** health checks (i.e., tests) that succeeded and failed in the last measurement period. The detailed diagnosis of the **AD Checks** test will provide detailed information pertaining to tests that failed, and thus assists in troubleshooting.

Note:

For this test to run, the **DCDIAG.exe** should be available in the **<WINDOWS_INSTALL_DIR>\windows\system32** directory of the AD server to be monitored. The **DCDIAG** utility ships with the Windows Server 2003 Support Tools and is built into Windows 2008 R2 and Windows Server 2008. This utility may hence not be available in older versions of the Windows operating sytem. When monitoring the AD server on such Windows hosts, this test will run only if the **DCDIAG.exe** is copied from the **<WINDOWS_INSTALL_DIR>\windows\system32** directory on any Windows 2003 (or higher) host in the environment to the same directory on the target host.

This test is disabled by default. To enable the test, follow the *Agents -> Tests -> Enable/Disable* menu sequence, pick **Active Directory** as the **Component type**, select **Performance** as the **Test type**, select this test from the **DISABLED TESTS** list and click the **<<** button.

Purpose	Executes the **DCDIAG** command at configured intervals, and based on the output of the command, reports the count of **DCDIAG** health checks (i.e., tests) that succeeded and failed in the last measurement period
Target of the test	An Active Directory or Domain Controller on Windows 2008
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The IP address of the machine where the Active Directory is installed.
3. PORT – The port number through which the Active Directory communicates. The default port number is 389.
4. DOMAIN, USERNAME, PASSWORD, and CONFIRM PASSWORD - In order to execute the DCDIAG command, the eG agent has to be configured with a domain administrator's privileges. Therefore, specify the domain name and login credentials of the domain administrator in the DOMAIN, USERNAME and PASSWORD text boxes. Confirm the PASSWORD you provide by retyping it in the CONFIRM PASSWORD text box.
5. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Active Directory being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Passed tests:
Indicates the number of DCDIAG tests that succeeded during the last measurement period.</td><td>Number</td><td></td></tr>
<tr><td>Failed tests:
Indicates the number of DCDIAG tests that failed during the last measurement period.</td><td>Number</td><td>A non-zero value for this measure indicates the existence of one/more errors in the functioning of the domain controller. To know what these errors are, use the detailed diagnosis of this measure.</td></tr>
</table>

3.3.3 DNS Server Health Test

If the DNS component of the AD server is unable to provide domain name resolution services, then users may be denied access to their mission-critical servers managed by the AD server. Under such circumstances, you may want to quickly check what is stalling the operations of DNS, so that the source of the issue can be isolated and eliminated.

DCDIAG is a command-line tool that encapsulates detailed knowledge of how to identify abnormal behavior in the system. The tool analyzes the state of one or all domain controllers in a forest and reports any problems to assist in

troubleshooting. It consists of a framework for executing tests and a series of tests to verify different functional areas of the system.

DCDIAG also performs seven DNS-centric health checks to report on the overall DNS health of the domain controllers. To know the current status of each of these seven health checks, use the **DNS Server Health** test. The periodic health reports provided by the **DNS Server Health** test will enable administrators to proactively isolate potential DNS-related issues with their domain controllers, determine the reason for these issues, and work towards preventing them.

Purpose	Reports the current status of the seven DNS-related health checks that **DCDIAG** performs on the domain controllers		
Target of the test	An Active Directory or Domain Controller on Windows 2008		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389. 4. **DOMAIN**, **USERNAME**, **PASSWORD**, and **CONFIRM PASSWORD** - In order to execute the **DCDIAG** command, the eG agent has to be configured with a *domain administrator's* privileges. Therefore, specify the domain name and login credentials of the *domain administrator* in the **DOMAIN**, **USERNAME** and **PASSWORD** text boxes. Confirm the **PASSWORD** you provide by retyping it in the **CONFIRM PASSWORD** text box. 5. **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.		
Outputs of the test	One set of results for every Active Directory server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td>test</td>
<td>Authentication:

This test is run by default and checks the following:

• Are domain controllers registered in DNS?

• Can they be pinged?

• Do they have Lightweight Directory Access Protocol/Remote Procedure Call (LDAP/RPC)?

This measure reports the current status of the Authentication or Connectivity test.</td>
<td></td>
<td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Pass</td><td>1</td></tr><tr><td>Fail</td><td>0</td></tr><tr><td>Warning</td><td>2</td></tr></table>
Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Basic:

The basic DNS test confirms the following:

a. Whether the DNS client, Netlogon, KDC, and DNS Server services are running and available on domain controllers tested by dcdiag

b. Whether the DNS servers on all adapters are reachable.

c. Whether A record of each domain controller is registered on at least one of the DNS servers configured on the client.

d. If a domain controller is running the DNS Server service, whether the Active Directory domain zone and SOA record for the Active Directory domain zone are present.

e. Whether the root (.) zone is present.

This measure reports the current status of the Basic test.</td>
<td>Number</td>
<td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Pass</td><td>1</td></tr>
<tr><td>Fail</td><td>0</td></tr>
<tr><td>Warning</td><td>2</td></tr>
</table>

Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Forwarders:

The forwarder test determines whether recursion is enabled. If forwarders or root hints are configured, the forwarder test confirms that all forwarders or root hints on the DNS server are functioning, and also confirms that the _ldap._tcp.<Forest root domain> DC Locator record is resolved.</td>
<td></td>
<td rowspan="2">The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Pass</td><td>1</td></tr><tr><td>Fail</td><td>0</td></tr><tr><td>Warning</td><td>2</td></tr></table>
Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
<tr>
<td></td>
<td>This measure reports the current status of the Forwarder test.</td>
<td></td>
</tr>
<tr>
<td></td>
<td>Delegations:

The delegation test confirms that the delegated name server is a functioning DNS Server. The delegation test checks for broken delegations by ensuring that all NS records in the Active Directory domain zone in which the target domain controller resides have corresponding glue A records.

This measure reports the current status of the Delegation test.</td>
<td></td>
<td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Pass</td><td>1</td></tr><tr><td>Fail</td><td>0</td></tr><tr><td>Warning</td><td>2</td></tr></table>
Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Dynamic update:

The dynamic update test confirms that the Active Directory domain zone is configured for secure dynamic update and performs registration of a test record (_dcdiag_test_record).

This measure reports the current status of the Dynamic Update test.</td>
<td></td>
<td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Pass</td><td>1</td></tr>
<tr><td>Fail</td><td>0</td></tr>
<tr><td>Warning</td><td>2</td></tr>
</table>

Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
<tr>
<td></td>
<td>Record registration:

The record registration test verifies the registration of all essential DC Locator records on all DNS Servers configured on each adapter of the domain controllers.

This measure reports the current status of the Record Registration test.</td>
<td></td>
<td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Pass</td><td>1</td></tr>
<tr><td>Fail</td><td>0</td></tr>
<tr><td>Warning</td><td>2</td></tr>
</table>

Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>Resolve external name:

The external name resolution test verifies basic resolution of external DNS from a given client, using a sample Internet name (www.microsoft.com), or user-provided Internet name.

This measure reports the current status of the External name resolution test.</td>
<td></td>
<td>The values that this measure reports and their corresponding numeric values have been discussed in the table below:

<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Pass</td><td>1</td></tr>
<tr><td>Fail</td><td>0</td></tr>
<tr><td>Warning</td><td>2</td></tr>
</table>

Note:

By default, the measure reports the Measure Values listed in the table above to indicate the status of a DCDIAG health check. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

If the measure reports the value Fail or Warning, you can use the detailed diagnosis of this measure to know the reason for the failure/warning. This eases the pain involved in troubleshooting problem conditions.</td>
</tr>
</table>

3.3.4 Name Resolutions Test

Active Directory uses DNS as its domain controller location mechanism and leverages the namespace design of DNS in the design of Active Directory domain names. As a result, DNS is positioned within the discoverability and logical structure components of Active Directory technology components. If a user complaints of being unable to access an AD domain, then administrators should first check whether the DNS component of AD is available and is able to resolve the IP address of the domain to its corresponding domain name and vice-versa. This is where, the **Name Resolutions** test will be useful!

This test emulates a client accessing DNS to issue a query. The query can either request DNS to resolve a domain name to an IP address or vice versa. Based on the response reported by the server, measurements are made of the availability and responsiveness of the DNS component of the AD server.

Purpose	To measure the state of the DNS component of AD
Target of the test	An AD server
Agent deploying the test	An external agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port on which the specified host is listening 4. **TARGETS** - The IP address or host name to be resolved during the test. Multiple **TARGETS** can be specified as a comma-separated list. 5. **RECURSIVE** - DNS supports two types of queries. For a non-recursive query, DNS attempts to respond to the request based on its local cache only. For a recursive query, a DNS server may use other DNS servers to respond to a request. The **Recursive** flag can be used to determine the type of queries to be issued to DNS. 6. **DNS PORT** – Specify the port at which the DNS server listens.		
Outputs of the test	One set of results per **TARGET** configured		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	DNS availability: Whether a successful response is received from the DNS component of the target AD server in response to the emulated user request.	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 DNS.
	DNS response time: Time taken (in seconds) by DNS 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 DNS, a network problem, etc.

3.3.5 Windows DNS Test

This test measures the workload and processing ability of the DNS component of the AD server.

Purpose	Measures the workload and processing ability of the DNS component of the AD server		
Target of the test	An AD server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for the AD server being monitored		
Measurements	**Measurement**	**Measurement Unit**	**Interpretation**

made by the test	**Total queries:** The rate of queries received by DNS.	Reqs/sec	Indicates the workload of the DNS component of the AD server.
	Total responses: The rate of responses from DNS to clients.	Resp/sec	Ideally, the total responses should match the total queries. Significant differences between the two can indicate that DNS is not able to handle the current workload.
	Recursive queries: The rate of recursive queries successfully handled by DNS.	Reqs/sec	The ratio of recursive queries to total queries indicates the number of queries that required the DNS component on the AD server to communicate with other DNS servers to resolve the client requests.
	Recursive query failures: The rate of recursive queries that could not be resolved by DNS.	Reqs/sec	Query failures can happen due to various reasons - e.g., requests from clients to invalid domain names/IP addresses, failure in the external network link thereby preventing a DNS server from communicating with other DNS servers on the Internet, failure of a specific DNS server to which a DNS server is forwarding all its requests, etc. A small percentage of failures is to be expected in any production environment. If a significant percentage of failures are happening, this could result in application failures due to DNS errors.
	Recursive timeouts: The rate of recursive queries that failed because of timeouts.	Reqs/sec	Timeouts can happen because of a poor external link preventing a DNS server from communicating with others. In some cases, improper/invalid domain name resolution requests can also result in timeouts. DNS timeouts can adversely affect application performance and must be monitored continuously.
	Zone transfers received: The number of zone transfer requests received by DNS.	Reqs	Zone transfers are resource intensive. Moreover, zone transfers to unauthorized clients can make an IT environment vulnerable to security attacks. Hence, it is important to monitor the number of zone transfer requests and responses on a periodic basis.

	Zone transfers failed: The number of zone transfers that were not serviced by DNS in the last measurement period.	Reqs	Zone transfers may fail either because the DNS server does not have resources, or the request is not valid, or the client requesting the transfer is not authorized to receive the results.

3.4 The AD Replication Service Layer

The tests mapped to this layer report on the health of the AD replication service.

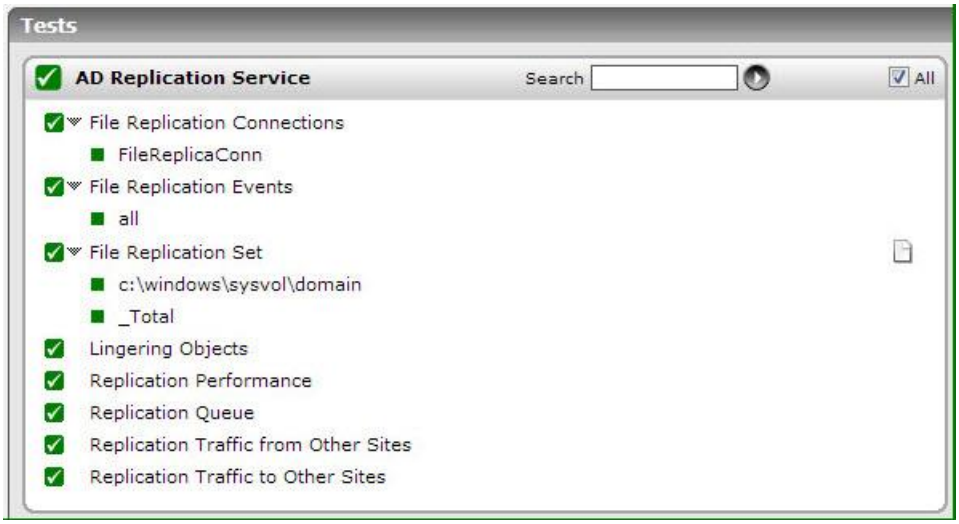

Figure 3.4: The tests mapped to the AD Replication Service layer

3.4.1 File Replication Connections Test

This test reports metrics related to the file replication connections to Distributed File System roots (DFS) in an Active Directory.

Purpose	Reports metrics related to the replica connections to Distributed File System roots (DFS) in an Active Directory
Target of the test	An Active Directory
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.

Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Authentications: Indicates the number of successful authentications that were performed.	Number	
	Bindings: Indicates the number of successful RPC bindings that were completed.	Number	
	Joins: Indicates the number of joins.	Number	After FRS discovers a connection from Active Directory, FRS establishes a connection session with the remote connection partner based on the information provided by the connection object. The connection is called "joined" when a connection session is successfully established.
	Unjoins: Indicates the number of unjoins.	Number	
	Local change orders sent: Indicates the number of local change orders that were sent.	Number	A **change order** is a message that contains information about a file or folder that has changed on a replica. A **local change order** is a change order that is created because of a change to a file or folder on the local server. The local server becomes the originator of the change order and constructs a **staging file –** this file is nothing but a backup of the changed file or folder.
	Packets: Indicates the packets that were sent.	Number	
	Remote change orders sent: Indicates the number of remote change orders that were sent.	Number	A remote change order refers to a change order received from an inbound (or upstream) partner that originated elsewhere in the replica set.
	Remote change orders received: Indicates the number of remote change orders that were received.	Number	

3.4.2 File Replication Events Test

This test reports statistical information about the File Replication Service events recorded in the File Replication Service event log. 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 *Active Directory* as the **Component type**, *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.

Purpose	Reports statistical information about the File Replication Service events recorded in the File Replication Service event log
Target of the test	An Active Directory server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the EventLog Service. Here it is null. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. The default value is *application*. 5. **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. 6. **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 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:all:none:all:none:all:none*. 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 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 **Click here** link appears just above the test configuration section, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the **Click here** link 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.

<table>
<tr><td></td><td colspan="3">7. USEWMI - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the USEWMI flag is YES, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows 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.
8. 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.
9. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the FILTER configured</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>File replication errors:
This refers to the number of File Replication Service events that were generated.</td><td>Number</td><td>A very low value (zero) indicates that the File Replication Service is in a healthy state without any potential problems.
An increasing trend or high value indicates the existence of problems like loss of functionality or data.
The detailed diagnosis capability, if enabled, lists the description of specific events.
Please check the Application Logs in the Event Log Viewer for more details.</td></tr>
<tr><td>File replication information count:
This refers to the number of File Replication Service information events generated when the test was last executed.</td><td>Number</td><td>A change in the value of this measure may indicate infrequent but successful operations performed by the File Replication Service.
The detailed diagnosis capability, if enabled, lists the description of specific events.</td></tr>
</table>

	File replication warnings: This refers to the number of warnings that were generated when the test was last executed.	Number	A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems in the File Replication Service. The detailed diagnosis capability, if enabled, lists the description of specific events.
	File replication critical errors: Indicates the number of critical events that were generated when the test was last executed.	Number	This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. A high value of this measure indicates that too many events have occurred, which the File Replication Service cannot automatically recover from. The detailed diagnosis capability, if enabled, provides the description of specific events.
	File replication verbose count: Indicates the number of verbose events that were generated when the test was last executed.	Number	This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better. The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.

3.4.3 File Replication Set Test

In the FRS, the replication of files and directories is according to a predefined topology and schedule on a specific folder. The topology and schedule are collectively called a replica set. A replica set contains a set of replicas, one for each machine that participates in replication.

This test reports statistics related to the health of the replication service provided by every replication set on an AD server.

Purpose	Reports statistics related to the health of the replication service provided by every replication set on an AD server
Target of the test	An Active Directory
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.

Outputs of the test	One set of results for every replication set on the Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Change orders received: Indicates the number of change orders that were currently received by this replica set.	Number	A **change order** is a message that contains information about a file or folder that has changed on a replica. These measures therefore serve as good indicators of the workload on the replica set.
	Change orders sent: Indicates the number of change orders that were currently sent by this replica set by this replica set.	Number	
	Files installed: Indicates the number of file installations.	Number	Installation is the process by which FRS applies a change order to the local file system to restore the file or folder as it is in the upstream partner. If the change order is for a deletion, the file or folder in the local file system is deleted (staging file is not needed). If the change order is for a renaming, the file or folder in the local file system is renamed (staging file is needed). If the change order is for a copying or creation, the file or folder is copied or created (staging file is needed). Installing a file or folder may fail if the file or folder is already opened by another process. If the installation failed, FRS retries installing the file or folder at a later time.
	Packets received: Indicates the number of packets received currently.	Number	In an idle state, there should be no packets received unless a computer is having trouble joining with other computers in the replica set.
	Packets sent: Indicates the number of of packets sent currently.	Number	

	USN records accepted: Indicates the number of USN records that were currently accepted.	Number	Active Directory replication does not primarily depend on time to determine what changes need to be propagated. Instead it uses update sequence numbers (USNs) that are assigned by a counter that is local to each domain controller. Because these USN counters are local, it is easy to ensure that they are reliable and never run backward (that is, they cannot decrease in value). Domain controllers use USNs to simplify recovery after a failure. When a domain controller is restored following a failure, it queries its replication partners for changes with USNs greater than the USN of the last change it received from each partner.
	Staging space free: Indicates the staging space that is currently free.	KB	The Staging Directory is an area where modified files are stored temporarily either before being propagated to other replication partners or after being received from other replication partners. FRS encapsulates the data and attributes associated with a replicated file or directory object in a staging file. FRS needs adequate disk space for the staging area on both upstream and downstream machines in order to replicate files. Typically, if the **Staging space free** measure reports the value *0*, or is found to be dangerously close to 0, it indicates that the staging directory is full. If the staging area is full, the FRS will stop functioning, and will resume only if disk space for the staging area becomes available or if the disk space limit for the staging area is increased. The staging area could get filled up owing to the following reasons: One or more downstream partners are not accepting changes. This could be a temporary condition due to the schedule being turned off and FRS waiting for it to open, or a permanent state because the service is turned off, or the downstream partner is in an error state. The rate of change in files exceeds the rate at which FRS can process them. A parent directory for files that have a large number of changes is failing to replicate, and so, all changes to subdirectories are blocked.
	Staging space in use: Indicates the staging space that is currently in use.	KB	

3.4.4 Replication Performance Test

Replication is the process by which the changes that are made on one domain controller are synchronized with all other domain controllers in the domain that store copies of the same information or replica.

Monitoring the replication operations on an AD server will shed light on the load generated by such operations and helps measure the ability of the AD server to process this load. The **Replication Performance** test does just that. In the process, the test points you to replication-related activities that could be contributing to processing delays (if any) and why. In addition, the test also promptly reports replication errors such as synchronization failures, and compels administrators to do what is necessary to ensure that no non-sync exists in the data that is replicated across the domain controllers in a forest.

This test applies only to Active Directory Servers installed on Windows 2008 or above.

Purpose	Monitors the replication operations on an AD server and sheds light on the load generated by such operations and helps measure the ability of the AD server to process this load		
Target of the test	An Active Directory or Domain Controller on Windows 2008 or above		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	DRA inbound full sync objects remaining: Indicates the number of object updates received in the current directory replication update packet that have not yet been applied to the local server..	Number	
	DRA inbound object updates remaining: Indicates the number of object updates received in the current directory replication update packet that have not yet been applied to the local server.	Number	The value of this measure should be low, with a higher value indicating that the hardware is incapable of adequately servicing replication (warranting a server upgrade).

	Pending replication operations: Indicates the total number of replication operations on the directory that are queued for this server but not yet performed.	Number	A steady increase in the value of this measure could indicate a processing bottleneck.
	Pending replication synchronizations: Indicates the number of directory synchronizations that are queued for this server but not yet processed.	Number	An unusually high value for a long duration may signify that the replication process is not being carried out at the desired rate. Forcing the replication activity may solve this problem.
	Sync failures on schema mismatch: Indicates the number of synchronization requests made to neighbours that failed because their schema are not synchronized.	Number	Ideally, the value of this measure should be 0.
	Sync requests made: Indicates the number of synchronization requests made to neighbors.	Number	
	Sync requests successful: Indicates the number of synchronization requests made to neighbors that were successfully returned.	Number	Ideally, the value of the *Sync requests made* measure should be equal to the value of the *Sync requests successful* measure - meaning, all sync request made should be successful, as one/more sync failures are a cause for concern.
	DRA inbound objects applied rate: Indicates the rate at which replication updates received from replication partners are applied by the local directory service. This counter excludes changes that are received but not applied (because, for example, the change has already been made). This indicates how much replication update activity is occurring on the server as a result of changes generated on other servers.	Appld/Sec	A low value may indicate one of the following • less changes to the objects in the other domains • this domain controller is not applying the changes to the objects at the desired rate. If the object changes are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory. Forcing the replication activity may solve this problem.

	DRA inbound properties applied rate: Indicates the number of properties that are updated due to the incoming property's winning the reconciliation logic that determines the final value to be replicated.	Appld/Sec	A low value may indicate one of the following less changes to the object properties in the other domains this domain controller is not applying the change to the object properties at the desired rate. If the object properties are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory. Forcing the replication activity may solve this problem.
	DRA inbound objects filtered rate: Indicates the number of objects received from inbound replication partners that contained no updates that needed to be applied.	Filtrd/Sec	A high value for this measure indicates that the objects are all static. This problem can be solved by increasing the replication frequency.
	DRA inbound properties filtered rate: Indicates the number of property changes (per second) already seen that were received during the replication.	Filtrd/Sec	A high value for this measure indicates that the properties are all static. This problem can be solved by increasing the replication frequency in the replicated domain.
	DRA inbound bytes total: Indicates the rate at which bytes were replicated in.	Total/Sec	This counter is the sum of the number of uncompressed bytes (never compressed) and the number of compressed bytes (after compression) per second.
	DRA outbound properties: Indicates the number of properties sent per second.	Properties/Sec	This counter tells you whether a source server is returning objects or not. Sometimes, the server might stop working correctly and not return objects quickly or at all.
	DRA outbound objects filtered rate: Indicates the number of objects per second that were determined by outbound replication to have no updates that the outbound partner did not already have.	Filtrd/Sec	A high value for this measure indicates that the objects are all static. This problem can be solved by increasing the replication frequency in the target domain.
	DRA outbound bytes total: Indicates the rate at which bytes were replicated out.	Total/Sec	This counter is the sum of the number of uncompressed bytes (never compressed) per second and the number of compressed bytes (after compression) per second.

3.4.5 Replication Traffic from Other Sites Test

Used in the Active Directory to express proximity of network connection, a **site** is defined as an IP subnetwork. A site consists of one or more subnets (unique network segments). Client machines use site information to find nearby DCs for logon operations. The Active Directory uses site information to help users find the closest machine that offers a needed network or a third-party service.

The Active Directory provides two methods of replication within the Active Directory environment: *intrasite replication* and *intersite replication*. *Intrasite replication* is replication within an Active Directory site. It is based assumption that the IP subnets within a site are well connected and that bandwidth is considered freely available and inexpensive. Because of this assumption, data is sent without compression.

Inter-site replication is replication between Active Directory sites. It is based on the assumption that the WAN is connected by slower links, so it is designed to minimize traffic rather than CPU cycles. Before being sent out, data is compressed to about 10% to 15% of original volume.

By monitoring the replication data flowing into each site, the **Replication Traffic from Other Sites** test helps determine the nature of the inbound traffic handled by every site - whether *inter-site* or *intrasite*, and reveals what type of inbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data compression is required at the source.

This test applies only to Active Directory Servers installed on Windows 2008 or above.

Purpose	By monitoring the replication data flowing into each site, the **Replication Traffic from Other Sites** test helps determine the nature of the inbound traffic handled by every site - whether *inter-site* or *intrasite*, and reveals what type of inbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data compression is required at the source.		
Target of the test	An Active Directory or Domain Controller on Windows 2008 or above		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory site being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**DRA inbound before bytes compression:** Indicates the original size of inbound compressed replication data (kilobytes per second before compression, from DSAs in other sites).	KB/Sec	
	DRA inbound after bytes compression: Indicates the compressed size of inbound replication data (kilobytes per second received after compression, before DSAs in other sites).	KB/Sec	To save bandwidth on the network connection, the bridgehead servers in each site compress the traffic at the expense of additional CPU usage. A high value for this measure indicates that the bridgehead server is receiving high *inter-site* inbound replication traffic. Replication traffic is compressed down to about 40 percent when replication traffic is more than 32 KB in size.
	DRA inbound bytes not compression: Indicates the number of incoming bytes replicated per second that were not compressed at the source (that is, from DSAs in the same site).	KB/Sec	A high value for this measure indicates that the *intra-site* replication traffic is high. Compressing the replication data adds an additional load on the domain controller server. Uncompressed replication traffic preserves server performance at the expense of network utilization.

3.4.6 Replication Traffic to Other Sites Test

Used in the Active Directory to express proximity of network connection, a **site** is defined as an IP subnetwork. A site consists of one or more subnets (unique network segments). Client machines use site information to find nearby DCs for logon operations. The Active Directory uses site information to help users find the closest machine that offers a needed network or a third-party service.

The Active Directory provides two methods of replication within the Active Directory environment: *intrasite replication* and *intersite replication*. *Intrasite replication* is replication within an Active Directory site. It is based assumption that the IP subnets within a site are well connected and that bandwidth is considered freely available and inexpensive. Because of this assumption, data is sent without compression.

Inter-site replication is replication between Active Directory sites. It is based on the assumption that the WAN is connected by slower links, so it is designed to minimize traffic rather than CPU cycles. Before being sent out, data is compressed to about 10% to 15% of original volume.

By monitoring the replication data flowing from each site, the **Replication Traffic to Other Sites** test helps determine the nature of the outbound traffic handled by every site - whether *inter-site* or *intrasite*, and reveals what type of outbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data is to be compressed by the bridgehead server on each site.

This test applies only to Active Directory Servers installed on Windows 2008 or above.

Purpose	By monitoring the replication data flowing from each site, the **Replication Traffic to Other Sites** test helps determine the nature of the outbound traffic handled by every site - whether *inter-site* or *intrasite*, and reveals what type of outbound traffic is high on a site. Using this information, administrators can determine whether or not the replication data has been compressed enough to optimize bandwidth usage, and accordingly decide if more data is to be compressed by the bridgehead server on each site.		
Target of the test	An Active Directory or Domain Controller on Windows 2008 or above		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory site being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	DRA outbound before bytes compression: Indicates the original size of outbound compressed replication data (kilobytes per second before compression, to DSAs in other sites).	KB/Sec	
	DRA outbound after bytes compression: Indicates the compressed size of outbound replication data (kilobytes per second sent after compression to DSAs in other sites).	KB/Sec	To save bandwidth on the network connection, the bridgehead servers in each site compress the traffic at the expense of additional CPU usage. A high value for this measure indicates that the bridgehead server is sending large high *inter-site* inbound replication traffic. Replication traffic is compressed down to about 40 percent when replication traffic is more than 32 KB in size.

	DRA outbound bytes not compression: Indicates the number of outgoing bytes replicated per second that were not compressed at the source (that is, to DSAs in the same site).	KB/Sec	A high value for this measure indicates that the *intra-site* replication traffic is high. Compressing the replication data adds an additional load on the domain controller server. Uncompressed replication traffic preserves server performance at the expense of network utilization.

3.4.7 Replication Queue Test

As the domain controller formulates change requests, either by a schedule being reached or from a notification, it adds a work item for each request to the end of the queue of pending synchronization requests. Each pending synchronization request represents one <source domain controller, directory partition> pair, such as "synchronize the schema directory partition from DC1," or "delete the ApplicationX directory partition."

When a work item has been received into the queue, the domain controller processes the item (begins synchronizing from that source) as soon as the item reaches the front of the queue, and continues until either the destination is fully synchronized with the source domain controller, an error occurs, or the synchronization is pre-empted by a higher-priority operation.

A long replication queue is often an indication that synchronization requests are not swiftly processed by the AD server. If the reasons for the abnormal queue length are not determined quickly and addressed promptly, replication of some changes may be stalled indefinitely causing the source and destination domain controllers to remain 'out-of-sync' for long durations; this in turn may result in users having to work with obsolete data! To prevent such an eventuality, you can use this test to continuously track the replication queue length, so that you can be alerted as soon as the number of work items in the queue crosses an acceptable limit. You can also use the detailed diagnostics of this test to know what type of synchronization requests are in queue, so that you can figure out why the requests are taking too long to be processed.

Purpose	Continuously tracks the replication queue length, so that you can be alerted as soon as the number of work items in the queue crosses an acceptable limit
Target of the test	An Active Directory or Domain Controller
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The IP address of the machine where the Active Directory is installed.
3. PORT – The port number through which the Active Directory communicates. The default port number is 389.
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Active Directory site being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Replication queue size:
Indicates the number of synchronization requests that are currently in the replication queue, awaiting processing.</td><td>Number</td><td>A high value for this measure is a cause for concern, as it indicates that too many synchronization requests are pending processing. This could be due to a severe processing bottleneck on the AD server. Very short replication schedules and large synchronization requests that require a lot of processing time are also factors that can increase the replication queue length.
You can use the detailed diagnosis of this measure to know which requests are yet to be processed, so that you can figure out why there is a delay (if any) in processing.</td></tr>
</table>

3.4.8 Lingering Objects Test

When restoring a backup file, Active Directory generally requires that the backup file be no more than 180 days old. If you attempt to restore a backup that has expired, you may encounter problems due to "lingering objects".

A lingering object is a deleted AD object that re-appears ("lingers") on the restored domain controller (DC) in its local copy of Active Directory. This can happen if, after the backup was made, the object was deleted on another DC more than than 180 days ago.

When a DC deletes an object it replaces the object with a **tombstone** object. The tombstone object is a placeholder that represents the deleted object. When replication occurs, the tombstone object is transmitted to the other DCs, which causes them to delete the AD object as well.

Tombstone objects are kept for 180 days, after which they are garbage-collected and removed.

If a DC is restored from a backup that contains an object deleted elsewhere, the object will re-appear on the restored

DC. Because the tombstone object on the other DCs has been removed, the restored DC will not receive the tombstone object (via replication), and so it will never be notified of the deletion. The deleted object will "linger" in the restored local copy of Active Directory.

Such lingering objects tend to create problems during replication. For instance, if the source domain controller has outdated objects that have been out of replication for more than one tombstone lifetime a failure event will be logged in the Windows event log at the time of replicating from the source. You will have to promptly capture such events, identify the lingering objects, and delete them to ensure that replication resumes. In order to achieve this, you can use the **Lingering Objects** test. This test scans the event logs for replication events related to lingering objects, and promptly alerts you upon the occurrence of such events. Using the detailed diagnosis of the test, you can easily determine the location of the lingering objects, so that you can immediately proceed to remove them. This way, the test ensures that the replication engine operates without a glitch.

This test works only on Active Directory servers that operate on Windows 2008 or above.

Purpose	Scans the event logs for replication events related to lingering objects, and promptly alerts you upon the occurrence of such events		
Target of the test	An Active Directory or Domain Controller on Windows 2008 or above		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389. 4. **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.		
Outputs of the test	One set of results for every Active Directory server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Lingering messages:** Indicates the number of messages that are currently logged in the event log, which contains references to *lingering objects*.	Number	This measure typically captures and reports the number of events with event IDs 1388 and 1988 in the event log. Event ID 1388 indicates that a destination domain controller that does not have strict replication consistency enabled received a request to update an object that does not reside in the local copy of the Active Directory database. In response, the destination domain controller requested the full object from the source replication partner. In this way, a lingering object was replicated to the destination domain controller. Therefore, the lingering object was reintroduced into the directory. Event ID 1988 indicates that a destination domain controller that has strict replication consistency enabled has received a request to update an object that does not exist in its local copy of the Active Directory database. In response, the destination domain controller blocked replication of the directory partition containing that object from that source domain controller. The detailed diagnosis of this test provides the complete description of the events with IDs 1388 and/or 1988 that are logged in the event log. The source domain controller and the lingering objects can be inferred from the event description. Using this information, you can run the **repadmin** command on the source domain controller to delete the lingering objects.

3.4.9 Replication Status Test

This test summarizes the replication state and relative health of an Active Directory forest by inventorying and contacting every domain controller in the forest, and collecting and reporting information such as replication deltas and replication failures. You can thus accurately identify the domain controllers that are prone to frequent failures.

Purpose	Summarizes the replication state and relative health of an Active Directory forest by inventorying and contacting every domain controller in the forest, and collecting and reporting information such as replication deltas and replication failures
Target of the test	An Active Directory or Domain Controller

<table>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The IP address of the machine where the Active Directory is installed.
3. PORT – The port number through which the Active Directory communicates. The default port number is 389.
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every domain controller in an Active Directory forest being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Total replication links:
Indicates the number of replica links for this domain controller.</td><td>Number</td><td>A replica link exists for each naming context on a domain controller. This measure is the sum total of such replica links per domain controller. Please note that this is not the connection objects or replication partners per domain controller.
You can use the detailed diagnosis of this measure to view the complete details of the replica links - this includes the source and destination sites, the source and destination domain controllers, the transport type, the number of link failures (if any), and details of the failures such as when the failure occurred and the failure status.</td></tr>
<tr><td>Replication links failure:
Indicates the total number of replica links on this domain controller that are failing to replicate for one reason or the other. This will never be greater than the Total field.</td><td>Number</td><td>Ideally, the value of this measure should be 0.</td></tr>
</table>

	Percent of replication links failure: Indicates the percentage of failures in relation to the total replica links on this domain controller.	Percent	A low value is desired for this measure. A value close to 100% is a cause for concern, as it indicates that almost all replica links are failing.
	Longest replication gap: Denotes the longest replication gap amongst all replication links on this domain controller.	Secs	Ideally, this value should be less than 1 hour.

3.4.10 Inter-Site Replication Test

Inter-site replication is based on the assumption that the WAN is connected by slower links or site links. It is designed to minimize traffic rather than CPU cycles. In inter-site replication, data is compressed and then sent out.

Bridgehead servers perform directory replication between sites. Only two designated domain controllers talk to each other. These domain controllers are called "Bridgehead servers".

After updates are replicated from one site to the bridgehead server in the other site, the updates are then replicated to other domain controllers within the site through intra-site replication process.

This test applies only to Active Directory Servers installed on Windows 2003.

Purpose	This test monitors the performance of the Active Directory Inter-site replication process in the target environment.		
Target of the test	An Active Directory or Domain Controller on Windows 2003		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**In rate:** This measure indicates the number of inbound kilobytes replicated between sites per second.	KB/Sec	A high value for this measure indicates that the bridgehead server is receiving high inter-site inbound replication traffic.
	Out rate: This measure indicates the number of outbound kilobytes replicated between sites per second.	KB/Sec	A high value indicates that bridgehead server is sending high inter-site outbound replication traffic.

3.4.11 Intra-Site Replication Test

Intra-site replication means replication happening between domain controllers in the same site. Intra-site replication attempts to complete in the fewest CPU cycles possible. Intra-site replication avoids unnecessary network traffic by introducing a change notification mechanism that replaces the usual polling of replication partners for updates. When a change is performed in its database, a domain controller waits for a configurable interval (default 5 minutes) and accepts more changes during this time. Then it sends a notification to its replication partners, which will pull the changes from the source. If no changes are performed for a configurable period (default 6 hours) the domain controller initiates a replication sequence anyway, just to make sure that it did not miss anything.

This test applies only to Active Directory Servers installed on Windows 2003.

Purpose	This test monitors the performance of the Active Directory Intra-site replication process in the target environment.		
Target of the test	An Active Directory or Domain Controller on Windows 2003		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	In rate: This measure indicates the number of inbound kilobytes replicated within the site per second.	KB/Sec	A high value for this measure indicates that the intra-site replication traffic is high.

	Out rate: This measure indicates the number of outbound kilobytes replicated within the site per second.	KB/Sec	A high value for this measure indicates that the intra-site outbound replication traffic is high.

3.4.12 Replication Test

As the number of domain controllers increase, the replication process consumes more network bandwidth. So, replication process should be monitored within the target environment.

This test applies only to Active Directory Servers installed on Windows 2003.

Purpose	This test monitors the performance of the Active Directory replication process in the target environment.		
Target of the test	An Active Directory or Domain Controller on Windows 2003		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	DRA inbound objects applied rate: This measure shows the number of replication updates applied per second that are occurring on this domain controller as a result of changes generated on other domain controllers.	Appld/Sec	A low value may indicate one of the following 1. less changes to the objects in the other domains 2. this domain controller is not applying the changes to the objects at the desired rate. If the object changes are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory. Forcing the replication activity may solve this problem.

	DRA inbound properties applied rate: This measure indicates the number of changes applied to object properties per second through inbound replication as a result of reconciliation logic. This logic is used to determine the final value to be replicated to the property.	AppId/Sec	A low value may indicate one of the following 1. less changes to the object properties in the other domains 2. this domain controller is not applying the change to the object properties at the desired rate. If the object properties are not applied at the desired rate, it may result in a loss of data integrity in the Active Directory. Forcing the replication activity may solve this problem.
	DRA inbound objects filtered rate: This measure indicates the number of inbound replication objects received per second from the replication partners that contained no updates that needed to be applied.	Filtrd/Sec	A high value for this measure indicates that the objects are all static. Increasing the replication frequency may solve this problem.
	DRA inbound properties filtered rate: This measure indicates the number of inbound replication properties received per second from the replication partners that did not contain any updates to be applied.	Filtrd/Sec	A high value for this measure indicates that the properties are all static. Increasing the replication frequency in the replicated domain may solve this problem.
	DRA outbound objects filtered rate: This measure indicates the number of outbound replication objects that have not yet been received by the outbound replication partner per second.	kerFiltrd/Sec	A high value for this measure indicates that the objects are all static. Increasing the replication frequency in the target domain may solve this problem.
	Pending replication synchronizations: This measure indicates the number of directory synchronizations that are queued per second for this domain controller but not yet processed.	Number	An unusually high value for a long duration may signify that the replication process is not being carried out at the desired rate. Forcing the replication activity may solve this problem.

3.4.13 AD Replications Test

Replication is the process by which the changes that are made on one domain controller are synchronized with all other domain controllers in the domain that store copies of the same information or replica. Given the various types of information that Active Directory can store, changes to Active Directory can swiftly accumulate across multiple domain controllers in a large organization. It is therefore necessary for Windows to frequently synchronize the domain controllers through the replication process. If replication fails, it causes Active Directory objects that represent the replication topology, replication schedule, domain controllers, users, computers, passwords, security groups, group memberships, and Group Policy to be inconsistent between domain controllers. Directory inconsistency causes either operational failures or inconsistent results, depending on the domain controller that is contacted for the operation at hand.

To avoid such inconsistencies, its best to capture failures promptly, isolate the source of failures, and fix them, The **AD Replications** test aids in this regard. This test closely monitors the replication activities on the domain controller and promptly reports replication failures, so that administrators can investigate such failures, discover the reasons for the same, fix them, and restore normalcy.

Purpose	Closely monitors the replication activities on the domain controller and promptly reports replication failures, so that administrators can investigate such failures, discover the reasons for the same, fix them, and restore normalcy		
Target of the test	An Active Directory or Domain Controller on Windows		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Replication failures: Indicates the number of replication failures in the target domain controller.	Number	Ideally, the value of this measure should be low.
	Total replications: Indicates the number of replication successes in the target domain controller.	Number	

	Percent replication failures: Indicates the percentage of replication failures in the target domain controller.	Percent	Ideally, the value of this measure should be low. A high value is indicative of too many replication failures. Active Directory replication problems can have several different sources. For example, Domain Name System (DNS) problems, networking issues, or security problems can all cause Active Directory replication to fail. • Network connectivity: The network connection might be unavailable or network settings are not configured properly. • Name resolution: DNS misconfigurations are a common cause for replication failures. • Authentication and authorization: Authentication and authorization problems cause "Access denied" errors when a domain controller tries to connect to its replication partner. • Directory database (store): The directory database might not be able to process transactions fast enough to keep up with replication timeouts. • Replication engine: If intersite replication schedules are too short, replication queues might be too large to process in the time that is required by the outbound replication schedule. In this case, replication of some changes can be stalled indefinitely — potentially, long enough to exceed the tombstone lifetime. • Replication topology: Domain controllers must have intersite links in Active Directory that map to real wide area network (WAN) or virtual private network (VPN) connections. If you create objects in Active Directory for the replication topology that are not supported by the actual site topology of your network, replication that requires the misconfigured topology fails.

3.4.14 Distributed File System Events Test

If you are suspecting that the DFS replication between members is failing, then use the **DFS Replication** log to confirm your suspicions or to negate them. This event log records events for the Distributed File System Replication services, such as when the DFS replication service started, and also captures service failures (if any). This way, the DFS Replication log serves as a rich source of information that is most useful when troubleshooting issues related to replication. By monitoring this event log, the **Distributed File System Events** test promptly alerts administrators to current replication problems and even warns them of probable replication failures.

Purpose	Monitors the **DFS Replication** log and promptly alerts administrators to current replication problems and even warns them of probable replication failures

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 EventLog Service. Here it is null. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. By default, this is set to *DFS Replication*. 5. **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. 6. **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 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:all:none:all:none:all:none*. 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 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 **Click here** link appears just above the test configuration section, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the **Click here** link 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.

7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs 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'.**

8. **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.

9. **EVENTS DURING RESTART** - By default, the **EVENTS DURING RESTART** flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.

10. **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**.

11. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.

12. **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**.

<table>
<tr><td></td><td colspan="3">13. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the FILTER configured</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Distributed file system information messages:

This refers to the number of information events that were captured by the DFS Replication log during the test's last execution.</td><td>Number</td><td>A change in value of this measure may indicate infrequent but successful replications.

Please check the DFS Replication log in the Event Log Viewer for more details.</td></tr>
<tr><td>Distributed file system warnings:

This refers to the number of warning events captured by the DFS Replication log during the test's last execution.</td><td>Number</td><td>A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems.

Please check the DFS Replication log in the Event Log Viewer for more details.</td></tr>
<tr><td></td><td>Distributed file system errors:

This refers to the number of error events captured by the DFS Replication log during the test's last execution.</td><td>Number</td><td>A very low value (zero) is desired for this measure, as it indicates good health.

An increasing trend or a high value indicates the existence of problems.

Please check the DFS Replication log in the Event Log Viewer for more details.</td></tr>
</table>

	Distributed file system critical errors: Indicates the number of critical events that were generated when the test was last executed.	Number	A critical event is one that the replication service cannot automatically recover from. This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. A very low value (zero) indicates that the service is in a healthy state and is running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems. The detailed diagnosis of this measure describes all the critical events captured by the *DFS Replication* log during the last measurement period. Please check the *DFS Replication* log in the Event Log Viewer for more details.
	Distributed file system verbose messages: Indicates the number of verbose events that were generated when the test was last executed.	Number	Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better. This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. The detailed diagnosis of this measure describes all the verbose events that were captured by the *DFS Replication* log during the last measurement period. Please check the *DFS Replication* log in the Event Log Viewer for more details.

3.5 The AD Service Layer

This layer tracks the health of the Active Directory in a Windows environment using the ActiveDirectory test shown in Figure 3.5.

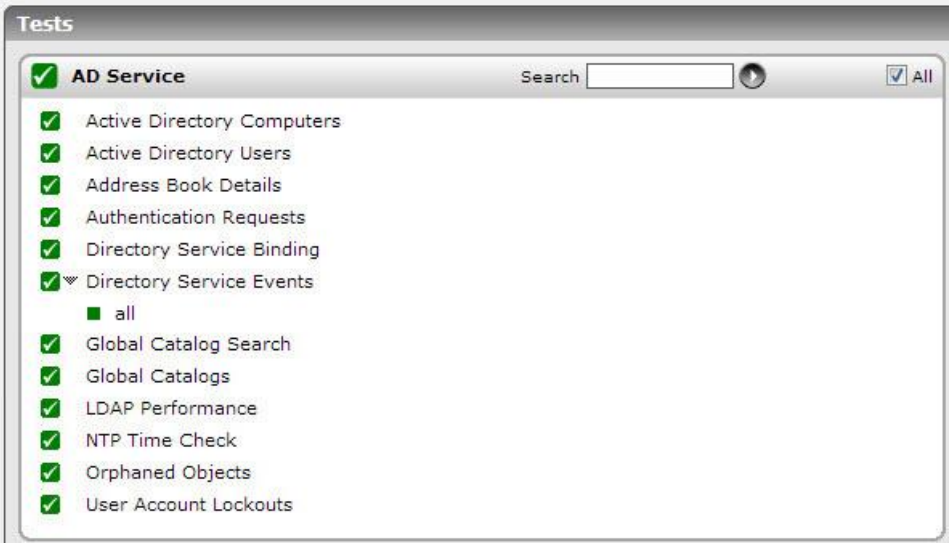

Figure 3.5: Tests mapping to the DC Service layer

3.5.1 Orphaned Objects Test

On a domain controller, the Lost and Found container contains Active Directory objects that have been orphaned. An object is orphaned when the object is created on one domain controller and the container in which the object is placed is deleted from the directory on another domain controller before the object has a chance to replicate. An orphaned object is automatically placed in the Lost and Found container where it can be found by an administrator, who must determine whether to move or delete the object.

The Orphaned Objects test periodically reports the number of orphaned objects on a domain controller.

Purpose	Periodically reports the number of orphaned objects on a domain controller
Target of the test	An AD server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Windows server
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every AD server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Orphaned objects:
Indicates the number of objects in the Lost and Found container.</td><td>Number</td><td>If the value of this measure is greater than 0, it indicates the existence of orphaned objects. In such a case, you can use the detailed diagnosis capability of this measure to view the complete details of the objects, and accordingly decide whether to move the object or delete it.</td></tr>
</table>

The detailed diagnosis of the *Orphaned objects* measure, if enabled, provides the complete details of the orphaned objects, which includes the named of the **Object class** and **Distinguished name**.

Detailed Diagnosis measure of Lost and Found Container in Active Directory

Time	Object class	Distinguished name
6.7.09 11:40:23	lostAndFound	CN=LostAndFound,DC=TESTSUB2,DC=TESTMAIN,DC=COM

Figure 3.6: The details of orphaned objects

3.5.2 Active Directory Status Test

This test tracks the performance of Active Directory existing in a Windows 2000 environment. Before getting into the details of this test, it is essential for the users to know that there are two choices for network authentication in a Windows 2000 environment. They are

1. **Kerberos Version 5.0**: This protocol is the default network authentication protocol for Windows 2000 servers.
2. **Windows NT LAN Manager (NTLM)**: The NTLM protocol was the default network authentication protocol for Windows NT 4.0 operating system. NTLM is also used to authenticate logons to standalone computers with Windows 2000.

When a user first authenticates to Kerberos, he/she talks to the Authentication Service (AS) on the Kerberos Key Distribution Center (KDC) to get a Ticket Granting Ticket (TGT). This ticket is encrypted with the user's password. When the user wants to talk to a Kerberized service, he/she uses the Ticket Granting Ticket (TGT) to talk to the Ticket Granting Service (TGS), which also runs on the KDC. The Ticket Granting Service then verifies the user's identity using the TGT and issues a ticket for the desired service. The reason the Ticket Granting Ticket exists is that a user doesn't have to enter their password every time they wish to connect to a Kerberized service.

The outputs of the ActiveDirectoryStatus Test are given below:

Purpose	This test monitors the performance of Active Directory in a Windows 2000 environment.		
Target of the test	An Active Directory or Domain Controller		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Schema cache hit ratio: This measure shows the percentage of object name lookups available in the Schema Cache. This cache is present in the Domain Controller. All changes made to the Active Directory are first validated against this schema cache.	Percent	A low value of this measure indicates that the Directory Service needs high disk read/write activity to perform its job. This results in poor response time of the components available in the Active Directory.
	Notify queue size: When any change in the Active Directory occurs, the originating domain controller sends an update notification requests to the other domain controllers. This measure shows the number of pending update notification requests that have been queued and not transmitted.	Number	A high value of this measure indicates that the Active Directory is changing frequently but the update notification requests have not been transmitted to the other domain controllers. This results in a loss of data integrity in the directory store. This problem can be corrected by forcing the replication process.

	Current threads: This measure shows the number of threads that are currently servicing the API calls by the users.	Number	A fluctuating value for this measure indicates a change in the load.
	Directory writes: This measure shows the number of successful write operations made by the directory service per second.	Writes/Sec	A high value for this measure indicates that the directory service has made write operations in the Active Directory. This results in the fragmentation of the Active Directory. This problem can be corrected by forcing the replication process.
	Kerberos requests: This measure shows the number of times per second that the user uses the user credentials to authenticate himself or herself with the domain controller that is being monitored.	Reqs/Sec	A high value for this measure indicates that the user requested some network resource, which requires authentication. Installing one or more Active Directory in the target environment can solve this problem
	NTLM requests: This measure shows the number of times per second that the user uses the user credentials to authenticate himself or herself with the domain controller, which is having the PDC emulator operation role.	Reqs/Sec	A high value for this measure indicates that the user requested some network resource, which basically belongs to the Windows NT network. Accessing this kind of resource needs authentication, which is serviced by the domain controller, who is having the PDC emulator operation role. Installing one or more domain controllers with PDC emulator operation role in the target environment can solve this problem.
	Ticket requests: This measure indicates the number of requests made by the Ticket Granting Service per second.	Reqs/Sec	A high value for this measure indicates that the user requested some network resources, which needs authentication. Installing one or more domain controllers in the target environment can solve this problem.
	Authentication requests: This measure indicates the number of requests made by the Authentication Server (to obtain the TGT) per second.	Reqs/Sec	A high value for this measure indicates that the user requested some network resources, which needs authentication. Installing one or more domain controllers in the target environment can solve this problem.

	Ldap sessions: This measure indicates the number of Ldap clients currently connected to the Active Directory.	Number	This measure is just an indicator of the number of Ldap clients connected to the Active Directory. A high or low value for this measure does not always denote an error situation.

3.5.3 Directory Service Events Test

This test reports statistical information about the Directory Service events recorded in the event log. 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 *Active Directory* as the **Component type**, *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.

Purpose	Reports statistical information about the Directory Service events recorded in the event log
Target of the test	An Active Directory server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the EventLog Service. Here it is null. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. The default value is *application*. 5. **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. 6. **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 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:all:none:all:none:all:none*. 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 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 **Click here** link appears just above the test configuration section, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the **Click here** link 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.

	7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows 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**. 8. **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**. 9. **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.
Outputs of the test	One set of results for the **FILTER** configured

Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Directory service errors: This refers to the number of Directory Service events that were generated.	Number	A very low value (zero) indicates that the Directory Service is in a healthy state without any potential problems. An increasing trend or high value indicates the existence of problems like loss of functionality or data. The detailed diagnosis capability, if enabled, lists the description of specific events. Please check the Application Logs in the Event Log Viewer for more details.
	Directory service information count: This refers to the number of Directory Service Service 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 the Directory Service. The detailed diagnosis capability, if enabled, lists the description of specific events.

	Directory service warnings: This refers to the number of warnings that were generated when the test was last executed.	Number	A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems in the Directory Service. The detailed diagnosis capability, if enabled, lists the description of specific events.
	Directory service critical errors: Indicates the number of critical events that were generated when the test was last executed.	Number	This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. A high value of this measure indicates that too many errors have occurred, which the Directory Service cannot automatically recover from. The detailed diagnosis capability, if enabled, provides the description of specific events.
	Directory service verbose count: Indicates the number of verbose events that were generated when the test was last executed.	Number	This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better. The detailed diagnosis of this measure describes all the verbose events that were generated during the last measurement period.

3.5.4 User Account Lockouts Test

Account lockout is a feature of password security that disables a user account when a certain number of failed logons occur due to wrong passwords within a certain interval of time. The purpose behind account lockout is to prevent attackers from brute-force attempts to guess a user's password.

Other ways accounts can get locked out include:

- Applications using cached credentials that are stale.
- Stale service account passwords cached by the Service Control Manager (SCM).
- Stale logon credentials cached by Stored User Names and Passwords in Control Panel.
- Scheduled tasks and persistent drive mappings that have stale credentials.
- Disconnected Terminal Service sessions that use stale credentials.
- Failure of Active Directory replication between domain controllers.
- Users logging into two or more computers at once and changing their password on one of them.

Any one of the above situations can trigger an account lockout condition, and the results can include applications behaving unpredictably and services inexplicably failing.

This is why, whenever a user complaints of inability to login to his/her desktop, help desk should be able to instantly figure out whether that user's account has been locked out, and if so, why. The **User Account Lockouts** test provides answers to these questions. This test, at configured intervals, reports the count of locked user accounts and names the users who have been affected by this anomaly.

Purpose	Reports the count of locked user accounts and names the users who have been affected by this anomaly
Target of the test	An Active Directory
Agent deploying the test	An internal agent**; this test cannot be run in an 'agentless' manner**

Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389. 4. **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.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td rowspan="3">test</td>
<td>Account lockout events:
Indicates the number of account lockouts that occurred during the last measurement period.</td>
<td>Number</td>
<td>A very high value for this measure could indicate a malicious attack, and may require further investigation.
If the high lockout rate is not due to any such attacks, then it is recommended that you alter the lockout policy in your environment to minimize the count and consequently, the impact of account lockouts. Microsoft recommends the following policies for high, medium, and low security environments:
<table>
<tr><th>Security Level</th><th>Lockout Policy</th></tr>
<tr><td>Low</td><td>Account Lockout Duration =Not Defined
Account Lockout Threshold = 0 (No lockout)
Reset account lockout counter after = Not Defined</td></tr>
<tr><td>Medium</td><td>Account Lockout Duration =30 minutes
Account Lockout Threshold = 10 invalid logon attempts
Reset account lockout counter after = 30 minutes</td></tr>
<tr><td>High</td><td>Account lockout duration = 0 (an administrator must unlock the account)
Account lockout threshold = 10 invalid logon attempts
Reset account lockout counter after = 30 minutes</td></tr>
</table>
</td>
</tr>
<tr>
<td>Unique users locked out:
Indicates the number of distinct users who were locked out during the last measurement period.</td>
<td>Number</td>
<td>Use the detailed diagnosis of this measure to view the names of these users.</td>
</tr>
<tr>
<td>Users currently locked out:
Indicates the number of users who are currently locked out.</td>
<td>Number</td>
<td>Use the detailed diagnosis of this measure to know which users are currently locked out.</td>
</tr>
</table>

3.5.5 Active Directory Lost and Found Test

On a domain controller, the Lost and Found container contains Active Directory objects that have been orphaned. An object is orphaned when the object is created on one domain controller and the container in which the object is placed is deleted from the directory on another domain controller before the object has a chance to replicate. An orphaned object is automatically placed in the Lost and Found container where it can be found by an administrator.

This test reports the number of orphaned objects currently in the Lost and Found container, provides the details of these objects, so that administrators can determine which objects to move and which ones to delete.

This test applies only to Active Directory Servers installed on Windows 2008.

This test is disabled by default. To enable the test, follow the *Agents -> Tests -> Enable/Disable* menu sequence, pick **Active Directory** as the **Component type**, select **Performance** as the **Test type**, select this test from the **DISABLED TESTS** list and click the **<<** button.

<table>
<tr><td>Purpose</td><td colspan="3">Reports the number of orphaned objects currently in the Lost and Found container, provides the details of these objects, so that administrators can determine which objects to move and which ones to delete</td></tr>
<tr><td>Target of the test</td><td colspan="3">An Active Directory or Domain Controller on Windows 2008</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The IP address of the machine where the Active Directory is installed.
3. PORT – The port number through which the Active Directory communicates. The default port number is 389.
4. 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
f. Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every Active Directory being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Lost and Found objects:
Indicates the number of objects currently available in the Lost and Found container.</td><td>Number</td><td>A non-zero value indicates the existence of orphaned objects. Use the detailed diagnosis of this measure to know which obects to move and which ones to delete.</td></tr>
</table>

3.5.6 Global Catalog Search Test

The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain Active Directory Domain Services (AD DS) forest. The global catalog is stored on domain controllers that have been designated as global catalog servers and is distributed through multimaster replication.

The Global Catalog enables searching for Active Directory objects in any domain in the forest without the need for subordinate referrals, and users can find objects of interest quickly without having to know what domain holds the object. The global catalog makes the directory structure within a forest transparent to users who perform a search. For example, if you search for all printers in a forest, a global catalog server processes the query in the global catalog and then returns the results. Without a global catalog server, this query would require a search of every domain in the forest.

This test reveals whether the server being monitored is a global catalog server or not. If it is, then the test attempts to search the global catalog server for a configured user and reports whether that user was found or not. The test also reports the time taken to search for that user. This information helps administrators assess how efficient the global catalog is in minimizing the time taken to locate a user across domains.

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	Reveals whether the server being monitored is a global catalog server or not. If it is, then the test attempts to search the global catalog server for a configured user and reports whether that user was found or not. The test also reports the time taken to search for that user. This information helps administrators assess how efficient the global catalog is in minimizing the time taken to locate a user across domains		
Target of the test	An Active Directory or Domain Controller on Windows 2008		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389. 4. **USERNAME** - Specify the name of the user who has to be searched in the global catalog.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Is it a global catalog server?:** Indicates whether the monitored server is a global catalog server or not.	Boolean	This measure reports the value *True* if the AD server being monitored is a global catalog server, and the value *False* if it is not. If this measure reports the value *False*, the remaining measures of the test will not report any values.
	Was user found? Indicates whether the configured **USERNAME** was found or not in the global catalog server.	Boolean	This measure reports the value *True* if the configured **USERNAME** was found in the global catalog server and the value *False* if the user name was not found.
	Catalog search time: Indicates the time taken by the global catalog server to search and find the configured **USERNAME**.	Secs	A high value for this measure would warrant an investigation.

3.5.7 Address Book Details Test

The Address Book is a client for the Active Directory database. It performs lookups and search operations on the Active Directory database to look for details such as account email ID, and so forth. Using the **Address Book Details** test, you can determine the number of Address Book clients currently connected to the AD database and the rate at which search operations are performed by each AD server. In the event that the AD database gets inundated with search queries, you can use this test to figure out whether or not the Address Book clients are contributing to the query load.

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	You can determine the number of Address Book clients currently connected to the AD database and the rate at which search operations are performed by each AD server. In the event that the AD database gets inundated with search queries, you can use this test to figure out whether or not the Address Book clients are contributing to the query load
Target of the test	An Active Directory or Domain Controller on Windows 2008
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.

Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Client sessions: Indicates the number of client sessions that are currently connected to the AD database.	Number	A high value is indicative of heavy load. A consistent increase in the value of this measure could indicate a potential overload condition.
	Search operations: Indicate the rate at which the key search operations are performed on the AD database.	Searche/Sec	If the value of this measure decreases while the number of **Client sessions** keeps increasing, it indicates that search queries are not been processed as quickly; this in turn is indicative of a processing bottleneck, which can consequently choke the AD server database.

3.5.8 ADAM LDAP Performance Test

The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that runs on a layer above the TCP/IP stack. It provides a mechanism used to connect to, search, and modify Internet directories. The LDAP directory service is based on a client-server model. The function of LDAP is to enable access to an existing directory. LDAP is one of the protocols used to query and modify items on the Active Directory server.

To monitor the interactions between clients and the AD server over LDAP, and to promptly capture slowdowns in LDAP searches and binds, use the **ADAM LDAP Performance** test.

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	To monitor the interactions between clients and the AD server over LDAP, and to promptly capture slowdowns in LDAP searches and binds		
Target of the test	An Active Directory or Domain Controller on Windows 2008		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Ldap searches:** Indicates the rate at which LDAP clients perform search operations.	Searches/Sec	This counter should show activity over time. If it does not, network problems are probably hindering the processing of client requests.
	Ldap writes: Indicates the rate at which clients perform write operations on the AD server.	Writes/Sec	
	Ldap active threads: Indicates the current number of threads in use by the LDAP subsystem of the local directory service.	Number	A high number indicates a high level of LDAP activity on the directory service.
	Ldap bind time: Indicates the time, in milliseconds, taken for the last successful LDAP bind.	Secs	In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as *binding.* When a programmatic object, such as an IADs or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be *bound to* the directory object. This measure should be as low as possible. If it is not, hardware or network-related problems are indicated.
	Ldap sessions: Indicates the number of currently connected LDAP client sessions.	Number	This measure is just an indicator of the number of Ldap clients connected to the Active Directory. A high or low value for this measure does not always denote an error situation.
	Ldap closed connections: Indicates the LDAP connections that have been closed in the last second.	Connections/Sec	
	Ldap new connections: Indicates the number of new LDAP connections that have arrived in the last second.	Connections/Sec	
	Ldap new ssl connections: Indicates the number of new SSL or TLS connections that arrived in the last second.	Connections/Sec	

	Ldap successful binds: Indicates the number of successful LDAP binds per second.	Binds/Sec	In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as *binding.* When a programmatic object, such as an IADs or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be *bound to* the directory object. A high value is desired for this measure. A very low value could indicate network problems.

3.5.9 Authentication Performance Test

Authentication of domain user logins is a core function of an Active Directory server. The default authentication protocol used by the AD server is **Kerberos**. Kerberos authentication is based on specially formatted data packets known as tickets. In Kerberos, these tickets pass through the network instead of passwords. Transmitting tickets instead of passwords makes the authentication process more resistant to attackers who can intercept the network traffic.

In a Kerberos environment, the authentication process begins at logon. The following steps describe the Kerberos authentication process:

1. When a user enters a user name and password, the computer sends the user name to the KDC (Key Distribution Center). The Key Distribution Center (KDC) maintains a database of account information for all security principals in the domain. The KDC stores a cryptographic key known only to the security principal and the KDC. This key is used in exchanges between the security principal and the KDC and is known as a long term key. The long term key is derived from a user's logon password.
2. Upon the receipt of a user name, the KDC looks up the user's master key (KA), which is based on the user's password. The KDC then creates two items: a session key (SA) to share with the user and a Ticket-Granting Ticket (TGT). The TGT includes a second copy of the SA, the user name, and an expiration time. The KDC encrypts this ticket by using its own master key (KKDC), which only the KDC knows.
3. The client computer receives the information from the KDC and runs the user's password through a one-way hashing function, which converts the password into the user's KA (i.e., master key). The client computer now has a session key and a TGT so that it can securely communicate with the KDC. The client is now authenticated to the domain and is ready to access other resources in the domain by using the Kerberos protocol.
3. When a Kerberos client needs to access resources on a server that is a member of the same domain, it contacts the KDC. The client will present its TGT and a timestamp encrypted with the session key that is already shared with the KDC. The KDC decrypts the TGT using its KKDC. The TGT contains the user name and a copy of the SA. The KDC uses the SA to decrypt the timestamp. The KDC can confirm that this request actually comes from the user because only the user can use the SA.
4. Next, the KDC creates a pair of tickets, one for the client and one for the server on which the client needs to access resources. Each ticket contains the name of the user requesting the service, the recipient of the request, a timestamp that declares when the ticket was created, and a time duration that says how long the tickets are valid. Both tickets also contain a new key (KAB) that will be shared between the client and the server so they can securely communicate.
5. The KDC takes the server's ticket and encrypts it using the server master key (KB). Then the KDC nests the

server's ticket inside the client's ticket, which also contains the KAB. The KDC encrypts the whole thing using the session key that it shares with the user from the logon process. The KDC then sends all the information to the user.

6. When the user receives the ticket, the user decrypts it using the SA. This exposes the KAB to the client and also exposes the server's ticket. The user cannot read the server's ticket. The user will encrypt the timestamp by using the KAB and send the timestamp and the server's ticket to the server on which the client wants to access resources. When it receives these two items, the server first decrypts its own ticket by using its KB. This permits access to the KAB, which can then decrypt the timestamp from the client.

In situations where a domain controller is not available or is unreachable, **NTLM** (the **NT LAN Manager**) is used as the authentication protocol. For example, NTLM would be used if a client is not Kerberos capable, the server is not joined to a domain, or the user is remotely authenticating over the web.

In some other environments **Digest** authentication is supported. Digest authentication offers the same functionality as Basic authentication; however, Digest authentication provides a security improvement because a user's credentials are not sent across the network in plaintext. Digest authentication sends credentials across the network as a Message Digest 5 (MD5) hash, which is also known as the MD5 message digest, in which the credentials cannot be deciphered from the hash.

Regardless of the protocol/authentication mode used, the quality of a user's experience with the AD server largely relies on how fast his/her login is authenticated by the AD server. The slightest of delays will hence not be tolerated! Administrators therefore need to keep their eyes open at all times for authentication-related latencies, isolate their source, and fix the problems, so that users are able to login to their systems quickly. The **Authentication Performance** test helps administrators in this regard.

This test reports the rate at which Kerberos, NTLM, and Digest authentication requests are serviced by the AD server and thus promptly reveals delays in authentication (if any). Where latencies are noticed in Kerberos requests, the test goes one step further and indicates the probable source of the latencies - could it be because the KDC took too long to grant TGTs to the clients? or is it because the KDC took too long to process the TGTs and grant the clients access to authorized resources?

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	Reports the rate at which Kerberos and NTLM authentication requests are serviced by the AD server and thus promptly reveals delays in authentication (if any). Where latencies are noticed in Kerberos requests, the test goes one step further and indicates the probable source of the latencies - could it be because the KDC took too long to grant TGTs to the clients? or is it because the KDC took too long to process the TGTs and grant the clients access to authorized resources?
Target of the test	An Active Directory or Domain Controller on Windows 2008
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.

Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Kerberos requests: Indicates the number of times per second that clients use a ticket to authenticate to the domain controller.	Reqs/Sec	A low value indicates a bottleneck when processing Kerberos requests.
	Digest requests: Indicates the rate at which requests from a potential user were received by a network server and then sent to a domain controller.	Reqs/Sec	A low value indicates a bottleneck when processing Digest requests.
	Ntlm requests: Indicates the rate at which NTLM authentication requests were serviced by the domain controller.	Reqs/Sec	A low value indicates a bottleneck when processing NTLM requests. A high value for this measure indicates that the user requested some network resource, which basically belongs to the Windows NT network. Accessing this kind of resource needs authentication, which is serviced by the domain controller, who is having the PDC emulator operation role. Installing one or more domain controllers with PDC emulator operation role in the target environment can solve this problem.
	Authentication requests: Indicates the number of Authentication Server (AS) requests serviced by the Kerberos Key Distribution Center (KDC) per second.	Reqs/Sec	AS requests are used by the client to obtain a ticket-granting ticket. If the AD server appears to be taking too long to process Kerberos requests - i.e., if the value of the *Kerberos requests* measure is too high - then you can compare the value of this measure with that of the *Ticket requests* measure to know where the request spent too much time - when granting TGTs to clients? or when processing the TGTs to allow users access to a resource?

	Ticket requests: Indicates the number of Ticket Granting Server (TGS) requests serviced by the KDC per second.	Reqs/Sec	TGS requests are used by the client to obtain a ticket to a resource. If the AD server appears to be taking too long to process Kerberos requests - i.e., if the value of the *Kerberos requests* measure is too high - then you can compare the value of this measure with that of the *Authentication requests* measure to know where the request spent too much time - when granting TGTs to clients? or when processing the TGTs to allow users access to a resource?

3.5.10 ADAM Binding Test

In Active Directory Domain Services, the act of associating a programmatic object with a specific Active Directory Domain Services object is known as *binding.* When a programmatic object, such as an IADs or DirectoryEntry object, is associated with a specific directory object, the programmatic object is considered to be *bound to* the directory object.

This test reports the type of binds that exist in an AD environment, and for each bind type, reports how fast the AD server bound the programmatic objects to the directory object.

This test applies only to Active Directory Servers installed on Windows 2008.

Purpose	Reports the type of binds that exist in an AD environment, and for each bind type, reports how fast the AD server bound the programmatic object to the directory object		
Target of the test	An Active Directory or Domain Controller on Windows 2008		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Ntlm binds: Indicates the rate at which programmatic and directory objects were bound to one another using *NTLM binds*.	Binds/Sec	

	Simple binds: Indicates the rate at which programmatic and directory objects were bound to one another using *Simple binds*.	Binds/Sec	In a simple bind, the client either binds anonymously, that is, with an empty bind Distinguished Name, or by providing a Distinguished Name and a password.
	External binds: Indicates the rate at which programmatic and directory objects were bound to one another using *External binds*.	Binds/Sec	
	Fast binds: Indicates the rate at which programmatic and directory objects were bound to one another using *Fast binds*.	Binds/Sec	Fast bind mode allows a client to use the LDAP bind request to simply validate credentials and authenticate the client without the overhead of establishing the authorization information.
	Negotiated binds: Indicates the rate at which programmatic and directory objects were bound to one another using *Negotiated binds*.	Binds/Sec	

3.5.11 Global Catalogs Test

The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain active directory domain services (AD DS) forest. The global catalog is stored on domain controllers that have been designated as global catalog servers and is distributed through multimaster replication.

The global catalog enables searching for active directory objects in any domain in the forest without the need for subordinate referrals, and users can find objects of interest quickly without having to know what domain holds the object. The global catalog makes the directory structure within a forest transparent to users who perform a search. For example, if you search for all printers in a forest, a global catalog server processes the query in the global catalog and then returns the results. Without a global catalog server, this query would require a search of every domain in the forest.

This test monitors the global catalogs in the target domain controller and reports the number of catalogs that are currently available and unavailable. This way, the test enables administrators to determine whether/not adequate global catalogs are available in the domain controller to handle the request load.

Purpose	Monitors the global catalogs in the target domain controller and reports the number of catalogs that are currently available and unavailable
Target of the test	An Active Directory or Domain Controller on Windows
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total global catalogs: Indicates the total number of global catalogs on the domain controller being monitored.	Number	
	Available global catalogs: Indicates the number of global catalogs that are currently available on the domain controller.	Number	
	Unavailable global catalogs: Indicates the number of global catalogs that are currently unavailable on the domain controller.	Number	If the value of this measure is equal to the value of the *Total global gatalogs* measure or is higher than that of the *Available global catalogs* measure, it indicates that enough global catalogs may not be available on the domain controller to process user logon requests and search requests. As a result, requests may fail.
	Percent unavailable global catalogs: Indicates percentage of global catalogs that are currently unavailable.	Percent	A high value indicates that too many global catalogs are unavailable for request processing. This in turn can cause many user logon and search requests to the domain controller to fail. Ideally therefore, the value of this measure should be very low.

3.5.12 Active Directory Users

This test reports the status of user accounts configured in the Active Directory server and thus, quickly points you to 'unused' accounts that can be deleted to make room for those that are actively used.

Purpose	Reports the status of user accounts configured in the Active Directory server and thus, quickly points you to 'unused' accounts that can be deleted to make room for those that are actively used
Target of the test	An Active Directory or Domain Controller on Windows
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Never logged on users: Indicates the indicates the number of AD users who have never logged on to the network.	Number	A healthy AD server is one that has no or very few 'unused' user accounts. A high value is therefore not desired for this measure. To know who these users are, use the detailed diagnosis of this measure.
	Inactive users: Indicates the number of users who are currently inactive in the AD server.	Number	To identify the inactive users, use the detailed diagnosis of this measure. The detailed diagnosis displays the **Distinguished Name** of the user, **the** date/time he/she logged in last, and the date/time at which the user account was created. This will help you in figuring out how long that user has been inactive. If you think that the user will never again become active, you can proceed to delete that user account.
	Disabled users: Indicates the number of user accounts that are currently disabled on the AD server.	Number	To identify the disabled users, use the detailed diagnosis of this measure. The detailed diagnosis displays the **Distinguished Name** of the user and the date/time at which the user account was created. This will help you in figuring out how long each user account has remained disabled. If you think that the user will never again become active, you can proceed to delete that user account.

3.5.13 Account Management Events Test

The addition of new users/computers/groups to an Active Directory domain, changes to existing user/computer/group accounts, and deletion of accounts are important to verify that they were performed only by authorized personnel and with no malicious intent. To track such operations, "Audit account management events" provides specific event IDs. Using the **Account Management Events** test, you can continuously track events with the event IDs grouped under *Audit account management events*, and be proactively alerted to the sudden addition/modificiation/deletion of users/groups/computers in the Active Directory. You can also use the detailed diagnosis of the test to know which user performed the addition/modification/deletion and when.

Purpose	Continuously tracks events with the event IDs grouped under *Audit account management events*, and be proactively alerted to the sudden addition/modificiation/deletion of

	users/groups/computers in the Active Directory
Target of the test	An Active Directory server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the EventLog Service. Here it is null. 4. **SUCCESSEVENTSINDD** - By default, this parameter displays *none*, indicating that by default none of the successful log audits will be reflected in the detailed diagnosis. If you set this parameter to, say 10, then the test will display only the 10 most recent successful log audits in the detailed diagnosis page. Setting this parameter to *all*, on the other hand will make sure that all successful log audits are listed in the detailed diagnosis. 5. **FAILUREEVENTSINDD** - By default, this parameter displays *all*, indicating that by default all the failed log audits will be reflected in the detailed diagnosis. If you set this parameter to, say 10, then the test will display only the 10 most recent log audits that failed, in the detailed diagnosis page. Setting this parameter to *none*, on the other hand will make sure that none of the failed log audits are listed in the detailed diagnosis. 6. **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**. 7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs are used. This option is provided because on some Windows 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**. 8. **EVENTS DURING RESTART** - By default, the **EVENTS DURING RESTART** flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available. 9. **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.

10. **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 users 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. 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. This is the default selection.

11. **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 users 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}:{users_to_be_included}:{users_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 provide the IDs as a comma-separated list. Likewise, specifying *all* makes sure that all the event IDs are excluded from monitoring.

- In the same way, you can also ensure that events generated by specific users on the target host are alone tracked by providing a comma-separated list of users to be monitored – for example, *john,elvis*. In our example however, *all* is specified, indicating that *all* users need be monitored.
- You can similarly indicate if specific users need to be excluded from monitoring. In our example however, *none* is provided to ensure that no users are excluded from monitoring.
- By default, the **FILTER** parameter contains the value*: all:all:none:all:none:all:none*. 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 users 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}:{users_to_be_included}:{users_to_be_excluded}

To monitor a specific combination of event sources, event IDs, and users, 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 **Click here** link appears just above the test configuration section, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the **Click here** link 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.

12. **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.

Outputs of the test	One set of results for the server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	User password reset by administrator: Indicates the number of times the user password was changed by the administrator since the last measurement period.	Number	Typically, such an event occurs when the administrator attempts to change some other user's password in response to a 'forgot password' call. You can use the detailed diagnosis of this measure to know which admin user attempted the password change on which computer.
	User password reset by users: Indicates the number of times the user password was changed by the users themselves since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user attempted the password change on which computer.
	User accounts created: Indicates the number of user accounts that have been created since the last measurement period.	Number	New user accounts are important to audit to verify that they correspond to a legitimate employee, contractor or application. Outside intruders often create new user accounts to facilitate continued access to the penetrated system. Therefore, you need to eye any sudden increase in the value of this measure with suspicion. You can use the detailed diagnosis of this measure to know which user created new users on which computer.
	User accounts deleted: Indicates the number of user accounts that have been deleted since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user deleted user accounts on which computer.
	User account changed: Indicates the number of times the user account has been changed since the last measurement period.	Number	Certain changes to user accounts are important to audit since they can be a tip-off to compromised accounts. For instance, both insider and outsider computer criminals often gain access to a system by socially engineering the help desk to a user's password. Or a previously disabled account being re-enabled may be suspicious depending on the history and type of the account. You can use the detailed diagnosis of this measure to know which user made changes to user accounts on which computer.

	Computer accounts created: Indicates the number of times computer accounts have been created since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user created computer accounts on which computer.
	Computer accounts deleted: Indicates the number of computer accounts that have been deleted since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user deleted computer accounts on which computer.
	Computer accounts changed: Indicates the number of times the computer accounts that have been changed since the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user changed computer accounts on which computer.
	User/Computer object disabled: Indicates the number of times the user/computer object was disabled during the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user disabled user/computer objects on which computer.
	User/Computer object enabled: Indicates the number of times the user/computer object was enabled during the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user enabled user/computer objects on which computer.

	User added to security group: Indicates the number of users who were added to the security group during the last measurement period.	Number	Group changes, especially changes to the group's membership, are very useful to track since groups are used to control access to resources, link security policies and control wireless and remote access all over a Windows network. Security groups are the only group type that you can assign permissions and rights. Security groups are referred to as "security enabled" groups in the security log. You can use the detailed diagnosis of this measure to know which user added users to the security group on which computer.
	Security groups deleted: Indicates the number of security groups that were deleted during the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user deleted security groups on which computer.
	Security groups created: Indicates the number of security groups that were created during the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user created security groups on which computer.
	Security groups changed: Indicates the number of security groups that were changed during the last measurement period.	Number	You can use the detailed diagnosis of this measure to know which user changed security groups on which computer.

Note:

The **STATELESS ALERTING** capability is currently available for the following tests alone, by default:

- EventLog test
- ApplicationEventLog test
- SystemEventLog test
- ApplicationEvents test
- SystemEvents test
- SecurityLog test
- Account Management Events test

If need be, you can enable the **stateless alerting** capability for other tests. To achieve this, follow the steps given below:

- Login to the eG manager host.
- Edit the **eg_specs.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- Locate the test for which the **Stateless Alarms** flag has to be enabled.
- Insert the entry, **-statelessAlerts yes**, into the test specification as depicted below:

```
EventLogTest::$hostName:$portNo=$hostName, -auto, -host $hostName -port
$portNo -eventhost $hostIp -eventsrc all -excludedSrc none -useWmi yes -
statelessAlerts yes -ddFreq 1:1 -rptName $hostName, 300
```

- Finally, save the file.
- If need be, you can change the status of the **statelessAlerts** flag by reconfiguring the test in the eG administrative interface.

Once the **stateless alerting capability** is enabled for a test (as discussed above), you will find that everytime the test reports a problem, the eG manager does the following:

- Closes the alarm that pre-exists for that problem;
- Sends out a normal alert indicating the closure of the old problem;
- Opens a new alarm and assigns a new alarm ID to it;
- Sends out a fresh email alert to the configured users, intimating them of the new issue.

In a redundant manager setup, the secondary manager automatically downloads the updated **eg_specs.ini** file from the primary manager, and determines whether the stateless alerting capability has been enabled for any of the tests reporting metrics to it. If so, everytime a threshold violation is detected by such a test, the secondary manager will perform the tasks discussed above for the problem reported by that test. Similarly, the primary manager will check whether the stateless alert flag has been switched on for any of the tests reporting to it, and if so, will automatically perform the above-mentioned tasks whenever those tests report a deviation from the norm.

Note:

- Since alerts will be closed after every measurement period, alarm escalation will no longer be relevant for tests that have **statelessAlerts** set to **yes**.
- For tests with **statelessAlerts** set to **yes**, **statelessAlerts** will apply for all measurements of that test (i.e., it will not be possible to only have one of the measurements with stateless alerts and others without).
- If **statelessAlerts** is set to **yes** for a test, an alarm will be opened during one measurement period (if a threshold violation happens) and will be closed prior to the next measurement period. This way, if a threshold violation happens in successive measurement periods, there will be one alarm per measurement period. This will reflect in all the corresponding places in the eG Enterprise system. For example, multiple alerts in successive measurement periods will result in multiple trouble tickets being opened (one for each measurement period). Likewise, the alarm history will also show alarms being opened during a measurement period and closed during the next measurement period.

3.5.14 Active Directory Computers Test

This test takes stock of the total number of computers managed by the AD server and the status of these computers, so that administrators can determine from a single glance which computers are inactive/unused.

Purpose	Takes stock of the total number of computers managed by the AD server and the status of these computers, so that administrators can determine from a single glance which computers are inactive/unused		
Target of the test	An Active Directory or Domain Controller on Windows		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD -** How often should the test be executed 2. **HOST -** The IP address of the machine where the Active Directory is installed. 3. **PORT –** The port number through which the Active Directory communicates. The default port number is 389.		
Outputs of the test	One set of results for every Active Directory being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Never logged on computers:** Indicates the number of computers to which no user has ever logged in.	Number	To know which computers are unused, use the detailed diagnosis of this measure. You can consider removing such computers to reduce the workload of the AD server.
	Inactive computers: Indicates the number of computers that are currently inactive.	Number	To identify the inactive computers, use the detailed diagnosis of this measure. The detailed diagnosis displays the **Distinguished Name** of the computer, the age of the computer, and the date/time at which the computer was created. This will help you in figuring out how long that computer has been inactive. If the computer has been inactive for too long, you may think about deleting it from the AD server.
	Disabled computers: Indicates the number of computers that are currently disabled on the AD server.	Number	To identify the disabled computers, use the detailed diagnosis of this measure. The detailed diagnosis displays the **Distinguished Name** of the computer and the date/time at which the computer was created.
	Total computers: Indicates the total number of computers managed by the AD server.	Number	Use the detailed diagnosis of this measure to know the **Distinguished Names** of the computers.

3.5.15 Key Management Events Test

The Key Management Service (KMS) activates computers on a local network, eliminating the need for individual computers to connect to Microsoft. To do this, KMS uses a client–server topology. KMS client computers can locate KMS host computers by using Domain Name System (DNS) or a static configuration. KMS clients contact the KMS host by using remote procedure call (RPC). A KMS host responds to each valid activation request from a KMS client with the count of how many computers have contacted the KMS host for activation. Clients that receive a count below their activation threshold are not activated. If a computer running Windows Server 2008 or Windows Server 2008 R2 receives an activation count that is ≥5, it is activated. If a computer running Windows 7 receives an activation count ≥25, it is activated.

If users to a Windows server are having trouble logging on, administrators may want to check the *Key Management Service* event log to see if it is owing to an issue with KMS. This event log tracks events related to Kerberos key distribution, when a server functions as a key distribution center. To enable administrators to rapidly capture error/warning events captured by this event log and troubleshoot logon issues that occur, administrators can run the **Key Management Events** test. This test monitors the *Key Management Service* event log and reports the count and details of errors and warning events captured by that log.

Purpose	Monitors the *Key Management Service* event log and reports the count and details of errors and warning events captured by that log

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 EventLog Service. Here it is null. 4. **LOGTYPE** – Refers to the type of event logs to be monitored. By default, this is set to *Key Management Service*. 5. **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. 6. **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 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:all:none:all:none:all:none*. 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 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 **Click here** link appears just above the test configuration section, once the **YES** option is chosen against **POLICY BASED FILTER**. Clicking on the **Click here** link 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.

7. **USEWMI** - The eG agent can either use WMI to extract event log statistics or directly parse the event logs using event log APIs. If the **USEWMI** flag is **YES**, then WMI is used. If not, the event log APIs 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'.**

8. **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.

9. **EVENTS DURING RESTART** - By default, the **EVENTS DURING RESTART** flag is set to **Yes**. This ensures that whenever the agent is stopped and later started, the events that might have occurred during the period of non-availability of the agent are included in the number of events reported by the agent. Setting the flag to **No** ensures that the agent, when restarted, ignores the events that occurred during the time it was not available.

10. **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**.

11. **DDFORWARNING** – To ensure that the test does not generate and store detailed measures for warning events, set the **DDFORWARNING** flag to **No**.

12. **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**.

<table>
<tr><td></td><td colspan="3">13. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the FILTER configured</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Key management event information messages:

This refers to the number of information events that were captured by the Key Management Service log during the test's last execution.</td><td>Number</td><td>A change in value of this measure may indicate infrequent but successful operations.

Please check the Key Management Service log in the Event Log Viewer for more details.</td></tr>
<tr><td>Key management event warnings:

This refers to the number of warning events captured by the Key Management Service log during the test's last execution.</td><td>Number</td><td>A high value of this measure indicates problems that may not have an immediate impact, but may cause future problems.

Please check the Key Management Service log in the Event Log Viewer for more details.</td></tr>
<tr><td></td><td>Key management event errors:

This refers to the number of error events captured by the Key Management Service log during the test's last execution.</td><td>Number</td><td>A very low value (zero) is desired for this measure, as it indicates good health.

An increasing trend or a high value indicates the existence of problems.

Please check the Key Management Service log in the Event Log Viewer for more details.</td></tr>
</table>

	Key management event critical errors: Indicates the number of critical events that were generated when the test was last executed.	Number	A critical event is one that the KMS cannot automatically recover from. This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. A very low value (zero) indicates that the service is in a healthy state and is running smoothly without any potential problems. An increasing trend or high value indicates the existence of fatal/irrepairable problems. The detailed diagnosis of this measure describes all the critical events captured by the *Key Management Service* log during the last measurement period. Please check the *Key Management Service* log in the Event Log Viewer for more details.
	Distributed file system verbose messages: Indicates the number of verbose events that were generated when the test was last executed.	Number	Verbose logging provides more details in the log entry, which will enable you to troubleshoot issues better. This measure is applicable only for Windows 2008/Windows Vista/Windows 7 systems. The detailed diagnosis of this measure describes all the verbose events that were captured by the *Key Management Service* log during the last measurement period. Please check the *Key Management Service* log in the Event Log Viewer for more details.

Chapter 4

Monitoring the BizTalk Server

Microsoft BizTalk server provides a powerful web-based development and execution environment that integrates loosely coupled, long-running business processes, both within and between businesses. The server provides a standard gateway for sending and receiving documents across the Internet, as well as providing a range of services that ensure data integrity, delivery, security, and support for the BizTalk Framework and other key document formats.

As mission-critical business processes are integrated via the BizTalk server, it is imperative that the BizTalk server itself stays in good health at all times. To ensure the continuous availability and smooth functioning of the BizTalk server, you need to constantly monitor the server, and promptly detect performance issues, so that the issues can be fixed before they prove fatal to the critical business processes that ride on the server.

eG Enterprise offers two dedicated models for monitoring the BizTalk server - one each for BizTalk Server 2000 and BizTalk Server 2010. Both these models are capable of monitoring the entire pipeline of the processes happening within the BizTalk server. This chapter takes a closer look at both the models.

4.1 Monitoring the BizTalk Server 2000

BizTalk server 2000 includes a document interchange engine, a business process execution engine, a business document editor, a business document mapper, and a set of business document and server management tools. Initially, an agreement should be made between the organizations, to determine the following:

- the source and destination locations of the business documents
- the transportation medium to be used,
- the source and destination formats of the business documents

After the agreement, the business process diagram should be drawn by using the VISIO style-drawing tool. The business process diagram is then compiled to a XLANG file using XLANG Scheduler tool given by the BizTalk Server environment. The XLANG engine loads the XLANG file at runtime environment.

The sender application (say Application 1 of Organization A) is responsible for generating business documents in well-defined XML format (for e.g., a purchase order). This business document is submitted to the BizTalk server. Then, the business document has to be transformed using Schema transformations. Here, a mapping is done to transform the business document from the source organization's native representation to the representation requested by the destination organization (for e.g., the source organization may submit an XML document, but the destination organization may require the document in EDI format). The source XML document is parsed to determine the well-defined XML standard. Encoding and encryption is done when specified. Until this stage, the documents are

available in the **work queue**. Then, the document is serialized to the standard that is ready for transmission. The document in the interchange form will be available in the **scheduled queue**. By using the specified transportation medium, the document interchanges are transmitted to the destination location that has been specified in the agreement. Decryption and decoding of the business document is done at the receiving end (Application 2 of Organization B) if necessary. At this stage, the business document is in the target representation form. It is received by the target application that is running in Organization B. The business documents and interchanges will be in the **retry queue** when the BizTalk server is overloaded. In this case, the documents and interchanges are re-submitted to the BizTalk server automatically. When any error happens during the above stages, the documents and interchanges are moved to the **suspended queue** and cannot be re-submitted to the BizTalk server.

Since a BizTalk server acts as a bridge between systems having heterogeneous inputs, it is critical for the BizTalk server to perform optimally so as not to choke the performance of the system being integrated. The eG Enterprise suite of products is capable of monitoring the BizTalk server 2000 inside out. The *BizTalk* monitoring model that is used by the eG Enterprise suite for monitoring the BizTalk server is shown in Figure 4.11.

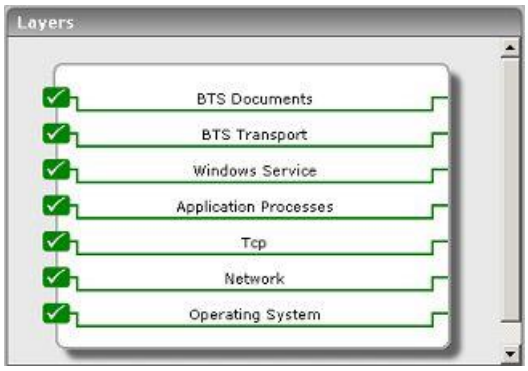

Figure 4.1: Layer model of a BizTalk server

Each layer of Figure 4.1 is mapped to tests that report a wide variety of metrics revealing the internal health of the BizTalk Server 2000. Using the metrics so reported, administrators can find quick and easy answers for many persistent performance queries, such as the following:

- Is the rate of interchange decodes and interchange decrypts unusually low?
- How is the transport mechanism functioning? Could problems in this mechanism be causing a slowdown in the reception and transmission of the interchange?
- Can the BizTalk server encode, encrypt, and serialize interchanges?
- Are applications able to receive and submit documents quickly to the BizTalk server?
- Is the BizTalk server experiencing any delays in document processing?
- Is the BizTalk server able to map documents?

The details about the 5 layers at the bottom of Figure 4.1 are available in the *Monitoring Unix and Windows Servers* document. The sections to come will therefore discuss the top 2 layers only.

4.1.1 The BTS Transport Layer

This layer monitors the transportation of the BizTalk documents and interchanges using the InterChangeRcvd test and InterChangeXmit test shown in Figure 4.2. A business document is an XML document containing the business transaction data. This transaction data may represent a purchase order, invoice, sales forecast, or any other

business information. A BizTalk document is a combination of one or more business documents, and zero or more binary file(s). BizTalk interchanges refer to a collection of one or more document instances that comprises a single transmission. This is exchanged from application to application within an organization or from one trading partner to another.

Figure 4.2: Tests mapping to the BTS Transport layer

4.1.1.1 Inter Changes Received Test

BizTalk messaging service enables the administrator to send, receive, parse, and verify the integrity of the documents, track interchanges and documents, and provide secure methods for exchanging documents with trading partners and applications. This test tracks the performance of the messaging service while receiving interchanges from the BizTalk server.

Purpose	This test measures the performance of the messaging service while receiving the interchanges from the BizTalk server.		
Target of the test	A BizTalk server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – This indicates how often should the test be executed 2. **HOST** – The IP address of the machine where BizTalk has been installed. 3. **PORT** - Not applicable to this test. Set to NULL.		
Outputs of the test	One set of results for every BizTalk server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Decode rate:** This measure indicates the number of interchanges being decoded per second by the runtime process.	Intchanges/Sec	A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. In case of an unusually low value, verify the status of the interchange in the suspended queue using the BizTalk server administration.
	Decrypt rate: This measure indicates the number of interchanges being decrypted per second by the runtime process.	Intchanges/Sec	A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. If the value of this measure is unusually low, then it indicates that the certificate might have expired. Verify the validity of the certificate in the Certificate Microsoft Management Console Snap-in.
	Receive rate: This measure indicates the number of interchanges received by the BizTalk messaging service between trading partners.	Intchanges/Sec	A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. If the value of this measure is unusually low, then it indicates that the transport mechanism (HTTP/MSMQ/FTP) used may not be functioning.

4.1.1.2 Inter Changes Transmitted Test

This test tracks the performance of the messaging service while receiving interchanges from the BizTalk server. The outputs of the test are given below:

Purpose	This test measures the performance of the messaging service while transmitting the interchanges to the BizTalk server.
Target of the test	A BizTalk server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** – This indicates how often should the test be executed 2. **HOST** – The IP address of the machine where BizTalk has been installed. 3. **PORT** - Not applicable to this test. Set to NULL.
Outputs of the test	One set of results for every BizTalk server being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Encode rate: This measure indicates the number of interchanges being encoded per second by the runtime process.	Intchanges/Sec	A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. In case of an unusually low value, verify the status of the interchange in the suspended queue using the BizTalk server administration. If the status corresponding to an interchange is **Encoding,** then it implies that the BizTalk server could not encode the interchange. Resubmitting the interchange to the BizTalk server may solve this problem.
	Encrypt rate: This measure indicates the number of interchanges being encrypted per second by the runtime process.	Intchanges/Sec	A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. In case of an unusually low value, verify the status of the interchange in the suspended queue using the BizTalk server administration. If the status corresponding to the interchange is **Encrypting** ,then it signifies that the BizTalk server could not encrypt this interchange. Also, verify the expiration of the certificate in the Certificate Microsoft Management Console snap-in.
	Serialize rate: This measure indicates the number of interchanges being serialized per second by the BizTalk runtime process.	Intchanges/Sec	A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. In case of an unusually low value, verify the status of the interchange in the suspended queue. If the status corresponding to the interchange is **Serializing**, then it implies that the BizTalk server could not convert the interchange to its native format. Resubmitting the interchange can solve this problem.

	Transmit rate: This measure indicates the number of interchanges being transmitted per second by the BizTalk messaging service.	Intchanges/Sec	A high value over a period may indicate that transmission took a long time to attain completion. A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. If the value of this measure is unusually low, verify the transport address in the channel. Correct the problem in the channel and resubmit the interchange. Alternatively, the BizTalk server might have taken a long time to transmit the interchange. Verify the transport mechanism used. Another reason could be that the BizTalk administrator might have moved the interchange to the suspended queue, resubmitted the interchange from the suspended queue. Alternatively, the computer on which the BizTalk server could be running out of memory, restart the server and resubmit all the interchanges in the suspended queue.

4.1.2 The BTS Documents Layer

This layer reports the statistics about the various attributes of the documents being handled by the BizTalk server using the DocReceive test and DocSubmit test shown in Figure 4.3.

Figure 4.3: Tests mapping to the BTS Documents layer

4.1.2.1 Documents Received Test

This test tracks the performance of the messaging service while it is receiving documents from the BizTalk server.

Purpose	This test measures the performance of the messaging service while receiving documents from the

	BizTalk server.		
Target of the test	A BizTalk server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – This indicates how often should the test be executed 2. **HOST** – The IP address of the machine where BizTalk has been installed. 3. **PORT** - Not applicable to this test. Set to NULL.		
Outputs of the test	One set of results for every BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Receive rate: This measure indicates the number of documents being received per second by the application from the BizTalk server.	Docs/Sec	A value of –1 for this measure indicates that either the BizTalk messaging service or XLANG Schedule Restart Service or Distributed Transaction Coordinator (MSDTC) may not be running. If the value of this measure is unusually low, then verify the interface between the BizTalk server and the application.

4.1.2.2 Documents Submitted Test

The DocSubmitTest tracks the performance of the messaging service while submitting documents to the BizTalk server.

Purpose	This test measures the performance of the messaging service while submitting documents to the BizTalk server.		
Target of the test	A BizTalk server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – This indicates how often should the test be executed 2. **HOST** – The IP address of the machine where BizTalk has been installed. 3. **PORT** - Not applicable to this test. Set to NULL.		
Outputs of the test	One set of results for every BizTalk server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Submit rate:** This measure shows the number of documents submitted asynchronously per second to the BizTalk server from the application. Once submitted, the BizTalk server holds the documents in the work queue for further processing.	Docs/Sec	A sudden increase in the value of this measure denotes a change in the workload. A value of –1 for this measure indicates that either the BizTalk messaging service or XLANG Schedule Restart Service or Distributed Transaction Coordinator (MSDTC) may not be running. If the value of this measure is unusually low, then verify the interface between the BizTalk server and the application or check the event log entry in the BizTalk server administration.
	Map rate: BizTalk runtime process maps the actual document content from one structural form to another. This measure shows the number of documents that have been mapped per second by the runtime process.	Docs/Sec	A sudden increase in the value of this measure indicates that the BizTalk runtime process is mapping larger number of documents. This scenario indicates an increased workload. A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. For other reasons, verify the status of the document available in the suspended queue using the BizTalk server administration. If the status corresponding to a document is **Mapping** then it indicates that the document has been failed to map. To rectify this problem, delete the document from the suspended queue, correct the map and resubmit the document.
	Parse rate: This measure shows the number of documents in the work queue that is being parsed per second by the appropriate parser.	Docs/Sec	A sudden increase in the value of this measure indicates that the parser is parsing larger number of documents. This scenario may be due to the deletion of large number of documents from the suspended queue, which affects the performance of the parser. A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. Incase of an unusually low value, verify the status of the documents available in the suspended queue using the BizTalk server administration. If the status corresponding to the document is **parsing** then it indicates that the BizTalk server was unable to parse the data. The other reasons could be that the timestamp of the document is no longer valid, or the document does not contain enough information to locate the channel.

	Process rate: This measure indicates the number of documents being processed successfully (necessary changes to the document) per second by the runtime process.	Docs/Sec	A high value for this measure over a period may indicate that the runtime system is processing larger number of documents. This scenario may indicate a change in the workload. A value of –1 for this measure implies that the BizTalk messaging service or Distributed Transaction Coordinator (MSDTC) or XLANG Schedule Restart Service may not be running. Incase of an unusually low value, verify the status of Microsoft SQL server in the Service Manager tool available in the Microsoft SQL server environment. Also check the status of the document available in the suspended queue using the BizTalk server administration.

4.2 Monitoring the BizTalk Server 2010

BizTalk Server is Microsoft's Integration and connectivity server solution. A mature product on its seventh release, BizTalk Server 2010 provides a solution that allows organizations to more easily connect disparate systems. Including over 25 multi-platform adapters and a robust messaging infrastructure, BizTalk Server provides connectivity between core systems both inside and outside your organization. In addition to integration functionality, BizTalk also provides strong durable messaging, a rules engine, EDI connectivity, Business Activity Monitoring (BAM), RFID capabilities and IBM Host/Mainframe connectivity.

The BizTalk Server includes a range of technologies. The figure below illustrates the product's major components.

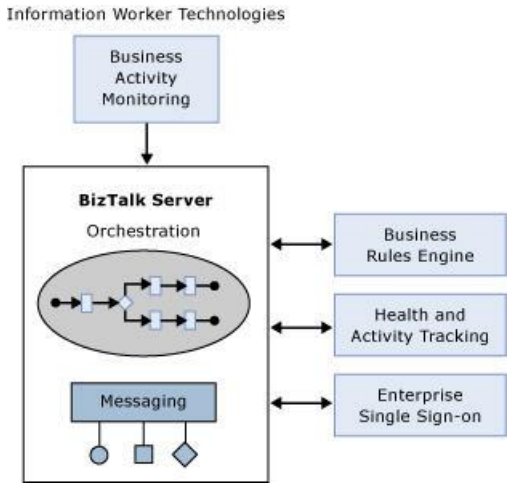

Figure 4.4: The major components of a BizTalk server

As the figure suggests, the heart of the product is the BizTalk Server Engine. The engine has two main parts:

- A messaging component that provides the ability to communicate with a range of other software. By relying on adapters for different kinds of communication, the engine can support a variety of protocols and data formats, including Web services and many others.
- Support for creating and running graphically-defined processes called orchestrations. Built on top of the engine's messaging components, orchestrations implement the logic that drives all or part of a business

process.

Several other BizTalk components can also be used in concert with the engine, including:

- A Business Rule Engine that evaluates complex sets of rules.
- A Group Hub that lets developers and administrators monitor and manage the engine and the orchestrations it runs.
- An Enterprise Single Sign-On (SSO) facility that provides the ability to map authentication information between Windows and non-Windows systems.

On top of this foundation, BizTalk Server includes Business Activity Monitoring, which information workers use to monitor a running business process. The information is displayed in business rather than technical terms, and business users determine what information is displayed.

As the present era is all about business process management, the BizTalk server plays a vital role in connecting and communicating with disparate business processes that may be operating within an organization or across organizations. If this 'connector' malfunctions, it could break the only link that exists between the processes, thereby significantly affecting the way the enterprise functions. All software-dependent activities of the enterprise - from the performance of simple, routine operations to the execution of critical business transactions - could either experience delays or could come to a virtual standstill. If such adversities are to be avoided, the BizTalk server has to be monitored 24x7.

eG Enterprise provides a *BizTalk 2010* monitoring model that provides in-depth monitoring of the BizTalk Server 2010. Each layer of this model is mapped to a series of tests that report issues in the overall health of the adapters and protocols supported by the BizTalk server, thus shedding light on applications with which the server is unable to communicate.

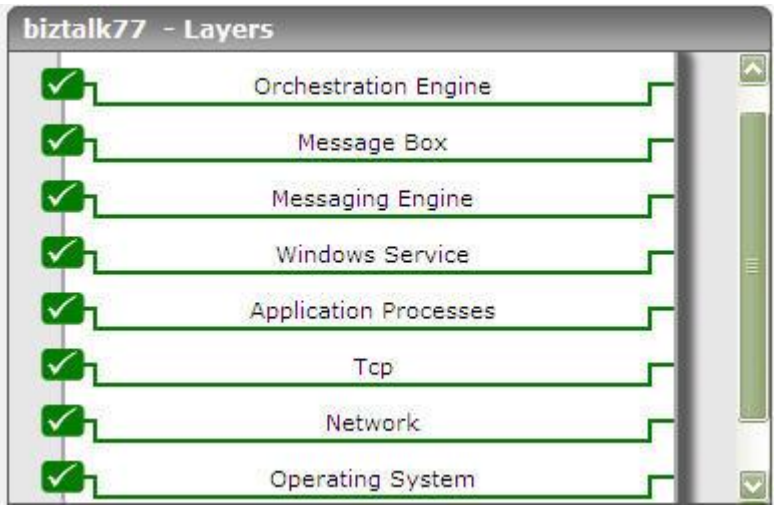

Figure 4.5: The layer model of the BizTalk Server 2010

The metrics extracted by these tests enable administrators to find answers to persistent performance queries such as the following:

- Which host instance is heavily loaded in terms of documents processed?
- Is any host instance experiencing processing bottlenecks?
- Have any documents been suspended by a host instance? If so, which host instance is it?
- Have any request messages timed out without response messages?

- How are the receive and send adapters on a host instance handling the load? Is any receive/send adapter experiencing a slowdown in processing? Which adapter is it - the file adapter, FTP adapter, HTTP adapter, Msmq adapter, POP3 adapter, SMTP adapter, SOAP adapter, or the SQL adapter?
- Is the messaging engine experiencing any latencies - if so, where did the delay originate? while delivering messages to the MessageBox, or while delivering messages to a target application?
- Are too many messages pending processing in the host queue?
- Are any SQL agent jobs taking too long to complete? If so, which ones?
- Is the depth of the spool table optimal, or is it growing continuously?
- Is the tracking data table growing uncontrollably in size?
- Have too many orchestrations been suspended or discarded?
- What is the rate at which dehydrations and rehydrations take place?
- Have the orchestrations acknowledged all the messages they received, or are there too many pending messages?
- Is there a contention for physical memory resources on any host instance?
- Have any BAM (Business Activity Monitoring) events failed?
- Has the tracking data decode service failed to process any batches?
- How is the host throttling mechanism functioning? Are message processing and/or message publishing throttled? Were any delays imposed on the message processing/publishing rates?
- Has process memory consumption exceeded its threshold?
- Has thread count exceeded its threshold?

The sections that follow will discuss each layer of Figure 4.5 in great detail.

4.2.1 The Messaging Engine Layer

The BizTalk Server Messaging engine enables users to create business processes that spans multiple applications by providing two primary things:

- A way to specify and implement the logic driving that business process
- A mechanism for communicating across the applications that the business process uses

The figure below illustrates the main components of the engine that address these two problems.

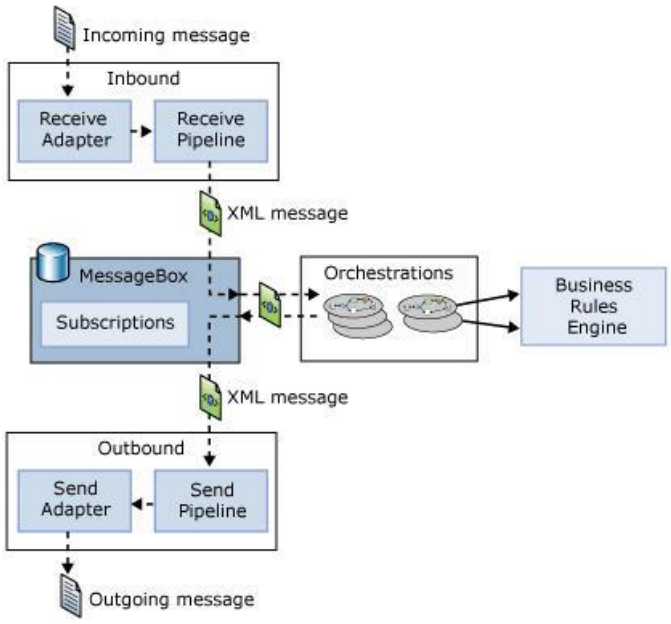

Figure 4.6: Messaging architecture

As the diagram shows, a message is received through a **receive adapter**. Different adapters provide different communication mechanisms, so a message might be acquired by accessing a Web service, reading from a file, or in some other way. The message is then processed through a **receive pipeline**. This pipeline can contain various components that do things such as converting the message from its native format into an XML document, validating a message's digital signature, and more. The message is then delivered into a database called the **MessageBox**, which is implemented using Microsoft SQL Server.

The logic that drives a business process is implemented as one or more **orchestrations**, each of which consists of executable code. These orchestrations are not created by writing code in a language such as C#, however. Instead, a business analyst or (more likely) a developer uses an appropriate tool to graphically organize a defined group of shapes to express conditions, loops, and other behavior. Orchestrations can optionally use the **Business Rule Engine**, which provides a simpler and more easily modified way to express complex sets of rules in a business process.

Each orchestration creates **subscriptions** to indicate the kinds of messages it wants to receive. When an appropriate message arrives in the MessageBox, that message is dispatched to its target orchestration, which takes whatever action the business process requires. The result of this processing is typically another message, produced by the orchestration and saved in the MessageBox. This message, in turn, is processed by a **send pipeline**, which may convert it from the internal XML format used by BizTalk Server to the format required by its destination, add a digital signature, and more. The message is then sent out using a **send adapter**, which uses an appropriate mechanism to communicate with the application for which this message is destined.

This layer monitors the messaging engine of the BizTalk server, measures the load on the engine, reports how quickly every send and receive adapter processes the message load, and sheds light on current / potential processing bottlenecks (if any) in the engine. All the tests mapped to this layer report metrics for each host instance on the BizTalk server. A *host* is a logical representation of a Microsoft Windows process that executes BizTalk Server artifacts such as send ports and orchestrations. A *host instance* is the physical representation of a host on a specific server.

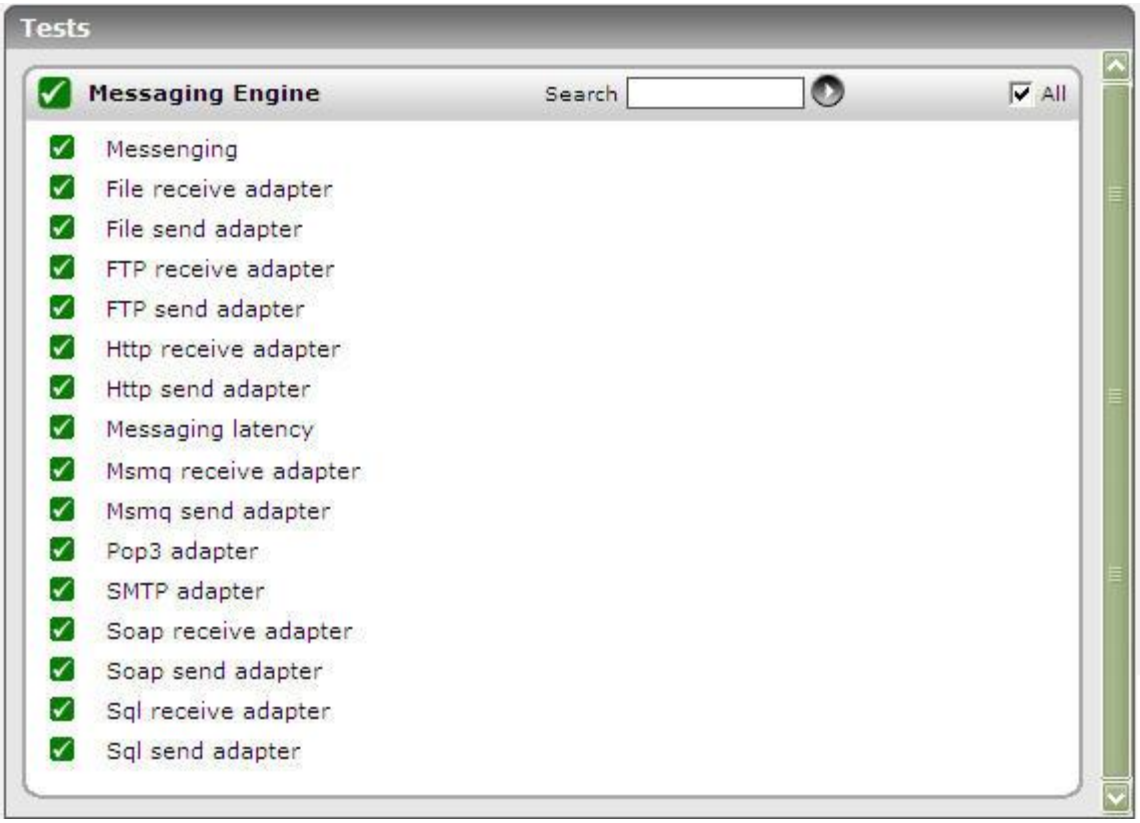

Figure 4.7: The tests mapped to the Messaging Engine layer

4.2.1.1 BT Messenging Test

This test monitors the documents received and sent by each host instance on the BizTalk server, and reports the load on that host instance and delays experienced by the host instance while processing the documents.

Using this test, administrators can easily isolate host instances that are overloaded or are experiencing bottlenecks in processing.

Purpose	Monitors the documents received and sent by each host instance on the BizTalk server, and reports the load on that host instance and delays experienced by the host instance while processing the documents
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.

Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active receive locations: Indicates the number of receive locations currently enabled in this host instance.	Number	A *receive location* is the configuration of a single endpoint (URL) to receive messages.
	Documents processed: Indicates the number of documents processed by this host instance.	Number	This is a good indicator of the load handled by a host instance. Comparing the value of this measure across host instances will reveal which instance is currently overloaded.
	Documents processed: Indicates the rate at which this host instance processed documents.	Docs/Sec	A very low value or a consistent decrease in the value of this measure indicates a slowdown in the corresponding host instance.
	Documents received: Indicates the number of documents received by this host instance from a target source.	Number	This is a good indicator of the load handled by a host instance. Comparing the value of this measure across host instances will reveal which instance is currently overloaded.
	Documents received: Indicates the rate at which documents were received by this host instance.	Number	

	Documents suspended: Indicates the number of documents that have been suspended by this host instance.	Number	By default, the BizTalk server places failed messages/documents in the Suspended queue. The value of this measure indicates the number of documents in the Suspended queue. A message failure can occur in one of the following instances: • **Failures in the disassembly phase:** Processing might also fail during the disassembly phase; that is, failure in one of the pipeline components. For example, decryption failed due to absence of decryption cert on the processing server, or parsing failure due to problem either in the schema or in the message. • **Failures in routing:** After a message disassembles successfully, the next potential failure point is routing; for example, users enable a corresponding receive location of an orchestration and forget to enlist the orchestration. In this case, the message picked up from the receive location fails routing and the MessageBox database generates a Routing Failure report. Routing Failure reports are listed in the BizTalk Server Administration Console as non-resumable suspended messages. Each Routing Failure report contains a message property snap shot taken when the routing failure occurred. You can use the information in each report to determine why routing failed for its associated message. If the associated message is resumable, you can correct the routing problem and resume the message so that processing continues.

			• **Failures during the transformation phase:** When a message is received from Receive Location, the message is disassembled (for example, decrypted and parsed), the message might optionally be transformed to a different format via an Inbound Map specified on a receive Port, and published to the MessageBox for routing to an orchestration or a Send Port. In this case, processing may fail during transformation phase due to incorrect Inbound Map, or problems in the schema or in the message received. When a message is to be sent to a Send Location, an Outbound Map configured on Send Port might optionally transform the message. Then the transformed message is assembled and handed to the adapter for final transmission to the Send Location. In this case, processing may fail during transformation phase due to incorrect Outbound Map or problem in schema or source message. • **Failures in the message assembly phase:** Processing can also fail during message assembly phase – in other words, failing in pipeline component. After a message successfully assembles, the next potential failure point becomes transmission to Send Location; for example, the Send Location (which belongs to the partner) might be down or not exist.
	Documents suspended: Indicates the rate at which documents were suspended by this host instance.	Docs/Sec	

	Request/Response timeouts: Indicates the number of request messages that have not received a response message within the time limit specified by the adapter associated with this host instance.	Number	A high value of this measure could indicate that too many messages are getting timed out. You may want to consider reconfiguring the timeout period.

4.2.1.2 BT Messaging Agents Test

Most of the processing that takes place on a BizTalk server occurs within a logical entity known as a BizTalk Server host instance, which is a process running as a Windows service or an isolated host process on the BizTalk server. To manage the use of resources by a host instance process, BizTalk Server utilizes an adjustable throttling mechanism that governs the flow and processing of messages through a host instance.

The throttling mechanism moderates the workload of the host instance to ensure that the workload does not exceed the capacity of the host instance or any downstream host instances. The throttling mechanism also prevents a condition known as resource contention that can lower the overall performance of the host instance process or other system processes. Resource contention occurs when one or more processes consume a limited resource to the detriment of themselves and/or another process. For example, the consumption of excessive memory or threads can lead to memory allocation failure or high thread context-switches, which can impact the performance of the process. Resource contention like this can be detrimental to the overall performance of BizTalk Server.

The host throttling mechanism also detects when available resources are being underutilized. If available resources are underutilized then the throttling mechanism allows additional messages to be processed by a host instance. The host throttling mechanism continually monitors if available resources are being over or underutilized and adjusts message flow through the host instance accordingly.

The BizTalk Server host throttling mechanism helps to ensure that the system operates at an optimal and sustainable level.

This test measures the efficiency of the host throttling mechanism.

Purpose	Measures the efficiency of the host throttling mechanism
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Publishing delay: Indicates the current delay imposed on each message publishing batch.	MilliSec	This measure is applicable only if the message publishing is throttled and if the message publishing batch is not exempted from throttling. *Message publishing throttling* in BizTalk Server, is applied to host instances that contain receive adapters or orchestrations that publish messages to the MessageBox database. An inbound host throttling condition can be triggered under the following conditions: • The amount of memory, the number of threads, or the number of database connections used by the host instance exceeds the throttling thresholds defined • Downstream hosts are unable to process the messages that are published. • The **Message publishing incoming rate** for the host instance exceeds the **Message publishing outgoing rate *** the specified **Rate overdrive factor (percent)** value. • The default throttling behavior has been modified by setting a registry value or values to control the throttling behavior of a host process. Depending on the severity of the throttling condition, the following actions are taken: • A progressive delay in the processing logic of the host instance is implemented. The delay may be implemented when an End Point Manager (EPM) thread receives a batch of messages from the transport adapter, and/or when the EPM submits a batch of messages to be published into the MessageBox database. Both the duration of the processing delay and the rate at which the duration is incremented scale with the severity of the throttling condition.

			• The number of threads that are available to the End Point Manager (EPM) is restricted. The EPM receives batches of messages from adapters and publishes the messages to the MessageBox database. By default, the EPM is configured to use 20 threads per CPU. If the host throttling mechanism detects a stress condition for inbound processing then it can temporarily reduce the number of threads available to the EPM until the stress condition is eliminated. The EPM cannot process messages from transport adapters or deliver message batches to the MessageBox database unless an EPM thread is available to service the inbound message batch. • The use of memory and other resources is reduced as applicable. BizTalk Server can send instructions to other service classes to limit memory use by dehydrating running schedules, shrinking memory cache size, and by limiting the usage of memory-intensive threads.
	Publishing incoming rate: Indicates the rate at which the messages are being sent by the message agent to the database of this host instance for publishing.	Msgs/Sec	A *message publishing throttling condition* is also triggered when the **Message publishing incoming rate** for the host instance exceeds the **Message publishing outgoing rate** * the specified **Rate overdrive factor (percent)** value. The **Rate overdrive factor (percent)** value is defined on the **Message Publishing Throttling Settings** dialog box available from the **Advanced** page of the **Host Properties** dialog box.
	Publishing outgoing rate: Indicates the rate at which the messages are actually published by the message agent in the database of this host instance.	Msgs/Sec	

<table>
<tr>
<td></td>
<td>Publishing throttling state:

Indicates whether the system is throttling the message publishing i.e., indicates whether the XLANG message processing and inbound transports are affected.</td>
<td>Number</td>
<td>This measure indicates any one of the following values while indicating whether the system is throttling the message publishing or not.
<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Not throttling</td></tr>
<tr><td>2</td><td>Throttling due to imbalanced message publishing rate (input rate exceeds output rate)</td></tr>
<tr><td>4</td><td>Throttling due to process memory pressure</td></tr>
<tr><td>5</td><td>Throttling due to system memory pressure</td></tr>
<tr><td>6</td><td>Throttling due to database growth</td></tr>
<tr><td>8</td><td>Throttling due to high session count</td></tr>
<tr><td>9</td><td>Throttling due to high thread count</td></tr>
<tr><td>11</td><td>Throttling due to user override on publishing</td></tr>
</table>
</td>
</tr>
</table>

	Delivery delay: Indicates the current delay imposed on each message delivery batch.	MilliSec	This measure is applicable only if message delivery is throttled. *Message processing throttling* in BizTalk Server, is applied to host instances that contain orchestrations or send adapters that receive and deliver or process messages that have been published to the MessageBox. An outbound host throttling condition can be triggered under the following conditions: • The amount of memory, the number of threads, or the number of database connections used by the host instance exceeds the throttling thresholds defined • The **Message delivery incoming rate** for the host instance exceeds the **Message delivery outgoing rate * the specified Rate overdrive factor (percent) value**. • The number of messages being processed concurrently by the host instance exceeds the **In-process messages per CPU * the number of CPUs available on the box**. • The default throttling behavior has been modified by setting a registry value or values to control the throttling behavior of a host process. Depending upon the severity of the throttling condition, the following actions are taken: • A progressive delay in the processing logic of the host instance is implemented before delivering the messages to the outbound transport adapter or the orchestration engine for processing the messages. Both the duration of the processing logic delay and the rate at which the duration is incremented scale with the severity of the throttling condition.

			• The number of messages that can be held by the in-memory queue is limited. The in-memory queue serves as a temporary placeholder for delivering messages from the MessageBox to the Message Agent which in turn delivers messages to XLANG and send adapters. By default, the in-memory queue is set to hold 100 messages per CPU. When the queue is full, no more messages are de-queued from the MessageBox until the in-memory queue is freed up. • The size of the Message Agent thread pool is limited. By limiting the Message Agent thread pool size, the host throttling mechanism effectively reduces the amount of messages that are delivered to XLANG and adapters. • The use of memory and other resources is reduced as applicable. BizTalk Server can send instructions to other service classes to limit memory use by dehydrating running schedules, shrinking memory cache size, and by limiting the usage of memory intensive threads.
	Delivery incoming rate: Indicates the rate at which the messages are delivered to the Orchestration engine or the Messaging engine of this host instance.	Msgs/Sec	A *message processing throttling* condition can also be triggered if the message **Delivery incoming rate** for the host instance exceeds the message **Delivery outgoing rate** * the specified **Rate overdrive factor (percent)** value. The **Rate overdrive factor (percent)** value is defined on the **Message Processing Throttling Settings** dialog box available from the **Advanced** page of the **Host Properties** dialog box.
	Delivery outgoing rate: Indicates the rate at which the messages are processed and sent to the recipients by the Orchestration engine or the Messaging engine of this host instance.	Msgs/Sec	

<table>
<tr>
<td></td>
<td>Delivery throttling state:

Indicates whether the system is throttling the message delivery i.e., indicates whether the XLANG message processing and outbound transports are affected or not.</td>
<td>Number</td>
<td>Indicates whether the system is throttling the message delivery i.e., indicates whether the XLANG message processing and outbound transports are affected or not.
<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Not throttling</td></tr>
<tr><td>1</td><td>Throttling due to imbalanced message delivery rate (input rate exceeds output rate)</td></tr>
<tr><td>3</td><td>Throttling due to high in-process message count</td></tr>
<tr><td>4</td><td>Throttling due to process memory pressure</td></tr>
<tr><td>5</td><td>Throttling due to process memory pressure</td></tr>
<tr><td>9</td><td>Throttling due to high thread count</td></tr>
<tr><td>10</td><td>Throttling due to user override on delivery</td></tr>
</table>
</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>High database session:

Indicates whether the database session is within normal limits or not for this host instance.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the database session is within normal limits or not.

<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>Database session count exceeds threshold</td></tr>
</table>

The database session count is nothing but the number of concurrent MessageBox database connections being used. The threshold for database session count is initially set to the value specified for Database connections per CPU on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. This value is auto-tuned based on the database session usage of the process. If the number of concurrent database sessions exceeds this threshold at any time, host throttling is implemented.</td>
</tr>
<tr>
<td></td>
<td>High database size:

Indicates whether the size of the database is within normal limits or not for this host instance.</td>
<td>Number</td>
<td>This measure indicates any one of the following values while indicating whether the size of the database is within normal limits or not.

<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>Database size has grown beyond threshold</td></tr>
</table>

Database size is represented by the number of messages in the database queues that a host instance has published. This value is measured by the number of items in the queue tables for all hosts and the number of items in the spool and tracking tables. If a process is publishing to multiple queues, this counter reflects the weighted average of all the queues. If the threshold set for database size is violated, then throttling is implemented.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>High in-process message count:

Indicates whether the In-process message count is within normal limits or not.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the In-process message count is within normal limits or not.

<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>In-process message count exceeds limit</td></tr>
</table>

The in-process message count indicates the number of in-memory messages delivered to the XLANG engine or the outbound messaging engine that are not yet processed.</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>High message delivery rate:

Indicates whether the message delivery rate is within normal limits or not.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the message delivery rate is within normal limits or not.

<table><tr><th>Value</th><th>State</th></tr><tr><td>0</td><td>Normal</td></tr><tr><td>1</td><td>Message delivery rate exceeds the message processing rate</td></tr></table></td>
</tr>
<tr>
<td></td>
<td>High message publishing rate:

Indicates whether the message publishing rate is within normal limits or not.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the message publishing rate is within normal limits or not.

<table><tr><th>Value</th><th>State</th></tr><tr><td>0</td><td>Normal</td></tr><tr><td>1</td><td>Message delivery rate exceeds the message processing rate</td></tr></table></td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>High process memory:
Indicates whether the process memory is within normal limits or not.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the process memory is within normal limits or not.

<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>Process memory exceeds threshold</td></tr>
</table>

Process memory consumption is the maximum of the process's working set size and the total space allocated for the page file for the process. The threshold for process memory consumption s initially set to the value specified for Process memory usage on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. If a percentage value is specified, it is computed based on the available memory to commit.</td>
</tr>
<tr>
<td></td>
<td>High system memory:
Indicates whether the system memory is within normal limits or not.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the system memory is within normal limits or not.
<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>System memory exceeds threshold</td></tr>
</table>
</td>
</tr>
</table>

<table>
<tr>
<td></td>
<td>High thread count:
Indicates whether the thread count is within normal limits or not for this host instance.</td>
<td>Number</td>
<td>This measure reports any one of the following values to indicate whether the thread count is within normal limits or not.

<table>
<tr><th>Value</th><th>State</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>Thread count exceeds threshold</td></tr>
</table>

The thread count indicates the number of threads being used in the process. The threshold for this count is initially set to the value specified for Threads per CPU on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. This value is auto-tuned depending on the thread requirements of the current process. If the number of threads in the process exceeds this threshold value at any point in time, host throttling is implemented.</td>
</tr>
<tr>
<td></td>
<td>Thread count:
Indicates the number of thread being used in the process.</td>
<td>Number</td>
<td></td>
</tr>
<tr>
<td></td>
<td>Thread count threshold:
Indicates the current threshold for the number of threads in the process.</td>
<td>Number</td>
<td>The threshold for the thread count is initially set to the value specified for Threads per CPU on the Throttling Thresholds dialog available from the Advanced page of the Host Properties dialog box. This value is auto-tuned depending on the thread requirements of the current process. If the number of threads in the process exceeds this threshold value at any point in time, host throttling is implemented.</td>
</tr>
<tr>
<td></td>
<td>Database size:
Indicates the number of messages in the database queues that this process has published.</td>
<td>Number</td>
<td>This value is measured by the number of items in the queue tables for all hosts and the number of items in the spool and tracking tables. If a process is publishing to multiple queues, this counter reflects the weighted average of all the queues.</td>
</tr>
</table>

	Database session: Indicates the number of concurrent message box database connections that is being used.	Number	
	Process memory usage: Indicates the memory consumption of the process.	MB	
	Process memory usage threshold: Indicates the current threshold for the memory consumption of the process.	MB	This threshold value is initially set to the value specified for the process memory consumption on the **Throttling Thresholds** dialog available from the **Advanced page** of the **Host Properties** dialog box. If a percentage value is specified, the threshold value is computed based on the available memory to commit.

4.2.1.3 BT File Receive Adapter Test

The file receive adapter is used to read messages from files and submit them to the server. The receive adapter reads the file and creates a BizTalk message object, so that BizTalk server can process the message. While reading from the file, the adapter locks the file to ensure that no modifications can be made to the file content. The file receive adapter does **not** pick up read-only files or system files.

This test reports how efficient the file receive adapter on each host instance is. The test monitors the inflow of messages to the file receive adapter, measures the load on the adapter, and reveals how well the adapter handled the load; lock failures encountered by the adapter while attempting to read files are also revealed by this test, so that reasons for the same can be diagnosed.

Purpose	This test reports how efficient the file receive adapter on each host instance is. The test monitors the inflow of messages to the file receive adapter, measures the load on the adapter, and reveals how well the adapter handled the load; lock failures encountered by the adapter while attempting to read files are also revealed by this test, so that reasons for the same can be diagnosed.
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes received: Indicates the total number of bytes received by the file receive adapter on this host instance.	Bytes	The counter is incremented after a message is completely read by the adapter from the file system.
	Bytes received: Indicates the rate at which bytes were received by the file receive adapter on this host instance.	Bytes/Sec	
	Messages received: Indicates the number of messages received by the file receive adapter on this host instance.	Number	The counter is incremented after a message is completely read by the file receive adapter from the file system.
	Messages received: Indicates the rate at which messages were received by the file receive adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have been completely read by the file receive adapter from the file system. Ideally, the value of this measure should be high. A low value indicates that the file receive adapter is not reading files quickly. Further investigation may be required to diagnose the root-cause of the slowdown.
	Lock failures: Indicates the number of times the file receive adapter on this host instance failed to lock the file.	Number	Ideally, the value of this measure should be 0. A non-zero value indicates a lock failure. This in turn implies that the adapter could not prevent changes from being made to one/more files that were being read.

4.2.1.4 BT File Send Adapter Test

The File send adapter transmits messages from the message box database to a specified destination address (URL). You define the URL, which is a file path and file name, by using wildcard characters related to the message context properties. The File send adapter resolves the wildcard characters to the actual file name before writing the message to the file.

When writing a message to a file, the File send adapter gets the message content from the body part of the BizTalk message object. The File send adapter ignores other message parts in the BizTalk Message object. After the File adapter writes the message to a file, it deletes the message from the MessageBox database. The File adapter writes files to the file system either directly or by using the file system cache, which can improve performance, particularly for large files.

This test monitors the outflow of data and messages from the file send adapter on each host instance and reports the load on the adapter and the slowdowns (if any) suffered by the adapter while processing the load.

Purpose	Monitors the outflow of data and messages from the file send adapter on each host instance and reports the load on the adapter and the slowdowns (if any) suffered by the adapter while processing the load		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes sent: Indicates the total number of bytes sent by the file send adapter on this host instance.	Bytes	The counter is incremented only for messages that have been completely written to file system.
	Bytes sent: Indicates the rate at which bytes were sent by the file send adapter on this host instance.	Bytes/Sec	The counter applies only to messages that have been completely written to file system.

	Messages sent: Indicates the number of messages sent by the file send adapter on this host instance.	Number	The counter is incremented only for messages that have been completely written to file system.
	Messages sent: Indicates the rate at which messages were sent by the file send adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have been completely written to file system. Ideally, the value of this measure should be high. A low value indicates that the file send adapter is experiencing delays while writing files to the file system. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.5 BT FTP Receive Adapter Test

The FTP receive adapter enables you to move data from an FTP server to BizTalk Server.

Key features of the FTP receive adapter are:

- Pulling files from the FTP server on demand
- Running polls based on a configurable schedule
- Polling the FTP server and sending data directly to BizTalk Server
- Specifying the FTP server as an IP address, port, password, and host name
- Guaranteed file delivery

The FTP receive adapter also works with the BizTalk Administration console and BizTalk Explorer to configure and administer each receive function, which is composed of the following configuration items:

- Poll interval to run an FTP command (for example, 60 minutes).
- Information with which to route the document to a specific BizTalk send port or receive location.

The FTP receive adapter does **not** support receiving files from a **partitioned data set**.

With the help of this test, you can measure the current load on the FTP receive adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing.

Purpose	Measure the current load on the FTP receive adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes received: Indicates the total number of bytes received by the FTP receive adapter on this host instance.	Bytes	The counter is incremented after a message is completely read by the FTP receive adapter from the FTP server.
	Bytes received: Indicates the rate at which bytes were received by the FTP receive adapter on this host instance.	Bytes/Sec	The counter applies only to messages that have been completely read by the FTP receive adapter from the FTP server.
	Messages received: Indicates the number of messages received by the FTP receive adapter on this host instance.	Number	The counter applies only to messages that have been completely read by the FTP receive adapter from the FTP server. This measure is a good indicator of the load on the FTP receive adapter.
	Messages received: Indicates the rate at which messages were received by the FTP receive adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have been completely read by the FTP receive adapter from the FTP server. Ideally, a value of this measure should be high. A low value indicates that the FTP receive adapter is experiencing delays while moving files and data from the FTP server to the BizTalk server. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.6 BT FTP Send Adapter Test

The FTP send adapter enables you to move data from BizTalk Server to an FTP server.

Key features of the FTP send adapter are:

- Ability to run sends on demand

- Guaranteed delivery

With the help of this test, you can measure the current load on the FTP send adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing.

Purpose	Measure the current load on the FTP send adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes sent: Indicates the total number of bytes sent by the FTP send adapter on this host instance.	Bytes	The counter is incremented only for messages that have been written to the destination FTP server.
	Bytes sent: Indicates the rate at which bytes were sent by the FTP send adapter on this host instance.	Bytes/Sec	The counter applies only to messages that have been written to the destination FTP server.
	Messages sent: Indicates the number of messages sent by the FTP send adapter on this host instance.	Number	The counter is incremented only for messages that have been written to the destination FTP server.

	Messages sent: Indicates the rate at which messages were sent by the FTP send adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have been written to destination FTP server. Ideally, a value of this measure should be high. A low value indicates that the FTP send adapter is experiencing delays while writing files to the destination FTP server. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.7 BT Http Receive Adapter Test

The HTTP adapter is used to exchange information between the BizTalk server and an application by means of the HTTP protocol. HTTP is the primary protocol for inter-business message exchange. Applications can send messages to a server by sending HTTP POST or HTTP GET requests to a specified HTTP URL. The HTTP receive adapter is an Internet Information Services (IIS) Internet Server Application Programming Interface (ISAPI) extension that the IIS process hosts, and controls the receive locations that use the HTTP adapter. The receive location for the HTTP receive adapter is a distinct URL configured through BizTalk Explorer.

Using this test, you can monitor the flow of messages to and from the HTTP receive adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP receive adapter of a host instance, and evaluate the load processing ability of that adapter.

Purpose	Monitor the flow of messages to and from the HTTP receive adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP receive adapter of a host instance, and evaluate the load processing ability of that adapter		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Messages received: Indicates the total number of HTTP requests received by the HTTP receive adapter on this host instance.	Number	The counter is incremented after a request message is completely read by the HTTP receive adapter from the HTTP client.
	Messages received: Indicates the rate at which the HTTP requests are received by the HTTP receive adapter on this host instance.	Msgs/Sec	The counter applies only to request messages that have been completely read by the HTTP receive adapter from the HTTP client. Ideally, the value of this measure should be high. A low value indicates that the HTTP receive adapter is experiencing delays while accepting requests from the HTTP client. Further investigation may be required to diagnose the root-cause of the slowdown.
	Messages sent: Indicates the total number of HTTP responses sent by the HTTP receive adapter on this host instance.	Number	The counter is incremented only for response messages that have been successfully sent to HTTP clients.
	Messages sent: Indicates the rate at which messages were sent by the FTP send adapter on this host instance on this host instance.	Number	The counter applies only to response messages that have been successfully sent to HTTP clients. Ideally, the value of this measure should be high. A low value indicates that the HTTP receive adapter is experiencing delays while responding to requests from HTTP clients. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.8 BT Http Send Adapter Test

The HTTP send adapter gets messages from BizTalk Server and sends them to a destination URL on an HTTP POST request. The HTTP send adapter gets the message content from the body part of the BizTalk Message object. The HTTP send adapter ignores all other parts of the BizTalk Message object.

Using this test, you can monitor the flow of messages to and from the HTTP send adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP send adapter of a host instance, and evaluate the load processing ability of that adapter.

Purpose	Monitor the flow of messages to and from the HTTP send adapter for each host instance on the BizTalk server. In the process, you can determine the current workload of the HTTP send adapter of a host instance, and evaluate the load processing ability of that adapter
Target of the test	A BizTalk Server 2010
Agent	An internal agent

deploying the test			
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages received: Indicates the total number of HTTP response messages received by the HTTP send adapter on this host instance.	Number	The counter is incremented after a response message is completely read by the HTTP send adapter from HTTP servers.
	Messages received: Indicates the rate at which HTTP response messages are received by the HTTP send adapter on this host instance.	Msgs/Sec	The counter applies only to response messages that have been completely read by the HTTP send adapter. Ideally, the value of this measure should be high. A low value indicates that the HTTP send adapter is experiencing delays while receiving messages from the BizTalk server. Further investigation may be required to diagnose the root-cause of the slowdown.
	Messages sent: Indicates the total number of HTTP requests sent by the HTTP send adapter on this host instance to the destination URL.	Number	The counter is incremented only for request messages that have reached the destination URL.
	Messages sent: Indicates the rate at which HTTP requests were sent by the HTTP send adapter on this host instance to the destination URL.	Msgs/Sec	The counter applies only to request messages that have reached the destination URL. Ideally, the value of this measure should be high. A low value indicates that the HTTP send adapter is experiencing delays while sending messages to the destination URL on an HTTP POST request. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.9 BT Messaging Latency Test

One of the key services provided by the BizTalk server messaging engine is the mechanism for communicating across the applications that a business process uses. As the first steps towards enabling this communication, the messaging engine receives messages from a source application through a **receive adapter**. The message is then processed through a **receive pipeline** and delivered into a database called the **MessageBox**. Depending upon the nature of the messages delivered to the MessageBox, the messaging engine dispatches the messages to their appropriate orchestrations; each orchestration then takes whatever action the business process requires. The result of this processing is typically another message, produced by the orchestration and saved in the MessageBox. This message, in turn, is processed by a **send pipeline**, and sent out to the application for which it is destined, using a **send adapter**.

The health of the messaging engine relies heavily on how quickly messages are processed at each step of the electronic data exchange that has been described above. Administrators should be promptly notified of even the slightest of latencies in this communication, so that they can take the measures to curb it before it causes a significant delay in the delivery of messages to the target. The **Messaging latency** test serves this purpose.

The test closely observes the time taken by the messaging engine to send a message to the MessageBox and to send a message in the MessageBox to the target destination, and proactively alerts administrators to delays; this way, administrators will not only be able to promptly detect latencies experienced by the messaging engine, but will also be able to pin-point where the delay originated - while delivering messages to the MessageBox? or while delivering messages to the destination application?

Purpose	Closely observes the time taken by the messaging engine to send a message to the MessageBox and to send a message in the MessageBox to the target destination, and proactively alerts administrators to delays; this way, administrators will not only be able to promptly detect latencies experienced by the messaging engine, but will also be able to pin-point where the delay originated - while delivering messages to the MessageBox? or while delivering messages to the destination application?		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Inbound Latency: Indicates the time taken by the messaging engine to receive a document from the receive adapter and publish it to the MessageBox.	Secs	Ideally, the value of this measure should be low. A very high value is indicative of a slowdown while publishing documents to the MessageBox.
	Outbound Latency: Indicates the time taken by the messaging engine to receive a document from the MessageBox and send it to the adapter.	Secs	Ideally, the value of this measure should be low. A very high value is indicative of a slowdown in publishing documents to the destination.
	Request-Response Latency: Indicates the total time taken by the messaging engine to process a request document from the adapter and send back a response document to the adapter.	Secs	Ideally, the value of this measure should be low. A very high value is indicative of delays in communication across applications.

4.2.1.10 BT Msmq Receive Adapter Test

With the BizTalk Server Adapter for MSMQ (the MSMQ adapter), you can send and receive messages to Microsoft Message Queuing (also known as MSMQ) queues using Microsoft BizTalk Server. The MSMQ adapter works with transactional and non-transactional, public and private, and local and remote queues. Additionally, the MSMQ adapter provides large (greater than 4 MB) message support and gives you access to Message Queuing features such as messaging over HTTP and multi-cast messaging. The key features of the MSMQ adapter are:

- Can be configured to deliver messages in order.
- Provides large message support by breaking the message into parts, accumulating the parts in memory, and delivering the parts in order to the destination (more memory intensive than MSMQT).
- Provides better performance than MSMQT.
- Enables other non-BizTalk applications to use MSMQ services at the same time on the same computer.
- Requires intermediate storage of MSMQ queues. Inbound messages are written to the MSMQ queue and then picked up from the MSMQ queue by the MSMQ adapter.

By continuously tracking the messages and data received and processed by the MSMQ receive adapter for every host instance on the BizTalk server, administrators can receive an overview of the load on the adapter, and will be able to accurately judge its processing ability. This test does just that.

Purpose	Continuously tracks the messages and data received and processed by the MSMQ receive adapter for every host instance on the BizTalk server, and provides administrators with an

	overview of the load on the adapter, and will be able to accurately judge its processing ability		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes received: Indicates the total number of bytes received by the MSMQ receive adapter on this host instance.	Bytes	The counter is incremented after a message is completely read by the MSMQ receive adapter from the source queue.
	Bytes received: Indicates the rate at which bytes were received by the MSMQ receive adapter on this host instance.	Bytes/Sec	The counter applies only to messages that have been completely read by the MSMQ receive adapter from the source queue.
	Messages received: Indicates the number of messages received by the MSMQ receive adapter on this host instance.	Number	The counter is incremented after a message is completely read by the MSMQ receive adapter from the source queue. This measure is a good indicator of the load on the MSMQ receive adapter.
	Messages received: Indicates the rate at which messages were received by the MSMQ receive adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have been completely read by the MSMQ receive adapter from the source queue. Ideally, the value of this measure should be high. A low value indicates that the MSMQ receive adapter is experiencing delays while reading messages from the source queue. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.11 BT Msmq Send Adapter Test

With the BizTalk Server Adapter for MSMQ (the MSMQ adapter), you can send and receive messages to Microsoft Message Queuing (also known as MSMQ) queues using Microsoft BizTalk Server. The MSMQ adapter works with transactional and non-transactional, public and private, and local and remote queues. Additionally, the MSMQ adapter provides large (greater than 4 MB) message support and gives you access to Message Queuing features such as messaging over HTTP and multi-cast messaging. The key features of the MSMQ adapter are:

- Can be configured to deliver messages in order.
- Provides large message support by breaking the message into parts, accumulating the parts in memory, and delivering the parts in order to the destination (more memory intensive than MSMQT).
- Provides better performance than MSMQT.
- Enables other non-BizTalk applications to use MSMQ services at the same time on the same computer.
- Requires intermediate storage of MSMQ queues. Inbound messages are written to the MSMQ queue and then picked up from the MSMQ queue by the MSMQ adapter.

By continuously tracking the messages and data sent by the MSMQ send adapter for every host instance on the BizTalk server, administrators can receive an overview of the load on the adapter, and will be able to accurately judge its processing ability. This test does just that.

Purpose	Measure the current load on the MSMQ send adapter for each host instance of the BizTalk server, and isolate bottlenecks in load processing		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Bytes sent: Indicates the total number of bytes sent by the MSMQ send adapter on this host instance.	Bytes	The counter is incremented only for messages that have reached the destination queue.
	Bytes sent: Indicates the rate at which bytes were sent by the MSMQ send adapter on this host instance.	Bytes/Sec	The counter applies only to messages that have reached the destination queue.
	Messages sent: Indicates the number of messages sent by the MSMQ send adapter on this host instance.	Number	The counter is incremented only for messages that have reached the destination queue.
	Messages sent: Indicates the rate at which messages were sent by the MSMQ send adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have reached the destination queue. Ideally, the value of this measure should be high. A low value indicates that the MSMQ send adapter is experiencing delays while sending messages to the destination queue. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.12 BT Pop3 Adapter Test

The Post Office Protocol 3 (POP3) adapter is used to retrieve data from a server that houses POP3 mailboxes into a BizTalk Server by means of the POP3 protocol. The POP3 adapter consists of only one adapter, a receive adapter. This receive adapter controls the receive locations that use the POP3 adapter.

The POP3 receive adapter retrieves e-mail from a specified mailbox on a specified POP3 server. By default, the POP3 receive adapter applies MIME processing to the e-mail messages that it downloads and submits these messages to BizTalk Server as multipart BizTalk messages. The POP3 receive adapter can receive and process e-mail in the following formats:

- Plain text
- MIME encoded
- MIME encrypted
- MIME encoded and signed
- MIME encrypted and signed

To monitor the session and message load on the POP3 adapter so that, overload conditions and processing bottlenecks are accurately identified, use the **POP3 adapter** test.

Purpose	To monitor the session and message load on the POP3 adapter so that, overload conditions and processing bottlenecks are accurately identified		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active sessions: Indicates the number of open POP3 connections that the POP3 adapter on this host instance is currently managing.	Number	This is a good indicator of the session load on the adapter.
	Bytes received: Indicates the total number of bytes downloaded by the POP3 adapter on this host instance from a mail server.	Bytes	This is a good indicator of the data load on the adapter.
	Bytes received: Indicates the rate at which bytes that the POP3 adapter on this host instance downloaded from a mail server.	Bytes/Sec	A consistent decrease in this value could indicate a processing bottleneck.
	Messages received: Indicates the number of messages that the POP3 adapter on this host instance downloaded from the mail server.	Number	This is a good indicator of the load on the adapter.

	Messages received: Indicates the rate at which the POP3 adapter on this host instance downloaded messages from the mail server.	Msgs/Sec	A consistent decrease in this value could indicate a processing bottleneck.

4.2.1.13 BT SMTP Adapter Test

The Simple Mail Transfer Protocol (SMTP) adapter is used to exchange information between a BizTalk Server and other applications by means of the SMTP protocol. BizTalk Server can send messages to other applications by creating an e-mail message and delivering it to a specified e-mail address. The SMTP adapter consists of only one adapter, a send adapter. The send adapter controls the send ports that use the SMTP adapter. Internally, the SMTP send adapter creates an SMTP-based e-mail message and sends it to a target e-mail address. The target e-mail address is a property of the SMTP adapter. The SMTP send adapter gets messages from the server and posts them to an SMTP server that sends them to e-mail recipients.

Using this test, you can figure out how quickly the SMTP send adapter sends out messages to other applications, and thus promptly detect slowdowns in message delivery.

Purpose	Helps figure out how quickly the SMTP send adapter sends out messages to other applications, and thus promptly detects slowdowns in message delivery		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages sent: Indicates the total number of messages sent by the SMTP adapter on this host instance to the target e-mail address.	Number	The counter is incremented only for messages that have been transmitted to the SMTP server.

	Messages sent: Indicates the rate at which messages were sent by the SMTP adapter on this host instance to the target e-mail address.	Msgs/Sec	The counter applies only to messages that have been transmitted to the SMTP server. Ideally, the value of this measure should be high. A low value indicates that the SMTP send adapter is experiencing delays while sending messages to the target e-mail address. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.14 BT Soap Receive Adapter Test

The SOAP adapter enables you to publish orchestrations as Web services and consume external Web services. The SOAP adapter consists of two adapters — a send adapter and receive adapter.

The SOAP receive adapter is used to receive Web service requests. The SOAP receive adapter creates a BizTalk Message object, and promotes the associated properties to the message context.

This test enables you to determine the web service request load on the SOAP receive adapter at any given point in time, and helps you assess the processing capability of the adapter.

Purpose	Enables you to determine the web service request load on the SOAP receive adapter at any given point in time, and helps you assess the processing capability of the adapter		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Messages received: Indicates the total number of messages that are received by the SOAP receive adapter on this host instance.	Number	The counter is incremented after a request message is completely read by the adapter from the SOAP client.
	Messages received: Indicates the rate at which the messages are received by the SOAP receive adapter on this host instance.	Msgs/Sec	The counter applies only to request messages that have been completely read by the adapter from the SOAP client. Ideally, the value of this measure should be high. A consistent decrease in this value indicates that the SOAP receive adapter is experiencing delays while reading messages from the SOAP client. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.15 BT Soap Send Adapter Test

The SOAP send adapter is used to call a web service. The SOAP send adapter reads the message context on the BizTalk Message object to get the proxy name and calls the associated external Web service proxy.

Monitor the load on the SOAP send adapter and be proactively alerted to processing bottlenecks in the adapter with the help of the **Soap send adapter** test.

Purpose	Monitor the load on the SOAP send adapter and be proactively alerted to processing bottlenecks in the adapter
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Messages sent: Indicates the total number of messages that are sent by the SOAP send adapter on this host instance.	Number	The counter is incremented only for messages that have reached the destination URL.
	Messages sent: Indicates the rate at which the messages are sent by the SOAP sent adapter on this host instance.	Msgs/Sec	The counter applies only to messages that have reached the destination URL. Ideally, the value of this measure should be high. A consistent decrease in this value indicates that the SOAP send adapter is experiencing delays while writing messages to the destination URL. Further investigation may be required to diagnose the root-cause of the slowdown.

4.2.1.16 BT Sql Receive Adapter Test

The SQL adapter exchanges data between the BizTalk Server and a SQL Server database. You can use the SQL adapter to poll data from one or more data tables and transmit the data as one or more XML messages to BizTalk Server. You can also use the SQL adapter to move large amounts of data to or from the SQL Server database as part of a BizTalk Server messaging or orchestration solution. In addition, you can use the SQL adapter to insert, update, and delete data in SQL Server tables by using SQL updategrams or by invoking stored procedures.The SQL adapter consists of two adapters—a receive adapter and a send adapter.

The SQL receive adapter is a polling adapter that periodically polls for SQL result sets.

This test monitors the load on the **SQL Receive Adapter** and proactively alerts you to potential overload conditions.

Purpose	Monitors the load on the **SQL Receive Adapter** and proactively alerts you to potential overload conditions
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Messages received: Indicates the total number of messages that are read by the SQL receive adapter from the SQL server.	Number	A high value could indicate an overload condition.
	Messages received: Indicates the rate at which the messages are read by the SQL receive adapter from the SQL server.	Msgs/Sec	A consistent decrease in the value of this measure points you to current/potential bottlenecks in the processing of messages.

4.2.1.17 BT File Send Adapter Test

The SQL adapter exchanges data between the BizTalk Server and a SQL Server database. You can use the SQL adapter to poll data from one or more data tables and transmit the data as one or more XML messages to BizTalk Server. You can also use the SQL adapter to move large amounts of data to or from the SQL Server database as part of a BizTalk Server messaging or orchestration solution. In addition, you can use the SQL adapter to insert, update, and delete data in SQL Server tables by using SQL updategrams or by invoking stored procedures.The SQL adapter consists of two adapters—a receive adapter and a send adapter.

The SQL send adapter is used to send dynamically created updategrams or dynamically invoked stored procedures to SQL Server. An updategram is an XML fragment that inserts, updates, or deletes data in a SQL Server database by mapping XML nodes against database tables and columns. SQL Server returns an optional response document after the updategram completes, which contains the success status of the update. If a failure occurs during the update, the SQL adapter throws an exception that the BizTalk Messaging Engine handles. When the SQL send adapter is configured to invoke a stored procedure, it returns any results in the form of a single XML-formatted record set.

This test monitors the load on the **SQL Send Adapter** and proactively alerts you to potential overload conditions and processing bottlenecks.

Purpose	Monitors the load on the **SQL Send Adapter** and proactively alerts you to potential overload conditions and processing bottlenecks
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Messages sent: Indicates the total number of messages that are sent by the SQL send adapter to the destination SQL table in the SQL server database.	Number	
	Messages sent: Indicates the rate at which the messages are sent by the SQL send adapter to the destination SQL table in the SQL server database.	Msgs/Sec	A consistent decrease in the value of this measure points you to current/potential bottlenecks in the processing of messages.

4.2.2 The Message Box Layer

The heart of the publish/subscribe engine in Microsoft BizTalk Server is the MessageBox database. The MessageBox is made up of two components: one or more Microsoft SQL Server databases and the Messaging Agent. The SQL Server database provides the persistence store for many things including messages, message parts, message properties, subscriptions, orchestration state, tracking data, host queues for routing, and others. The BizTalk Server group may have one or more MessageBox databases into which it publishes messages and from which subscribers to those messages extract messages.

The database provides some of the logic related to routing messages and fulfilling subscriptions. The Message Agent, however, is the component that encapsulates and abstracts the database component and is the interface used by BizTalk Server to interact with the MessageBox. The Message Agent is a Component Object Model (COM) component that provides interfaces for publishing messages, subscribing to messages, retrieving messages, and so on. This interface is the only mechanism used by other BizTalk Server components, including the adapter framework and orchestrations, to interact with the MessageBox.

Using the tests mapped to this layer you can monitor the health of the BizTalk server MessageBox and the efficiency of the SQL Server agent jobs.

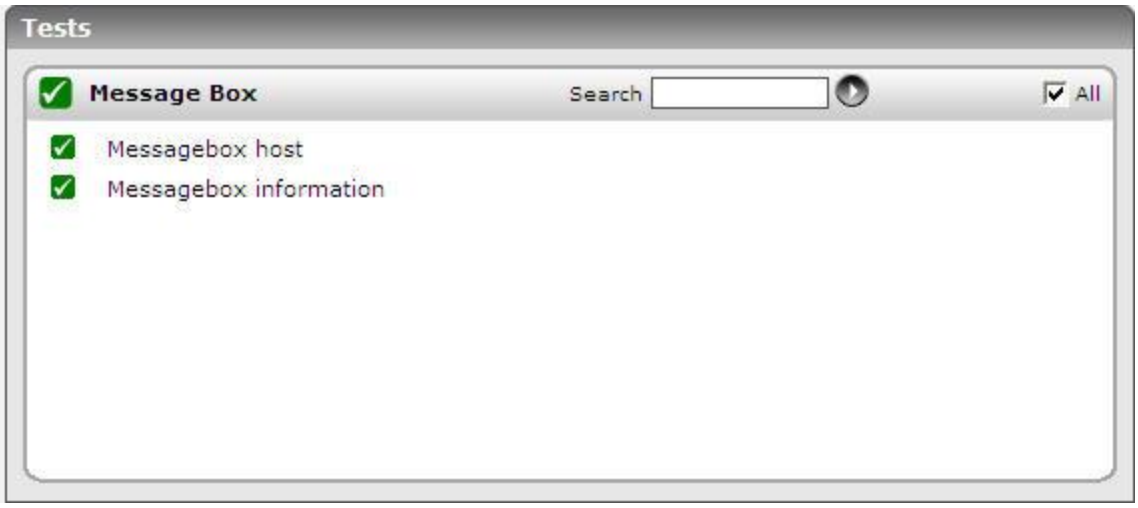

Figure 4.8: The tests mapped to the Message Box layer

4.2.2.1 BT Messagebox Host Test

The first time you configure a BizTalk server, the following set of tables are created in the MessageBox database for a BizTalkServerApplicationHost:

- The BizTalkApplicationQ
- The BizTalkServerApplicationQ_Suspended
- The BizTalkServerApplicationQ_Scheduled
- The InstanceStateMessageReferences_BizTalkServerApplication

BizTalk uses these tables to keep **references** of all the messages that are **"live"** in the system . That is: Messages with active subscriptions, suspended messages, and *awaiting messages* associated to each host.

The word ***references*** implies that the host tables are only pointers to the **Spool** table, but the real messages itself are saved in another set of tables (messageparts, parts and fragments).

This test monitors the number of message references in the host queue tables, and proactively alerts administrators to the following:

- A sudden/consistent increase in the length of the host queues
- Too many message references in the suspended queue

Purpose	This test monitors the number of message references in the host queue tables, and proactively alerts administrators to the following: • A sudden/consistent increase in the length of the host queues • Too many message references in the suspended queue
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Message references in Instance Queue: Indicates the number of message references in the instance state queue of this host instance.	Number	The State Queue table holds the list of messages that have been processed by an instance but **will be needed later.** When an orchestration uses the State Queue, it is usually because the orchestration performed some operations on a message, persisted the message, and might need the message later. This is normal operation, and you should take this into account when determining correct sizing of the State Queue
	Instances of Host Queue: Indicates the number of instances of the host queue for this host instance.	Number	
	Messages in Host Queue: Indicates the number of messages in the host queue of this host instance.	Number	Generally, this queue should not grow too large. The length of the queue indicates the number of messages waiting to be processed. A large number means you could have a backlog.
	Suspended Messages in Host Queue: Indicates the number of suspended messages for this host instance.	Number	When a message gets suspended it remains in the messagebox until resume or terminate actions occurs. So, if the suspended queue keeps growing, the performance of the BizTalk server will continue to get affected. A suspended message can be due, for example, to parsing errors, serialization errors, failed transmissions, or the inability to find a subscription.

4.2.2.2 BT Messagebox Information Test

The BizTalk server includes certain SQL agent jobs to assist administrators in managing the BizTalk server databases.

Using this test, you can monitor the time taken to perform each of these SQL agent jobs so that, jobs that took too long to complete can be instantly identified.

Purpose	Monitors the time taken to perform each of these SQL agent jobs so that, jobs that took too long to complete can be instantly identified		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Dead Processes Cleanup: Indicates the time taken by the MessageBox_DeadProcesses_Cleanup_BizTalkMsgBoxDb job for this host instance to complete.	Secs	This job detects when a BizTalk Server host instance (NT service) has stopped and releases all work that was being done by that host instance so that it can be worked on by another host instance.
	Cleanup Messages: Indicates the time taken by the MessageBox_Message_Cleanup_BizTalkMsgBoxDb job for this host instance to complete its work.	Secs	This job removes all messages that are no longer being referenced by any subscribers in the BizTalk MessageBox (BizTalkMsgBoxDb) database tables. **Note:** This is an unscheduled job which is started by the MessageBox_Message_ManageRefCountLog_BizTalkMsgBoxDb job. Do not manually start this job.
	Total Instances: Indicates the total number of host instances that exist within a message box.	Number	

	Cleanup Message Parts: Indicates the time taken by the MessageBox_Parts_Cleanup_BizTalkMsgBoxDb job for this host instance to complete its work.	Secs	This job removes all message parts that are no longer being referenced by any messages in the BizTalk MessageBox (BizTalkMsgBoxDb) database tables. All messages are made up of one or more message parts, which contain the actual message data.
	Spool Size: Indicates the size of the spool that is available on a particular message box in this host instance.	Number	The primary measure of sustainability over time is that a backlog is not allowed to grow indefinitely. In other words, over time, there must be a balance between the high and low peak throughput levels so that the MessageBox database is able to maintain a constant and manageable average backlog. The primary measure of backlog is the depth of the spool table. The message bodies are handled via a set of tables represented by the **spool table**. The Spool can start growing for multiple reasons. One reason for Spool growth is if the application queues are growing. **Application queues** host in-flight transition data. They could grow due to various reasons like downstream bottlenecks and/or resource contention. If the application queues are small and the Spool is still large, verify that the purge jobs are keeping up. Ensure that the SQL-Agent Service is running and then verify that the following jobs are successfully completing: • MessageBox_Message_Cleanup_BizTalkMessageBoxDb • MessageBox_Parts_Cleanup_BizTalk MessageBoxDb One reason for this is if the SQL-Server machine is experiencing severe CPU contention, impacting the ability of the purge jobs to keep up due to CPU starvation.
	Tracked Messages: Indicates the time taken by the DTA Purge and Archive job of this host instance to complete its execution.	Secs	This job automatically archives data in the BizTalk Tracking (BizTalkDTADb) database and purges obsolete data.

	Tracking Data Size: Indicates the size of the data table that is tracked from the message available for this host instance.	Number	As BizTalk Server processes more and more data on your system, the BizTalk Tracking (BizTalkDTADb) database continues to grow in size. Unchecked growth decreases system performance and may generate errors in the Tracking Data Decode Service (TDDS). In addition to general tracking data, tracked messages can also accumulate in the MessageBox database, causing poor disk performance. This implies that ideally the value of this measure should be low. By archiving and purging data from the BizTalk Tracking database, you can maintain a healthy system, as well as keep your tracking data archived for future use.
	Tracking Spool Cleanup: Indicates the time taken to purge inactive spools in the tracking database tables so as to free database space.	Secs	

4.2.3 The Orchestration Engine Layer

An orchestration is a flexible, powerful tool for representing an executable business process based on XLANG/s language. At run time, the BizTalk Orchestration Engine executes XLANG/s files that are produced by BizTalk Orchestration Designer. Orchestration Designer is a rich graphical tool for visually designing business processes. It generates XLANG/s files that have an .odx extension and contain additional visualization information in their headers and custom attribute information in their bodies.

The primary functions of the orchestration engine are:

- Persistence
- Hosting the .NET components
- Transactions
- Large message support
- Runtime validation
- Load throttling

Using the tests mapped to the **Orchestration Engine** layer you can monitor the orchestrations, the BAM interceptor, and the tracking data decode service offered by the Orchestration engine.

Figure 4.9: The tests mapped to the Orchestration Engine layer

4.2.3.1 BT Orchestrations Test

Orchestrations are executable business processes that can subscribe to (receive) and publish (send) messages through the MessageBox database. In addition, orchestrations can construct new messages. Messages are received using the subscription and routing infrastructure.

When subscriptions are filled for orchestrations, a new instance is activated and the message is delivered, or in the case of instance subscriptions, the instance is rehydrated if necessary and the message is then delivered. When messages are sent from an orchestration, they are published to the MessageBox in the same manner as a message arriving on a receive location with the appropriate properties getting inserted into the database for use in routing.

Messages that are constructed in an orchestration must be placed in the MessageBox database and referenced by the orchestration instance, but they should not be published because they have not yet been sent. The XLANG/s subservice makes calls to the Message Agent API to insert messages directly. This allows the orchestration engine to insert the message body into the MessageBox and have it directly associated with the running orchestration instance. The persistence of the constructed message in the MessageBox database is coordinated with persistence points in the orchestration as an additional optimization of database operations.

This test helps you determine the number of orchestrations that were created on each host instance, and also tracks the status of these orchestrations over time, thereby promptly alerting you when too many orchestrations are suspended or discarded. The test also tracks the memory usage of the orchestrations, and alerts you if excessive memory is being consumed.

Purpose	Helps you determine the number of orchestrations that were created on each host instance, and also tracks the status of these orchestrations over time, thereby promptly alerting you when too many orchestrations are suspended or discarded
Target of the test	A BizTalk Server 2010
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance on the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Idle orchestrations: Indicates the number of idle orchestration instances currently hosted by this host instance.	Number	This refers to orchestrations that are not making progress but are not dehydratable, as when the orchestration is blocked waiting for a receive, listen, or delay in an atomic transaction.
	Orchestrations created: Indicates the number of orchestration instances that were created since this host instance was started.	Number	
	Orchestrations created: Indicates the rate at which the orchestration instances were created on this host instance.	Orchestrations / Sec	
	Running orchestrations: Indicates the number of orchestration instances that are currently executing on this host instance.	Number	
	Orchestrations completed: Indicates the number of orchestration instances that were completed since this host instance was started.	Number	
	Orchestration completion rate: Indicates the rate at which the orchestration instances are completed.	Orchestrations/Sec	A high value is desired for this measure. A low value or a steady decline in the value of this measure could indicate an execution bottleneck.

	Orchestrations discarded: Indicates the number of orchestration instances discarded from memory since this host instance was started.	Number	An orchestration can be discarded if the engine fails to persist in its state.
	Orchestrations discarded: Indicates the rate at which orchestrations instances were discarded from the memory of this host instance.	Orchestrations/Sec	
	Orchestrations suspended: Indicates the number of orchestration instances that are suspended since this host instance was started.	Number	All failures encountered in orchestrations appear as exceptions. If an orchestration does not include any **CatchException** shape for an exception, the exception causes the orchestration to be Suspended, but not resumable. This means that message and service instance tracking, or a WMI script, cannot recover the instance. However, you can save all messages associated with the Suspended (not Resumable) instance using tracking (or WMI script) for diagnostic and manual retry. To diagnose the problem, use the Orchestration Debugger to see the last shape executed before the instance was suspended. You can also view exception details using the Orchestration Debugger.
	Orchestrations suspended: Indicates the rate at which orchestrations were suspended on this host instance.	Orchestrations/Sec	
	Orchestrations rehydrated: Indicates the number of orchestration instances that were rehydrated since this host instance was started.	Number	Rehydration is the process of deserializing the last running state of an orchestration from the database. The orchestration engine can be triggered to rehydrate an orchestration instance by the receipt of a message or by the expiration of a

	Orchestrations rehydrated: Indicates the rate at which orchestrations instances were rehydrated on this host instance.	Orchestrations/Sec	time-out specified in a Delay shape. It then loads the saved orchestration instance into memory, restores its state, and runs it from the point where it left off.
	Orchestrations dehydrated: Indicates the number of orchestration instances that were dehydrated since this host instance was started.	Number	Dehydration is the process of serializing the state of an orchestration into a SQL Server database. The orchestration engine might determine that an orchestration instance has been idle for a relatively long period of time. It calculates thresholds to determine how long it will wait for various actions to take place, and if those thresholds are exceeded, it dehydrates the instance. This can occur under the following circumstances: • When the orchestration is waiting to receive a message, and the wait is longer than a threshold determined by the engine. • When the orchestration is "listening" for a message, as it does when you use a **Listen** shape, and no branch is triggered before a threshold determined by the engine. The only exception to this is when the **Listen** shape contains an activation receive. • When a delay in the orchestration is longer than a threshold determined by the engine. The engine dehydrates the instance by saving the state, and frees up the memory required by the instance. By dehydrating dormant orchestration instances, the engine makes it possible for a large number of long-running business processes to run concurrently on the same computer. This implies that the larger the number and rate of dehydrations minimal will be the use of system resources.
	Orchestrations dehydrated: Indicates the rate at which orchestration instances were dehydrated on this host instance.	Orchestrations/Sec	

	Pending messages: Indicates the number of received messages for which receipt has not yet been acknowledged to the message box from the orchestration.	Number	A very large value could indicate a processing bottleneck.
	Pending work items: Indicates the number of code execution blocks that are scheduled for execution in the orchestration.	Number	
	Failure connections: Indicates the number of attempted database connections that has failed since this host instance was started.	Number	Ideally, the value of this measure should be 0.
	Database transactions: Indicates the number of database transactions performed since the host instance was started.	Number	
	Transactions / Sec: Indicates the rate of database transactions performed by the orchestrations hosted by this host instance.	Trans/Sec	
	Current Orchestrations Instances: Indicates the number of orchestration instances currently hosted by this host instance.	Number	
	Private memory: Indicates the allocated private memory for this host instance.	MB	This is the current size of memory that this process has allocated that cannot be shared with other processes.
	Virtual memory: Indicates the reserved virtual memory for this host instance.	MB	This is the current size of the virtual address space the process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.

	Total physical memory: Indicates the percentage of total physical memory used on this host instance.	Percent	The dehydration behavior of BizTalk Server depends entirely on how much memory is available and how much memory is being used. The dehydration behavior is different with different amounts of memory and differences in memory use between 32-bit and 64-bit hosts.

4.2.3.2 BT BAM Interceptor Test

Information workers need flexibility in looking at and evaluating business processes. A purchasing manager might need to see how many POs are approved and denied each day, for example, while a sales manager might want an hourly update on what products are being ordered. Meeting these diverse needs requires a general framework for tracking what's going on with a particular business process. This is exactly what the Business Activity Monitoring (BAM) component in Microsoft BizTalk Server provides.

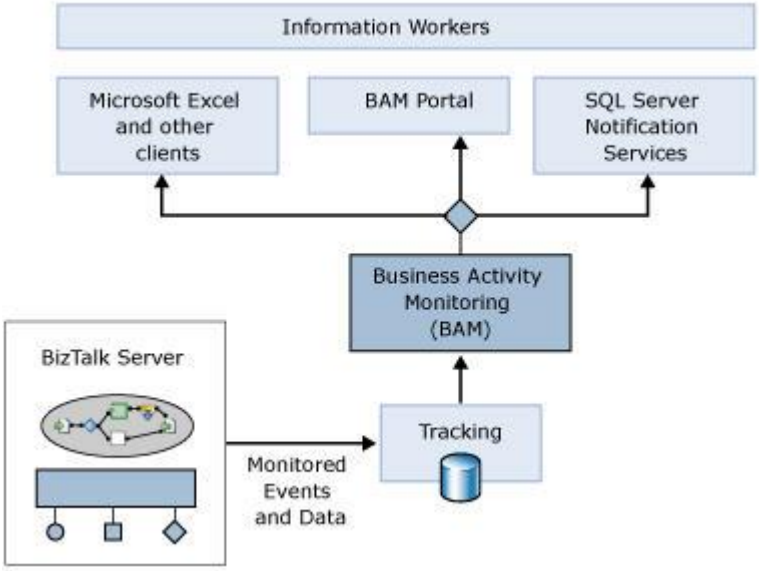

Figure 4.10: How does BAM work?

As the figure above illustrates, the BAM component allows monitoring of events and data produced by a BizTalk application. This information is made accessible using SOAP-callable Web services, and it can be accessed in several different ways, as follows:

- Through Microsoft Excel or other desktop clients, such as a custom dashboard application.
- Using a BAM portal, a component in BizTalk Server that enables business users to examine and configure BAM information.
- Through SQL Server Notification Services, allowing BAM information to be delivered as notifications.

The BAM Interceptor is an object that lets you instrument your application to capture data of interest. Using this test,

you can monitor the BAM interceptors, and swiftly spot the failure of BAM events.

Purpose	You can monitor the BAM interceptors, and swiftly spot the failure of BAM events		
Target of the test	A BizTalk Server 2010		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total Failed Events: Indicates the total number of failed BAM events that occured during data flush.	Number	

4.2.3.3 BT Tracking Data Decode Service Test

The BAM Event Bus Service, also known as the Tracking Data Decode Service (TDDS), processes tracking data (streams) stored in a source database and persists that data in such a way that it is easy to query it at a later date. The BAM Event Bus service moves Business intelligence data to the BAM Primary Import database and BizTalk Health Monitoring data to the DTA database.

This test reveals the processing power of the TDDS by reporting the number of batches, events, and records it processes, and also sheds light on failures experienced by the TDDS while processing.

Purpose	Reveals the processing power of the TDDS by reporting the number of batches, events, and records it processes, and also sheds light on failures experienced by the TDDS while processing
Target of the test	A BizTalk Server 2010
Agent deploying the	An internal agent

test			
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **ISPASSIVE** - If the **ISPASSIVE** parameter is set to **Yes**, then it means that, by default, all BizTalk servers being monitored by the eG system are the passive servers of a BizTalk cluster. No alerts will be generated if the servers are not running. Measures will be reported as "Not applicable" by the agent if the servers are not up.		
Outputs of the test	One set of results for each host instance of the BizTalk server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total Failed Batches: Indicates the total number of batches that the TDDS has failed to process on this host instance.	Number	
	Total Failed Events: Indicates the total number of batches that the TDDS has failed to process on this host instance.	Number	Ideally, this value should be 0.
	Total Events: Indicates the total number of events that are processed by the TDDS since you started it on this host instance.	Number	

	Total Records: Indicates the total number of records that are processed by the TDDS since you started it on this host instance.	Number	

Chapter 5

Monitoring DHCP Servers

The Microsoft® Windows® 2000 Server network operating system builds on the Microsoft support for Dynamic Host Configuration Protocol (DHCP).

Each host computer connected to a TCP/IP network must be assigned a unique IP address. The Microsoft DHCP server allows the network administrator to dynamically assign network settings to the clients that connect to a network.

The DHCP server offers the following features:

- Integration of DHCP with DNS.
- Dynamic assignment of IP addresses allows,address reuse through leases.
- Multicast address allocation.
- Automatic pushdown of configurations to clients allows configuration changes to be applied transparently.

If the DHCP server experiences an overload or a slowdown while processing requests, it is bound to delay the automatic discovery of additions (client / server) to the network and the assignment of identification (i.e., IP address) to them; consequently, users may be denied timely access to critical clients or servers. Continuous monitoring of the DHCP server can alone help administrators in promptly identifying and resolving such problem conditions.

eG Enterprise prescribes a unique *DHCP* monitoring model (see Figure 5.1) for the DHCP server, which keeps a watchful eye on the requests received and acknowledgements sent by the server to help administrators determine the following:

- How quickly is the DHCP server processing request packets? Were too many requests enqueued? Have too many packets expired?
- Is the hardware on the DHCP server adequately sized to facilitate swift processing of the request packets?
- Were any negative acknowledgement messages sent by the DHCP server?
- Were any DHCP decline messages received by the server?
- Have enough IP addresses been configured on the server for assignment to clients?

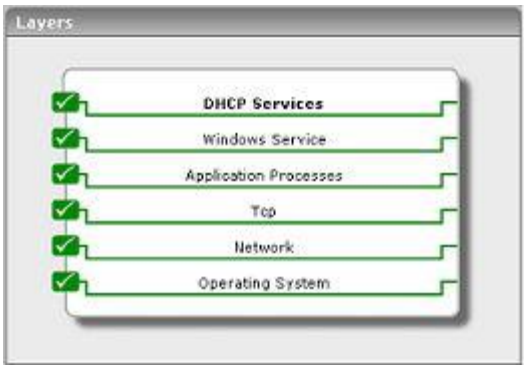

Figure 5.1: Layer model of a DHCP server

Every layer of Figure 5.1 above is mapped to a set of tests. The eG agent executing on the DHCP server runs these tests on the server, and extracts the metrics of interest.

Since the bottom 5 layers of Figure 5.1 have already been discussed in the *Monitoring Unix and Windows Servers* document, the section to come will discuss the **DHCP Services** layer only.

5.1 The DHCP Services Layer

The tests associated with this layer do the following:

- Track the overall responsiveness of the DHCP server to requests received from clients
- Verify the availability of free IP addresses on the server for assignment to clients

Figure 5.2: Tests associated with the DHCP Services layer

5.1.1 DHCP Performance Test

This test reports the performance statistics of the Microsoft 2000 DHCP server running on the network.

Purpose	Reports the performance statistics of the DHCP server on a Windows 2000 network.
Target of the test	Any DHCP server
Agent deploying the	An internal agent

test			
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The variable name of the host for which the test is to be configured. 3. **PORT** – Refers to the port used by the DHCP server		
Outputs of the test	One set of results for server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Avg packet rate: Refers to the average time in seconds used by the DHCP server to process each packet it receives.	Pkts/sec	This measure can vary depending on the server hardware and its I/O subsystem. A sudden or unusual increase might indicate a problem, either with the I/O subsystem becoming slower or because of an intrinsic processing overhead on the server computer.
	Current message queue length: Refers to the current length of the internal message queue of the DHCP server.	Number	A large value in this measure might indicate heavy server traffic.
	Request rate: Refers to the number of DHCP request messages received per second by the DHCP server from clients.	Reqs/sec	A sudden or unusual increase in this measure indicates a large number of clients trying to renew their leases with the DHCP server
	Request acks rate: Refers to the number of DHCP acknowledgement messages sent per second by the DHCP server to clients.	Reqs/sec	A sudden or unusual increase in this measure indicates that a large number of clients are being renewed by the DHCP server
	Request nacks rate: Refers to the number of negative acknowledgement messages sent per second by the DHCP server to clients.	Reqs/sec	A very high value might indicate potential network trouble in the form of misconfiguration of either the server or clients. When servers are misconfigured, one possible cause is a deactivated scope. For clients, a very high value could be caused by computers moving between subnets, such as laptop portables or other mobile devices.
	Request declines rate: Refers to the number of DHCP decline messages received per second by the DHCP server from clients.	Reqs/sec	A high value indicates that several clients have found their address to be in conflict, possibly indicating network trouble.

	Packets expired rate: Refers to the number of packets per second that expire and are dropped by the DHCP server.	Pkts/sec	A large value in this measure indicates that the server is either taking too long to process some packets while other packets are queued and becoming stale, or traffic on the network is too high for the server to manage.
	Packet drop rate: Refers to the number of duplicate packets per second dropped by the DHCP server.	Pkts/sec	This measure can be affected by multiple clients or network interfaces forwarding the same packet to the server. A large value in this measure indicates that either clients are probably timing out too fast or the server is not responding fast enough.
	Requests release rate: Refers to the number of DHCP release messages received per second by the DHCP server from clients.	Reqs/sec	This measure only exists if a DHCP client sends a release message to the server. This measure remains low for many DHCP network configurations .

5.1.2 DHCP Utilization Test

This test reports general statistics pertaining to the Microsoft 2000 DHCP server running on the network.

Purpose	Reports the statistics of the DHCP server on Windows 2000 network.		
Target of the test	Any DHCP server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured. 3. **PORT –** Refers to the port used by the DHCP server.		
Outputs of the test	One set of results for server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Current addresses in use: Refers to the number of IP addresses in use in the target network.	Number	This measure indicates the number of IP addresses assigned to clients in the target network.

	Free addresses: Refers to the number of free IP addresses available in the target network.	Number	This measure indicates the number of IP addresses available for allocation to clients in the target network.
	Total addresses: Indicates the total number of IP addresses allocated to the target network.	Number	
	Current address usage: Indicates the percentage of IP addresses that are currently used in the target network.	Percent	If the value of this measure reaches close to 100% then it is a cause for concern, which indicates excessive usage of the IP addresses in the target network.

Chapter 6

Monitoring the Windows Internet Name Service (WINS)

The Windows Internet Name Service (WINS) provides a distributed database for registering and querying dynamic mappings of NetBIOS names for computers and groups used on your network. WINS maps NetBIOS names to IP addresses and was designed to solve the problems arising from NetBIOS name resolution in routed environments. The main benefit of a WINS server is that it avoids the need for broadcasts to resolve computer names to IP addresses.

Typically, WINS servers use UDP port 137. **This port should be provided when you manually add a WINS server for monitoring**. The steps below highlight how WINS works:

- **Name Registration**: When a WINS client initializes, it registers its NetBIOS name by sending a name request to the configured WINS server. All services get registered as they are initialized in the WINS server database. If the WINS server is available and the name is not registered by another machine, the WINS server returns a successful registration message.
- If the NetBIOS name is already registered in the WINS database, the WINS server will send a challenge to the current registered owner. This request will be sent 3 times at 500ms intervals. If the current owner responds the WINS server will send a negative name resolution response to the WINS client attempting to register the name. If there is no response the registering client will receive a Name Registration response.
- **Name Renewal:** To continue using the same NetBIOS name, a client must renew its lease before it expires. If the client does not renew the lease, the WINS server makes it available to another WINS client. A WINS client will first attempt to refresh its name registration request after 1/8 of the TTL is completed. If the client is successful subsequent name registration requests will occur when ½ the TTL is expired.

  If the client is unsuccessful with lease renewal on the initial attempt the client will try every 2 minutes until ½ TTL is remaining. At ½ of TTL the client will revert to the secondary WINS server if configured in 1/8 TTL intervals. At completion of TTL lease, the WINS client will revert back to the primary WINS server and start the process all over again.
- **Name Release:** Before the expiry of its lease, a client can send an explicit request to release the name assigned to it.

If even one of these steps experience latencies, it could cause a significant delay in the entire process of resolving an IP address to its corresponding NetBIOS name. This could be much worse in large environments where the WINS server might have to handle hundreds of concurrent 'name resolution' requests; here, even a seemingly insignificant drop in the processing rate of the WINS server can grow in severity within minutes, and can bring the whole environment to a virtual standstill!

If such adverse consequences are to be prevented, it is recommended that you continuously monitor the processing ability of the WINS server, so that you are promptly alerted when there is any threat to its normal functioning.

eG Enterprise offers a 100% web-based *WINS* monitoring model (see Figure 6.1) that closely observes the performance of the WINS server in relation to real-time changes in load.

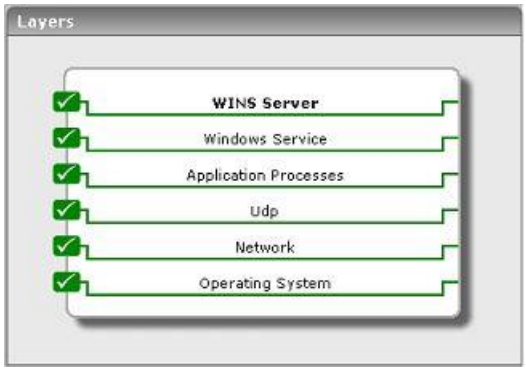

Figure 6.1: Layer model of a WINS server

Figure 6.1 comprises of a set of hierarchical layers, each of which is associated with one/more tests. The eG agent on the WINS server periodically executes these tests on the server, extracts performance data from the server, and instantly alerts administrators of an impending overload or a probable dip in the processing speed of the server.

The sections to come discuss the top layer of Figure 6.1 alone, as all other layers have been discussed in the *Monitoring Unix and Windows Servers* document.

6.1 The WINS Server Layer

Using the **Wins** test associated with it, this layer measures the rate at which the WINS server processes requests.

Figure 6.2: Test associated with the WINS server layer

6.1.1 Wins Test

This test reports general statistics pertaining to the Windows Internet Name Service (WINS).

Purpose	Reports general statistics pertaining to the WINS server
Target of the test	A WINS server
Agent deploying the	An internal agent

test			
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the WINS server		
Outputs of the test	One set of results for every server		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Queries: The total number of queries received by the WINS server	Queries/sec	This indicates the server workload. It is useful for capacity planning and to detect unusual usage situations.
	Failed queries: Total number of failed queries/sec	Failures/sec	The percentage of failed queries should be low. An unusually high number of failed queries can indicate a configuration problem, or a fault in the WINS server.
	Releases: The rate at which release requests are received and processed by the WINS server	Releases/sec	
	Failed releases: The rate of release failures	Failures/sec	Release failures could result in many names being unused for a period of time, and hence, should be minimized.
	Conflicts: The total rate of conflicts seen by the WINS server. This value includes both Unique and Group conflicts.	Conflicts/sec	
	Renewals: The total rate of renewal requests received by the WINS server. This value includes both Unique and Group renewals.	Renewals/sec	

Chapter 7

Monitoring MS Print Servers

Print servers are a popular mode of sharing printing resources in IT infrastructures. The Microsoft Windows operating system allows for specific servers to be designated and managed as print servers. Some of the key reasons for why IT administrators configure and use print servers include centralized management of print drivers, access control and prioritization of print jobs, central auditing capability or charging, etc. Since print servers are common resources for all the users of an IT infrastructure, IT administrators must continuously monitor the print servers to ensure high uptime, good performance, and scalability.

The eG Enterprise suite includes specialized monitoring capability for Microsoft Windows-based print servers. The layer model of a print server is given below (see Figure 7.1)

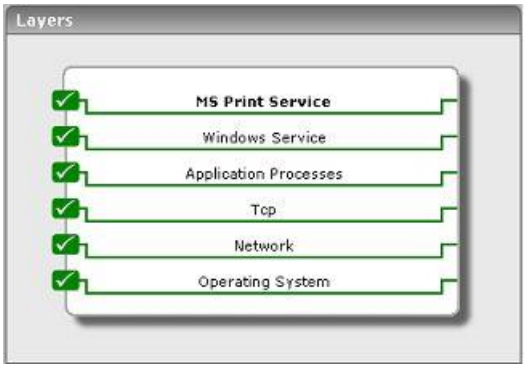

Figure 7.1: Layer model of an MS Print server

The section that follows discusses the **MS Print Service** layer only, as all other layers have been extensively discussed in the *Monitoring Unix and Windows Servers* document.

7.1 The MS Print Service Layer

This layer (see Figure 7.2) monitors the print queues on the print server and reports on their availability and overall health.

Figure 7.2: Tests associated with the MS Print Service layer

7.1.1 Print Server Test

This test auto-discovers the print queues of a print server and continuously tracks various key metrics relating to the availability and performance of each of the print queues.

Purpose	Tracks various key metrics relating to the availability and performance of each of the print queues of a print server		
Target	An MS Print server		
Agent deploying this test	An internal agent		
Configurable parameters for this test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** - The host for which the test is to be configured. 3. **PORT** - The port to which the specified **HOST** listens 4. **USEWMI** - If the **USEWMI** flag is **Yes**, then the test uses WMI to extract the statistics of interest. This option is provided because on some Windows 2000 systems (especially ones with service pack 3 or lower), the use of WMI access can cause the CPU usage of the WinMgmt process to shoot up. On such systems, set the **USEWMI** parameter value to **No**. The default is **No**.		
Outputs of the test	One set of results for every print queue monitored		
Measurements of the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Availability: Indicates whether or not the Print server is currently available.	Boolean	If the value of this measure is *1*, it indicates that the print server is available. The value *0* on the other hand, indicates that the print server is unavailable.
	Jobs services: The rate at which users' jobs are being processed over a print queue	Jobs/Sec	The value of this metric is a key indicator of a print queue's workload.

	Pages printed: The number of pages printed through a print queue during the last measurement period	Number	This is another key indicator of the workload of a print queue.
	Print traffic: Indicates the rate at which data is transmitted to a print queue for printing	KBytes/Sec	
	Current jobs: Shows the current number of jobs in a print queue.	Number	Use this counter to identify excessive use of a print queue.
	Print errors: The number of jobs to a print queue that resulted in errors during the last measurement period.	Number	This value includes the number of out of paper errors and printer not ready errors. Job errors can occur even if the connection to the printer has errors due to network problems.
	Spooled jobs: The current number of spooling jobs in a print queue	Number	
	Paper errors: The total number of out of paper errors that occurred in a print queue during the last measurement period	Number	
	Not ready errors: The total number of out of printer not ready errors that occurred in a print queue during the last measurement period	Number	

Chapter 8

Monitoring MS Proxy Servers

Microsoft Proxy Server 2.0 is an extensible firewall and content cache server, providing Internet security while improving network response time and efficiency by 50%, on average, for businesses of all sizes. It is the first firewall product to include high-performance content caching. Similarly, it is the first content cache server to provide firewall support. Microsoft Proxy Server 2.0 offers distributed (hierarchical and array-based) Web caching, providing unbeaten scalability, fault-tolerance and load balancing to meet even the rigorous demands of large enterprises and Internet Service Providers. MS Proxy Server acts as a gateway with firewall-class security between a LAN and the Internet. The product also blocks access to undesirable sites and provides other easy-to-use management features. It works with existing networks, including IPX networks, and supports several Internet protocols and services. It is therefore imperative that the MS Proxy server is continuously monitored, so that security risks to your environment are minimized, and business is transacted smoothly and efficiently.

eG Enterprise provides a specialized *Microsoft Proxy* monitoring model (see Figure 8.1 that monitors the internal health and external availability and responsiveness of the Microsoft Proxy server, and alerts administrators to potential performance issues.

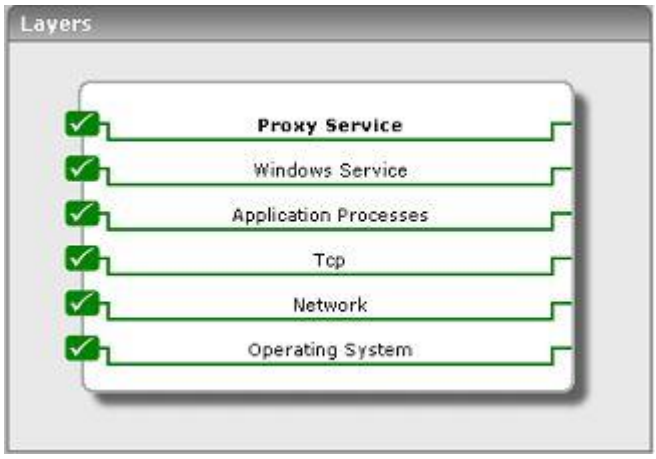

Figure 8.1: Layer model of an MS Proxy server

8.1 The Proxy Service Layer

The tests mapped to the **Proxy Service** layer monitors the performance of the following services executing on an MS Proxy server:

- The WinSock Proxy Service
- The Web Proxy Service
- The Caching service

Figure 8.2: Tests associated with the Proxy Service layer

8.1.1 Win Sock Test

The WinSock Proxy service supports Microsoft Windows operating systems using Windows Sockets. Windows Sockets is an interprocess communication mechanism derived from the Berkeley Sockets interface (originally designed for Unix systems). The Sockets interface was extended to support Windows-based clients running Microsoft implementations of TCP/IP. The name given to this Sockets interface for Windows was WinSock (for Windows Sockets). The WinSock Proxy Service support is available for both Transmission Control Protocol/Internet Protocol (TCP/IP) and Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) protocols. The WinSock Proxy service applies mainly to Windows clients, including Windows 3.x, Windows 95, and Windows NT.

This test reports the performance statistics pertaining to this WinSock Proxy Service.

Purpose	Reports the performance statistics pertaining to this WinSock Proxy service
Target of the test	An MS Proxy Server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the MS Proxy server

Outputs of the test	One set of results for every WinSock monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Accepting TCP connections: The number of TCP connection objects that will wait for TCP connections from WinSock proxy clients	Percent	A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted.
	Active sessions: The number of active sessions for the WinSock proxy service	Number	
	Active TCP connections: The total number of TCP connections that are currently transmitting data	Number	
	Active UDP connections: The number of active UDP connections	Number	
	Available worker threads: The number of available WinSock worker threads	Number	The high increase in the number may affect the performance of the host / applications.
	Data received: The rate at which data is received	KB/sec	A low value could indicate a network bottleneck
	Data transmitted: The rate at which data is submitted	KB/sec	A high value of this measure could result in a network congestion
	Failed DNS resolutions: The number of calls that have failed to resolve DNS domain name and IP address for WinSock proxy connections	Number	This value must be low; a high value indicates that there may be a network / WinSock service problem on the host.
	Pending DNS requests: The number of calls awaiting DNS domain name and IP address resolution for WinSock proxy connections	Number	This value must be low; a high value indicates that there may be a network / WinSock service problem on the host.

	Worker threads: The number of WinSock worker threads that are currently available or alive	Number	An increase in this value may affect the performance of the host / application.

8.1.2 Proxy Server Test

The Web Proxy service provides support for HTTP (a.k.a. Web publishing), FTP, Gopher, and secure (SSL) communications. The Web Proxy service works with any CERN-compliant Web browser, such as Internet Explorer or Netscape Navigator. Because the Web Proxy supports only these widely adopted Internet standard communication methods, it isn't operating system dependent. Clients running Unix, Macintosh, or Windows operating systems can communicate with the Web Proxy service as long as they're configured with a CERN-compliant Web browser.

This test reports the performance statistics pertaining to this Web Proxy service running on an MS Proxy server.

Purpose	Reports performance statistics pertaining to the Web Proxy service running on an MS Proxy server		
Target of the test	An MS Proxy Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the MS Proxy server		
Outputs of the test	One set of results for every web proxy service monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Cache hit ratio: The percentage of requests that have used cached data, to the total number of requests to the web proxy service	Percent	A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted.
	Client data receive rate: The number of active sessions for the web proxy service	Number	A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications.
	Client data transmit rate: The rate at which the data bytes are sent by the proxy server to the web proxy clients	Kb/sec	A high value could indicate a high data transfer from the proxy server to the web proxy client, which may result in congestion in network traffic

	Avg response time: The mean response time in seconds to service a request	Secs/req	High network traffic, low server performance are some of the factors that cause this measure to increase.
	Current users: The current number of users connected to the web proxy service.	Number	A high value can indicate an increase in the load on the web proxy service.
	DNS cache hits: This measure give the percentage of DNS domain names served from the proxy server cache, from the total DNS entries that are retrieved by the web proxy service.	Percent	A high value can indicate an increase in load on web proxy service.
	Failing requests: The rate of request that have completed with some error.	Reqs/Sec	The high value indicates possible problems in the web proxy service.
	FTP requests: The number of ftp requests that have been made to the web proxy service	Number	A high value can indicate an increase in the load on the web proxy service.
	HTTP requests: The number of http requests that have been made to the web proxy service.	Number	A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications.
	HTTPS sessions: The total number of HTTP-Secured sessions serviced by the SSL tunnel	Number	A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications on the server.
	Thread pool active sessions: The number of sessions being actively served by the pool of threads	Number	A high value can indicate an increase in the load on the web proxy service.
	Thread pool failures: The number of requests rejected, since the thread pool was overcommitted	Number	The high value indicates a possible problem in the thread pool of the web proxy service.

	Upstream receive rate: The rate at which the data is received by the web proxy service from remote servers on the internet/proxy servers surrounding the current proxy server	Kb/sec	A high value can indicate an increase in the load on the web proxy service from one or more remote servers.
	Upstream transmit rate: The rate at which the data is sent by the web proxy service to remote servers on the internet/proxy servers surrounding the current proxy server	Kb/sec	A high value can indicate an increase in the load of one or more remote servers.

8.1.3 Proxy Cache Test

Web site caching is an efficient use of resources and another benefit of the MS proxy server. Since you can use the MS proxy server as a common connection point to the Internet, you can also use it to cache frequently accessed resources. The proxy server allocates a portion of the server's hard disk space to store frequently accessed objects. Internet requests are more efficiently responded to through the use of fresh-cached data, which in the long run, helps in minimizing internet response times.

Caching can either be passive or active. Passive caching just stores objects as they are requested, so the cache is updated only when users request information. Active caching directs the server to refresh objects in the cache automatically.

You can selectively control the proxy server caching so that you can limit the size of cached objects, change the expiration limits (control the freshness of objects), and determine whether the server always caches, or always excludes from cache, certain content.

This test reports the performance statistics pertaining to this caching activity of the MS Proxy server.

Purpose	Reports the performance statistics pertaining to this caching activity of the MS Proxy server		
Target of the test	An MS Proxy Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the MS Proxy server		
Outputs of the test	One set of results for every web proxy server cache monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Active refreshes: The rate at which data is retrieved from the Internet to refresh popular URLs in the URL cache.	Kb/sec	A low value indicates low refresh rate and a possible network problem.
	Active URL refreshes: The rate at which the URLs in the URL cache are refreshed from the internet	URSs/sec	A low or 0(zero) indicates the non-availability of URLs or DNS servers from the internet.
	Cache size: The total number of bytes currently available in the URL Cache	Kb	A high value indicates possible high usage of virtual memory on web proxy cache.
	URL commits: The rate at which the URLs are committed to the URLs cache	URLs/sec	The low value or 0 (zero) indicates low URL commits, low network resource availability.
	URLs retrieved: The rate at which the URLs are retrieved from the URL cache.	URLs/sec	A low value indicates the low availability of the URLs from the proxy cache.
	URLs in cache The current number of URLs in the URL cache	Number	A high value indicates possible low availability of virtual memory.

8.1.4 Proxy Svc Test

This test can be executed from a location external to the proxy server, and presents an unbiased external perspective of the state of 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 *Microsoft Proxy* as the **Component type**, *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.

Purpose	This test measures the state of an MS proxy server
Target	An MS Proxy server
Agent deploying this test	An external agent executing on an eG server
Configurable parameters for this test	1. **TEST PERIOD** – How often should the test be executed 2. **URL** – The web page being accessed. While multiple URLs (separated by commas) can be provided, each URL should be of the format **URL name:URL value**. **URL name** is a unique name assigned to the URL, and the **URL value** is the value of the URL. For example, a URL can be specified as **HomePage:http://192.168.10.12:7077/**, where **HomePage** is the **URL name** and http://192.168.10.12:7077/ is the **URL value**.

<table>
<tr><td></td><td colspan="3">3. HOST - The host for which the test is to be configured.
4. PORT - The port to which the specified HOST listens
5. COOKIEFILE – Whether any cookies being returned by the MS Proxy server need to be saved locally and returned with subsequent requests
6. PROXYHOST – The host on which a web proxy server is running (in case a proxy server is to be used)
7. PROXYPORT – The port number on which the web proxy server is listening
8. PROXYUSERNAME – The user name of the proxy server
9. PROXYPASSWORD – The password of the proxy server
10. CONFIRM PASSWORD – Confirm the password by retyping it here.
11. CONTENT – Is a set of instruction:value pairs that are used to validate the content being returned by the test. If the CONTENT value is none:none, no validation is performed. The number of pairs specified in this text box, must be equal to the number of URLs being monitored. The instruction should be one of Inc or Exc. Inc tells the test that for the content returned by the MS Proxy server to be valid, the content must include the specified value (a simple string search is done in this case). An instruction of Exc instructs the test that the server's output is valid if it does not contain the specified value.
12. CREDENTIALS – The HttpTest supports HTTP authentication. The CREDENTIALS parameter is to be set if a specific user name / password has to be specified to login to a page. This parameter is a comma separated list of user name:password pairs, one pair for each URL being monitored. A value of none:none indicates that user authorization is not required. Please be sure to check if your web site requires HTTP authentication while configuring this parameter. HTTP authentication typically involves a separate pop-up window when you try to access the page. Many sites uses HTTP POST for obtaining the user name and password and validating the user login. In such cases, the username and password have to be provided as part of the POST information and NOT as part of the CREDENTIALS specification for the HttPTest.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of outputs for every URL being monitored</td></tr>
<tr><td>Measurements of the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Proxy service availability:
This measurement indicates whether the server was able to respond successfully to the query made by the test.</td><td>Percent</td><td>Availability failures could be caused by several factors such as the MS Proxy process(es) being down, the MS Proxy servers being misconfigured, a network failure, etc. Temporary unavailability may also occur if the proxy server is overloaded. Availability is determined based on the response code returned by the server. A response code between 200 to 300 indicates that the server is available.</td></tr>
</table>

	Total response time: This measurement indicates the time taken by the server to respond to the requests it receives.	Secs	Response time being high denotes a problem. Poor response times may be due to the server being overloaded or misconfigured. If the URL accessed involves the generation of dynamic content by the server, backend problems (e.g., an overload at the application server or a database failure) can also result in an increase in response time.
	TCP connection availability: This measure indicates whether the test managed to establish a TCP connection to the server.	Percent	Failure to establish a TCP connection may imply that either the MS proxy server process is not up, or that the process is not operating correctly. In some cases of extreme overload, the failure to establish a TCP connection may be a transient condition. As the load subsides, the server may start functioning properly again.
	TCP connection time: This measure quantifies the time for establishing a TCP connection to the MS proxy server host.	Secs	Typically, the TCP connection establishment must be very small (of the order of a few milliseconds). Since TCP connection establishment is handled at the OS-level, rather than by the application, an increase in this value signifies a system-level bottleneck on the host that supports the MS proxy server.
	Server response time: This measure indicates the time period between when the connection was established and when the server sent back a HTTP response header to the client.	Secs	While the total response time may depend on several factors, the server response time is typically, a very good indicator of a server bottleneck (e.g., because all the available server threads or processes are in use).
	Response code: The response code returned by the server for the simulated request	Number	A value between 200 and 300 indicates a good response. A 4xx value indicates a problem with the requested content (eg., page not found). A 5xx value indicates a server error.
	Content length: The size of the content returned by the server	Kbytes	Typically the content length returned by the server for a specific URL should be the same across time. Any change in this metric may indicate the need for further investigation on the server side.

	Content validity: This measure validates whether the server was successful in executing the request made to it.	Percent	A value of 100% indicates that the content returned by the test is valid. A value of 0% indicates that the content may not be valid. This capability for content validation is especially important for multi-tier web applications. For example, a user may not be able to login to the web site but the server may reply back with a valid HTML page where in the error message, say, "Invalid Login" is reported. In this case, the availability will be 100 % (since we got a valid HTML response). If the test is configured such that the content parameter should exclude the string "Invalid Login," in the above scenario content validity would have a value 0.

Chapter 9

Monitoring Windows Domain Controllers

Windows Domain Controllers are critical components of IT infrastructures. Users accessing resources in a Windows domain have to first be authenticated by the Domain Controller in order to get access. Any slowdown or failure of the domain controllers can severely impact users. Hence, 24x7 monitoring of domain controllers is critical.

The eG Enterprise suite provides a specialized *Domain Controller* monitoring model for the Windows domain controller (DC) (see Figure 9.1), using which key performance parameters related to the DC can be continuously monitored, and anomalies, instantly detected.

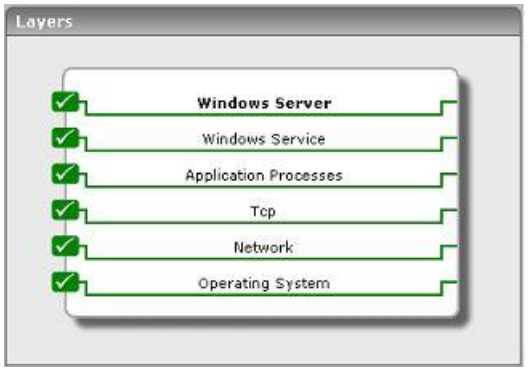

Figure 9.1: Layer model of a Windows Domain Controller

Each of the layers in this specialized model (see Figure 9.1) executes a wide variety of tests on the DC and extracts critical metrics, which help quantify the performance level achieved by the DC, and simplifies problem identification.

The *Monitoring Unix and Windows Servers* document deals extensively with the bottom 5 layers of Figure 9.1. In the section that follows, the **Windows Server** layer will be discussed.

9.1 The Windows Server Layer

Using the tests associated with this layer, administrators can gauge how effectively the DC authenticates login requests it receives.

Figure 9.2: Tests associated with the Windows Server layer

9.1.1 Windows Access Test

This test monitors the accesses to a Windows server.

Purpose	Monitors the accesses to the Windows server		
Target of the test	A Windows server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every Windows server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Blocking request rejects: The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items	Reqs/sec	If the number of blocking request rejects is high, you may need to adjust the MaxWorkItem or MinFreeWorkItems server parameters
	Permission errors: The number of times opens on behalf of clients have failed with STATUS_ACCESS_DENIED in the last measurement period	Number	Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected.

	File access denied errors: The number of times accesses to files opened successfully were denied in the last measurement period	Number	This number indicates attempts to access files without proper access authorization.
	Internal server errors: This value indicates the number of times an internal server error was detected in the last measurement period.	Number	Unexpected errors usually indicate a problem with the server.
	Data received: The rate at which the server has received data from the network	Kbytes/sec	This metric indicates how busy the server is.
	Data transmitted: The rate at which the server has sent data over the network	Kbytes/sec	This metric indicates how busy the server is.
	Resource shortage errors: The number of times STATUS_DATA_NOT_ACCEPTED was returned to clients in the last measurement period	Number	A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the InitWorkItems or MaxWorkItems server parameters might need to be adjusted.
	Avg response time: Average time taken by the server to respond to client requests	Secs	This is a critical measure of server health.

9.1.2 Windows Sessions Test

This test reports various session-related statistics for a Windows server.

Purpose	Reports various session-related statistics for a Windows server		
Target of the test	A Windows Domain Controller		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows server		
Outputs of the test	One set of results for every Windows server being monitored		
Measurements	**Measurement**	**Measurement Unit**	**Interpretation**

made by the test	**Logons:** Rate of logons to the server	Reqs/sec	This measure reports the rate of all interactive, network, and service logons to a windows server. The measure includes both successful and failed logons.
	Logon errors: Number of logons in the last measurement period that had errors	Number	This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server.
	Current sessions: The number of sessions currently active in a server	Number	This measure is one of the indicators of current server activity.
	Sessions with errors: The number of sessions in the last measurement period that were closed to unexpected error conditions	Number	Sessions can be closed with errors if the session duration reaches the autodisconnect timeout.
	Sessions forced off: The number of sessions in the last measurement period that have been forced to logoff	Number	This value indicates how many sessions were forced to logoff due to logon time constraints.
	Sessions logged off: The number of sessions in the last measurement period that were terminated normally	Number	Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low.
	Sessions timed out: The number of sessions that have been closed in the last measurement period due to their idle time exceeding the AutoDisconnect parameter for the server	Number	The number of session timed out gives an indication of whether the AutoDisconnect setting is helping to conserve server resources

9.1.3 Window Authentication Test

This test emulates a user logging into a Windows domain or local host and reports whether the login succeeded and how long it took.

Purpose	Emulates a user logging into a windows domain or a local host and reports whether the login succeeded and how long it took

<table>
<tr><td>Target of the test</td><td colspan="3">A Windows Domain Controller</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT – Refers to the port used by the Windows server
4. USERNAME and PASSWORD - This test emulates a user logging into a domain controller at the system-level. Therefore, specify the credentials of a user with both interactive logon and logon locally privileges against the USERNAME and PASSWORD fields
5. CONFIRM PASSWORD – Confirm the password by retyping it here.
6. 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/or multiple 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 Environments document.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every user account being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Authentication status:
Indicates whether the login was successful or not</td><td>Percent</td><td>A value of 100 % indicates that the login has succeeded. The value 0 is indicative of a failed login.</td></tr>
</table>

	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.

Chapter 10

Monitoring MS File Servers

In the client/server model, a file server is a computer responsible for the central storage and management of data files so that other computers on the same network can access the files. A file server allows users to share information over a network without having to physically transfer files. Any computer can be configured to be a host and act as a file server. In its simplest form, a file server may be an ordinary PC that handles requests for files and sends them over the network. In a more sophisticated network, a file server might be a dedicated network-attached storage device that also serves as a remote hard disk drive for other computers, allowing anyone on the network to store files on it as if to their own hard drive.

The true indicator of the efficiency of a File server is the speed with which it serves concurrent file requests. If users are unable to access important files stored on the file server as and when they need due to a temperorary break in connection to the server or because of a long request queue, it might severely hamper the productivity of the users, and might unnecessarily delay critical operations. If such a problem situation is to be averted, the file server needs to be monitored, and administrators promptly warned about probable performance issues.

eG Enterprise provides out-of-the-box a specialized *Microsoft File* server model (see Figure 10.1) that periodically runs diagnostic tests on the file server to ensure that it performs to peak capacity at all times.

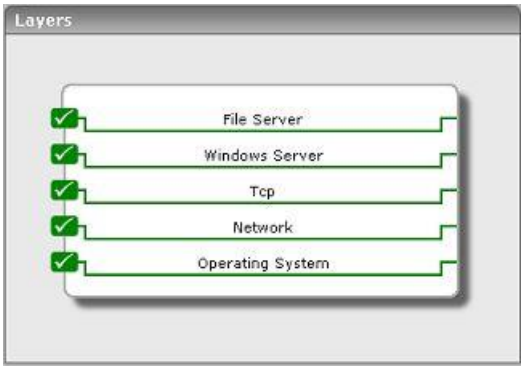

Figure 10.1: Layer model of an MS File server

The sections to come discuss the top 2 layers of Figure 10.1, since the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

10.1 The Windows Server Layer

Using the tests associated with the **Windows Server** layer, administrators can closely observe the user logins to and session behavior on the MS File server.

Figure 10.2: Tests associated with the Windows Server layer

10.1.1 Windows Access Test

This test monitors the accesses to the MS File server.

Purpose	Monitors the accesses to the MS File server		
Target of the test	An MS File server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the File server		
Outputs of the test	One set of results for every File server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Blocking request rejects: The number of times in the last measurement period that the server has rejected blocking requests due to insufficient count of free work items	Reqs/sec	If the number of blocking request rejects is high, you may need to adjust the MaxWorkItem or MinFreeWorkItems server parameters

	Permission errors: The number of times opens on behalf of clients have failed with STATUS_ACCESS_DENIED in the last measurement period	Number	Permission errors can occur if any client/user is randomly attempting to access files, looking for files that may not have been properly protected.
	File access denied errors: The number of times accesses to files opened successfully were denied in the last measurement period	Number	This number indicates attempts to access files without proper access authorization.
	Internal server errors: This value indicates the number of times an internal server error was detected in the last measurement period.	Number	Unexpected errors usually indicate a problem with the server.
	Data received: The rate at which the server has received data from the network	Kbytes/sec	This metric indicates how busy the server is.
	Data transmitted: The rate at which the server has sent data over the network	Kbytes/sec	This metric indicates how busy the server is.
	Resource shortage errors: The number of times STATUS_DATA_NOT_ACCEPTED was returned to clients in the last measurement period	Number	A resource shortage event occurs when no work item is available or can be allocated to service the incoming request. If many repeated resource shortage events occur, the InitWorkItems or MaxWorkItems server parameters might need to be adjusted.
	Avg response time: Average time taken by the server to respond to client requests	Secs	This is a critical measure of server health.

10.1.2 Windows Sessions Test

This test reports various session-related statistics for an MS File server.

Purpose	Reports various session-related statistics for an MS File server
Target of the test	An MS File server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the MS File server		
Outputs of the test	One set of results for every MS File server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Logons: Rate of logons to the server	Reqs/sec	This measure reports the rate of all interactive, network, and service logons to an MS File server. The measure includes both successful and failed logons.
	Logon errors: Number of logons in the last measurement period that had errors	Number	This measure reports the number of failed logon attempts to the server during the last measurement period. The number of failures can indicate whether password-guessing programs are being used to get into the server.
	Current sessions: The number of sessions currently active in a server	Number	This measure is one of the indicators of current server activity.
	Sessions with errors: The number of sessions in the last measurement period that were closed to unexpected error conditions	Number	Sessions can be closed with errors if the session duration reaches the autodisconnect timeout.
	Sessions forced off: The number of sessions in the last measurement period that have been forced to logoff	Number	This value indicates how many sessions were forced to logoff due to logon time constraints.
	Sessions logged off: The number of sessions in the last measurement period that were terminated normally	Number	Compare the number of sessions logged off to the number of sessions forced off, sessions with errors, or those that timed out. Typically, the percentage of abnormally terminated sessions should be low.
	Sessions timed out: The number of sessions that have been closed in the last measurement period due to their idle time exceeding the AutoDisconnect parameter for the server	Number	The number of session timed out gives an indication of whether the AutoDisconnect setting is helping to conserve server resources

10.2 The File Server Layer

With the help of the tests associated with this layer, administrators can:

- Accurately determine the current work load on the server in terms of the number of files currently accessed on the server and the current user traffic to the server
- Quickly identify locked files and the users who have acquired a lock on those files

Figure 10.3: Tests associated with the File server layer

10.2.1 MS File Stats Test

The MsFileTest tracks various statistics pertaining to open file connections at the host.

Purpose	Tracks various statistics pertaining to open file connections at the host
Target of the test	An MS File server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **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.
Outputs of the test	One set of results for every MS file server monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	File locks count: The number of files locked at the host	Number	A high value can indicate too many files being opened at the host. The detailed diagnosis of this measure, if enabled, lists the files that have been locked, the user who holds the lock, and the number of locks on the file.
	Unique users count: A unique count of users who have opened files at this host	Number	A high value can indicate too many users connected to the host.

10.2.2 Windows Usage Test

This test tracks various statistics pertaining to sessions open at the host.

Purpose	Tracks various statistics pertaining to sessions open at the host		
Target of the test	An MS File server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **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.		
Outputs of the test	One set of results for every MS file server monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Open files:** The number of files opened over the network by users connecting to the file server	Number	This measurement is an indicator of the workload on the file server. The detailed diagnosis of this measure, if enabled, provides the number of open sessions for every user, and the time for which the sessions have been idle. If the idle time displayed here is very high, then measures for closing the inactive open sessions can be initiated.
	Unique users: A unique count of users who have opened sessions at this host	Number	A high value can indicate too many users connected to the host.

Chapter 11

Monitoring ISA Proxy Servers

Microsoft Internet Security and Acceleration (ISA) Server can be deployed as a dedicated firewall that acts as the secure gateway to the Internet for internal clients. ISA Server protects all communication between internal computers and the Internet. In a simple firewall scenario, the ISA Server computer has two network interface cards, one connected to the local network and one connected to the Internet. By setting the security access policies, you prevent unauthorized access and malicious content from entering the network. You can also restrict what traffic is allowed for each user and group, application, destination, content type, and schedule.

To assure users of safe and secure access to the Internet, and to shield the network from malicious attacks, the availability and internal health of the ISA Proxy server should be constantly monitored.

The eG Enterprise suite's unique *ISA Proxy* monitoring model (see Figure 11.1) executes a wide variety of tests on the proxy server to enable administrators to determine the following:

- Does the server take too much time to service firewall requests?
- Is the Web Proxy server fuctioning optimally?
- Is the Web Proxy Cache utilized effectively?

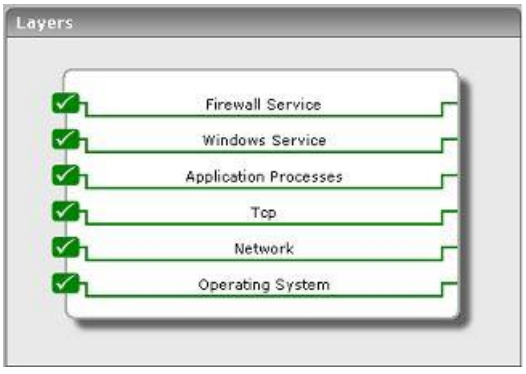

Figure 11.1: Layer model of an ISA Proxy server

The sections to come will discuss the tests associated with the **Firewall Service** layer only, since the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

11.1 The Firewall Service Layer

The tests associated with the **Firewall Service** layer monitor various critical firewall services provided by the ISA Proxy server.

Figure 11.2: The tests associated with the Firewall Service layer

11.1.1 ISA Cache Test

This test reports statistics pertaining to the ISA Proxy server cache.

Purpose	Reports statistics pertaining to the ISA Proxy server cache		
Target of the test	An ISA Proxy server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Data received from disk cache: Indicates the rate at which data is retrieved from the disk cache.	KB/Sec	

	Data received from memory cache: Indicates the rate at which data is retrieved from the memory cache.	KB/Sec	
	Disk failures: Indicates the rate at which I/O failures occurred since the Firewall service started.	Fails/Sec	An I/O failure occurs when the ISA server fails to read from or write to disk cache.
	Disk writes: Indicates the rate at which data was written to the disk cache.	Writes/Sec	
	Memory cache util: Indicates the percentage of fetches made from the memory.	Percent	A high percentage may indicate that it is worthwhile allocating more available memory resources to the cache.
	URLs in cache: Indicates the number of URLs currently stored in the cache.	Number	

11.1.2 ISA Firewall Test

This test reports statistics pertaining to the Firewall service of the ISA Proxy server 2004.

Purpose	Measures the firewall protection of the ISA proxy server		
Target of the test	An ISA Proxy server 2004		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	DNS cache hit ratio: Indicates the percentage of DNS domain names retrieved from the DNS cache.	Percent	

	Pending DNS resolutions: Indicates the number of gethostbyname and gethostbyaddr API calls pending resolution. These are calls used to resolve host DNS domain names and IP addresses for Firewall Service connections.	Number	
	Worker threads: Indicates the number of Firewall Service worker threads that are currently alive.	Number	A high value indicates that the current workload of the ISA Proxy Server is very high.

11.1.3 ISA Web Proxy Test

This test monitors the performance of the Web proxy service of the ISA Proxy server 2004.

Purpose	Measures the firewall protection of the ISA proxy server		
Target of the test	An ISA Proxy server 2004		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active web sessions: Indicates the number of active Web sessions currently connected to the ISA proxy Server.	Number	
	Connect errors: Indicates the total number of errors that occurred while connecting.	Number	

	DNS cache hit ratio: Indicates the percentage of DNS domain names served from the DNS cache.	Percent	
	Failed requests: Indicates the rate of requests that have failed because of some type of error.	Conns/Sec	A high failure rate, in comparison to the rate of incoming requests, will suggest that the ISA Proxy server is having difficulty in coping with all incoming requests. Connection settings for incoming Web requests may be incorrectly configured, or connection bandwidth may be insufficient.
	Inbound connections: Indicates the rate of incoming connections.	Conns/Sec	
	Outbound connections: Indicates the rate of outgoing connections.	Conns/Sec	
	Requests rate: Indicates the rate of requests to the Web Proxy filter.	Requests/Sec	A higher value means that more ISA Proxy server resources will be required to service incoming requests.

11.1.4 Packet Engine Test

The PacketEngine test reports statistics relating to the firewall packet engine of the ISA Proxy server 2004.

Purpose	Reports statistics relating to the firewall packet engine of the ISA Proxy server 2004		
Target of the test	An ISA Proxy server 2004		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active connections: Indicates the number of active connections currently transmitting data.	Number	A high value indicates that the current workload of the ISA Proxy Server is very high.

	Allowed packets rate: Indicates the number of packets allowed per second.	Packets/Sec	
	Data sent rate: Indicates the rate at which data was transmitted by the firewall packet engine driver.	KB/Sec	
	Dropped packets rate: Indicates the rate at which packets were dropped.	Packets/Sec	
	New connections rate: Indicates the rate at which connections were created.	Conns/Sec	
	Packets inspected rate: Indicates the rate at which the firewall packet engine driver inspects the packets.	Packets/Sec	

11.1.5 Proxy Server Test

The Web Proxy service provides support for HTTP (a.k.a. Web publishing), FTP, Gopher, and secure (SSL) communications. The Web Proxy service works with any CERN-compliant Web browser, such as Internet Explorer or Netscape Navigator. Because the Web Proxy supports only these widely adopted Internet standard communication methods, it isn't operating system dependent. Clients running Unix, Macintosh, or Windows operating systems can communicate with the Web Proxy service as long as they're configured with a CERN-compliant Web browser.

This test reports the performance statistics pertaining to this Web Proxy service running on an ISA Proxy server.

Purpose	Reports performance statistics pertaining to the Web Proxy service running on an ISA Proxy server		
Target of the test	An ISA Proxy Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every web proxy service monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Cache hit ratio The percentage of requests that have used cached data, to the total number of requests to the web proxy service	Percent	A high value could indicate an increase in the proxy server load, due to which lesser TCP connection requests are accepted.
	Client data receive rate: The number of active sessions for the web proxy service	Number	A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications.
	Client data transmit rate: The rate at which the data bytes are sent by the proxy server to the web proxy clients	Kb/sec	A high value could indicate a high data transfer from the proxy server to the web proxy client, which may result in congestion in network traffic
	Avg response time: The mean response time in seconds to service a request	Secs/req	High network traffic, low server performance are some of the factors that cause this measure to increase.
	Current users: The current number of users connected to the web proxy service.	Number	A high value can indicate an increase in the load on the web proxy service.
	DNS cache hits: This measure give the percentage of DNS domain names served from the proxy server cache, from the total DNS entries that are retrieved by the web proxy service.	Percent	A high value can indicate an increase in load on web proxy service.
	Failing requests: The rate of request that have completed with some error.	Reqs/Sec	The high value indicates possible problems in the web proxy service.
	FTP requests: The number of ftp requests that have been made to the web proxy service	Number	A high value can indicate an increase in the load on the web proxy service.
	HTTP requests: The number of http requests that have been made to the web proxy service.	Number	A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications.
	HTTPS sessions: The total number of HTTP-Secured sessions serviced by the SSL tunnel	Number	A high value can indicate an increase in the load on one or more applications, or a change in the characteristics of one or more applications on the server.

	Thread pool active sessions: The number of sessions being actively served by the pool of threads	Number	A high value can indicate an increase in the load on the web proxy service.
	Thread pool failures: The number of requests rejected, since the thread pool was overcommitted	Number	The high value indicates a possible problem in the thread pool of the web proxy service.
	Upstream receive rate: The rate at which the data is received by the web proxy service from remote servers on the internet/proxy servers surrounding the current proxy server	Kb/sec	A high value can indicate an increase in the load on the web proxy service from one or more remote servers.
	Upstream transmit rate: The rate at which the data is sent by the web proxy service to remote servers on the internet/proxy servers surrounding the current proxy server	Kb/sec	A high value can indicate an increase in the load of one or more remote servers.

11.1.6 Tests that are Disabled by Default

In addition to the tests discussed above, the **Firewall Service** layer is also mapped to a few tests that are disabled by default. To enable the test, go to the **ENABLE / DISABLE TESTS** page using the menu sequence : Agents -> Tests -> Enable/Disable, pick *ISA Proxy* as the **Component type**, *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.

The tests that are disabled by default have been discussed in the following sections.

11.1.6.1 Firewall Service Test

The FirewallService test measures the firewall protection of the ISA proxy server.

Purpose	Measures the firewall protection of the ISA proxy server
Target of the test	An ISA Proxy server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active sessions: Indicates the number of active sessions for the firewall service.	Number	Comparing this measure at both peak and off-peak times will provide you with valuable insight into the usage patterns of the ISA server.
	Active TCP connections: Indicates the number of active TCP connections transmitting data.	Number	
	Active UDP conections: Indicates the number of active UDP connections for the firewall service.	Number	
	Active threads: Indicates the number of firewall worker threads that are currently active.	Number	
	Read rate: Indicates the number of kilobytes read by the data-pump per second.	KB/Sec	A consistent decrease in the value of this measure may indicate a delay in servicing firewall requests.
	Write rate: Indicates the number of kilobytes written by the data-pump per second.	KB/Sec	A consistent decrease in the value of this measure may indicate a delay in servicing firewall requests.

11.1.6.2 Web Proxy Service Test

This test monitors the Web Proxy service. Requests from Web Proxy clients are directed to the Web Proxy service on the ISA server to determine if access is allowed.

Purpose	Monitors the Web Proxy service
Target of the test	An ISA Proxy server
Agent	An internal agent

deploying the test			
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Cache hit ratio: The percentage of successful web proxy client requests to the ISA Server.	Percent	This measure is a good indicator of the effectiveness of the cache. A higher percentage indicates that a number of requests are being serviced from the cache. This in turn is indicative of faster responsiveness. A zero value indicates that caching is not enabled, and a low value may indicate a configuration problem.
	Current users: Indicates the number of clients that are currently running the web proxy service.	Number	Monitoring this measure at both peak and off-peak times will enable users to assess the extent of server usage. This measure may also be useful if you need to temporarily stop ISA Server services.
	Read rate: Indicates the rate at which data bytes are received from Web Proxy clients.	KB/Sec	A consistent decrease in the value of this measure may indicate a delay in servicing requests.
	Active threads: Indicates the rate at which data bytes are sent to Web Proxy clients.	KB/Sec	A consistent decrease in the value of this measure may indicate a delay in servicing requests.
	Avg requests sec: Indicates the number of kilobytes read by the data-pump per second.	KB/Sec	A consistent decrease in the value of this measure may indicate a delay in servicing firewall requests.
	Write rate: Indicates the average amount of time required by the ISA server to process a request.	Secs/Request	This measure can be monitored at peak and off-peak times to receive a clear idea about how fast client requests are being serviced. A very high value of this measure might indicate that the ISA Server is having difficulty in handling all requests.
	Thread pool size: Indicates the number of threads in the thread pool	Number	

	Thread pool sessions: Indicates the number of sessions being actively serviced by thread pool threads.	Number	
	Thread pool failures: Indicates the number of requests rejected because the thread pool was full.	Number	

11.1.6.3 Web Proxy Cache Test

This test monitors the Web Proxy cache. The ISA server implements a cache of frequently-requested objects to improve network performance. You can configure the cache to ensure that it contains the data that is most frequently used by the organization or accessed by your Internet clients.

Purpose	Monitors the Web Proxy cache		
Target of the test	An ISA Proxy server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT –** Refers to the port used by the ISA Proxy server		
Outputs of the test	One set of results for every ISA Proxy server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Disk cache space: Indicates the amount of space used by the disk cache.	KB	If the value of this measure grows closer to or equal to the allocated disk cache space, it would indicate that subsequent cache requests might be rejected due to non-availability of adequate cache space. This, in turn, would increase the rate of direct disk accesses, which will consequently degrade system performance.
	Memory cache space: Indicates the amount of space used by the memory cache	KB	An excessive consumption of the memory cache space would result in slow-down of the system. This is because the lack of sufficient space in the memory cache would cause the real memory (RAM) to directly service the requests for objects.

	URL commit rate: Indicates the speed at which URLs are being written to the cache.	URLs/Sec	

Chapter

12

Monitoring Microsoft Radius Servers

NPS (Network Policy Server) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server, and as such, it performs connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections. NPS also functions as a health evaluation server for NAP (Network Access Protection).

The following illustration shows NPS as a RADIUS server for a variety of access clients and a RADIUS proxy. NPS uses an Active Directory domain for user credential authentication of incoming RADIUS Access-Request messages.

When NPS is used as a RADIUS server, RADIUS messages provide authentication, authorization, and accounting for network access connections in the following way:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the NPS server.
3. The NPS server evaluates the Access-Request message.
4. If required, the NPS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the NPS server.
5. The user credentials are checked and the dial-in properties of the user account are obtained by using a secure connection to a domain controller.
6. The connection attempt is authorized with both the dial-in properties of the user account and remote access policies.
7. If the connection attempt is both authenticated and authorized, the NPS server sends an Access-Accept message to the access server. If the connection attempt is either not authenticated or not authorized, the NPS server sends an Access-Reject message to the access server.
8. The access server completes the connection process with the access client and sends an Accounting-Request message to the NPS server, where the message is logged.
9. The NPS server sends an Accounting-Response to the access server.

Issues in the functioning of NPS, if not promptly isolated and resolved, might result in the complete collapse of the

remote authentication and authorization service provided by the Windows server. 24x7 monitoring of NPS, hence becomes imperative.

The eG Enterprise suite provides out-of-the-box monitoring support to the Windows Internet Authentication Service, and proactively alerts administrators of authentication, authorization, or accounting bottlenecks encountered by the NPS server. The specialized *Microsoft Radius* monitoring model (see Figure 12.1) offered by the eG Enterprise suite executes a variety of tests on the NPS server; these tests, in turn, use the perfmon utility of Windows to extract critical performance statistics pertaining to the services offered by the NPS server.

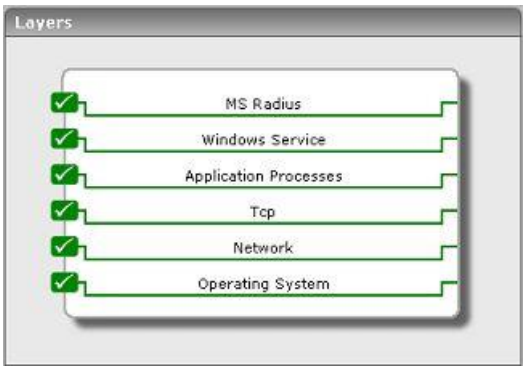

Figure 12.1: The layer model of the MS Radius server

This section will discuss the **MS Radius** layer alone, as all the other layers have been elaborately discussed in the *Monitoring Unix and Windows Servers* document.

12.1 The MS Radius Layer

This layer monitors the authentication, authorization, and accounting activities performed by the NPS server and clients.

Figure 12.2: The tests associated with the MS Radius layer

12.1.1 NPS Accounting Server Test

Besides providing remote authentication services to clients, the NPS also provides a central accounting recording service for all accounting requests that are sent by the clients. Once the NPS server completes the connection process initiated by a client, the access server which processed the connection request sends an Accounting-Request message to the NPS server, where the message is logged. The NPS then sends an Accounting-Response to the access server. In addition, the access server also sends Accounting-Request messages for the following:

- During the time in which the connection is established
- When the access client connection is closed
- When the access server is started and stopped

This test monitors the accounting-requests received and accounting-responses sent by the NPS to clients.

Purpose	Monitors the accounting-requests received and accounting-responses sent by the NPS to clients		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for every NPS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Packets sent: Indicates the rate at which packets were sent by the NPS server.	Packets/Sec	
	Packets received: Indicates the rate at which the NPS server received packets.	Packets/Sec	When viewed along with the *Packets sent* measure, this measure serves as a good indicator of the traffic on the server.
	Packets dropped: Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.

	Invalid requests: Indicates the rate at which packets were received from an unknown address.	Reqs/Sec	
	Malformed packets: Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
	Unknown packets: Indicates the rate at which packets of an unknown type were received.	Packets/Sec	
	No record packets: Indicates the rate at which RADIUS Accounting-Request packets were received and responded to but not recorded.	Records/Sec	
	Accounting requests: Indicates the rate at which RADIUS Accounting-Requests were received from this client on the accounting port.	Reqs/Sec	
	Accounting responses: Indicates the rate at which RADIUS Accounting-Response packets were sent to this client on the accounting port.	Reqs/Sec	The *Accounting requests* and *Accounting responses* measures serve as effective indicators of the workload on the NPS server.
	Duplicate requests: Indicates the rate at which duplicate RADIUS Accounting-Request packets were received from this client.	Reqs/Sec	
	Bad authenticators: Indicates the rate at which Accounting-Requests containing invalid signature attributes were received.	Reqs/Sec	

12.1.2 NPS Accounting Client Test

This test monitors the accounting-requests sent and accounting-responses received by the RADIUS clients from the NPS servers.

Purpose	Monitors the accounting-requests sent and accounting-responses received by the RADIUS clients from the NPS servers		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for every NPS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Packets sent: Indicates the rate at which packets were sent by this client.	Packets/Sec	
	Packets received: Indicates the rate at which this RADIUS client received packets.	Packets/Sec	When viewed along with the *Packets sent* measure, this measure serves as a good indicator of the traffic that originated from a client.
	Packets dropped: Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
	Malformed packets: Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
	Unknown packets: Indicates the rate at which packets of an unknown type were received.	Packets/Sec	

	No record packets: Indicates the rate at which RADIUS Accounting-Request packets were received and responded to but not recorded.	Records/Sec	
	Accounting requests: Indicates the rate at which RADIUS Accounting-Requests were sent by this client on the accounting port.	Reqs/Sec	
	Accounting responses: Indicates the rate at which RADIUS Accounting-Response packets were sent to this client on the accounting port.	Reqs/Sec	The *Accounting requests* and *Accounting responses* measures serve as effective indicators of the workload on the client.
	Duplicate requests: Indicates the rate at which duplicate RADIUS Accounting-Request packets were received from this client.	Reqs/Sec	
	Bad authenticators: Indicates the rate at which Accounting-Requests containing invalid signature attributes were received.	Reqs/Sec	

12.1.3 NPS Authentication Server Test

When NPS is used as a RADIUS server, it provides the a central authentication and authorization service for all access requests that are sent by RADIUS clients. NPS uses either a Microsoft® Windows NT® Server 4.0 domain, an Active Directory® domain, or the local Security Accounts Manager (SAM) to authenticate user credentials for a connection attempt. NPS uses the dial-in properties of the user account and remote access policies to authorize a connection.

This test measures how well the NPS server performs remote authentication and authorization.

Purpose	Measures how well the NPS server performs remote authentication and authorization
Target of the test	An NPS server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.

Outputs of the test	One set of results for every NPS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Packets sent: Indicates the rate at which packets were sent by the NPS server.	Packets/Sec	
	Packets received: Indicates the rate at which the NPS server received packets.	Packets/Sec	When viewed along with the *Packets sent* measure, this measure serves as a good indicator of the traffic on the server.
	Packets dropped: Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
	Invalid requests: Indicates the rate at which packets were received from an unknown address.	Reqs/Sec	
	Malformed packets: Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
	Unknown packets: Indicates the rate at which packets of an unknown type were received.	Packets/Sec	
	Access accepts: Indicates the rate at which RADIUS Access-Accept packets were sent by the NPS server to this client.	Accepts/Sec	
	Access challenges: Indicates the rate at which Access-Challenge messages are being processed.	Challenges/Sec	

	Access rejects: Indicates the rate at which Access-Reject messages are being processed.	Rejects/Sec	A very high value of this measure could warrant a review of the remote access policies.
	Access requests: Indicates the rate at which packets were received on an authentication port from this client.	Reqs/Sec	
	Duplicate requests: Indicates the rate at which duplicate RADIUS Access-Request packets were received from this client.	Reqs/Sec	

12.1.4 NPS Authentication Client Test

This test monitors the access-requests sent by access-clients to the NPS server, and indicates how many requests were accepted/rejected by the NPS server.

Purpose	Monitors the access-requests sent by access-clients to the NPS server, and indicates how many requests were accepted/rejected by the NPS server		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for every NPS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Packets sent: Indicates the rate at which packets were sent by this client to the NPS server.	Packets/Sec	
	Packets received: Indicates the rate at which this client received packets from the NPS server.	Packets/Sec	When viewed along with the *Packets sent* measure, this measure serves as a good indicator of the traffic on the client.

	Packets dropped: Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
	Malformed packets: Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
	Unknown packets: Indicates the rate at which packets of an unknown type were received.	Packets/Sec	
	Access accepts: Indicates the rate at which RADIUS Access-Accept packets were sent to this client.	Accepts/Sec	
	Access challenges: Indicates the rate at which Access-Challenge messages are being processed.	Challenges/Sec	
	Access rejects: Indicates the rate at which Access-Reject messages are being processed.	Rejects/Sec	A very high value of this measure could warrant a review of the remote access policies.
	Access requests: Indicates the rate at which packets were received on an authentication port from this client.	Reqs/Sec	
	Bad authenticators: Indicates the rate at which packets containing invalid signature attributes were received.	Reqs/Sec	

	Duplicate requests: Indicates the rate at which duplicate RADIUS Access-Request packets were received from this client.	Reqs/Sec	

12.1.5 NPS System Health Validators Test

NPS (Network Policy Server) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server, and as such, it performs connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections. NPS also functions as a health evaluation server for NAP (Network Access Protection).

System health validators (SHVs) in an NPS are server software counterparts to system health agents (SHAs) on NAP (Network Access Protection)-capable client computers. Each SHA on the client has a corresponding SHV in Network Policy Server (NPS). SHVs allow NPS to verify the statement of health (SoH) that is made by its corresponding SHA on the client computer. SHVs contain the details of the required configuration settings on client computers. For example, the Windows Security Health Validator (WSHV) is the counterpart to the Windows Security Health Agent (WSHA) on client computers. WSHV allows you to create a policy for the way in which settings on Network Access Protection (NAP)-capable client computers must be configured. If the settings on the client computer as reported in the SoH do not match the settings in the SHV on the server running NPS, it implies that the client computer is not compliant with the health policy requirements of the server. Once the system health validator validates the SoH from the client as either compliant or non-compliant, it marks the SoH with the relevant compliance status and sends it to the NPS.

By monitoring the statements of health issued by each system health validator, administrators can quickly capture non-compliances, investigate the reasons for the same, and can either fix it at the client side or fine-tune the access policies configured on the NPS to ensure secure access. This is exactly what the **NPS System Health Validators** test does. For each system health validator on NPS, this test reports the rate of compliances and non-compliances reported by that system health validator, thus shedding light on validations that often resulted in non-compliances. In addition, the test also reports the rate at which health statements could not be adjudged compliant/non-compliant, pinpoints the system health validators that sent out such statements, and reveals the reason for the same – is it owing to frequent server side failures? Or client side failures? Or is it because of other failures? The test also highlights 'slow' validators by measuring the responsiveness of every validator at pre-configured intervals.

Purpose	For each system health validator on NPS, this test reports the rate of compliances and non-compliances reported by that system health validator, thus shedding light on validations that often resulted in non-compliances. In addition, the test also reports the rate at which health statements could not be adjudged compliant/non-compliant, pinpoints the system health validators that sent out such statements, and reveals the reason for the same – is it owing to frequent server side failures? Or client side failures? Or is it because of other failures? The test also highlights 'slow' validators by measuring the responsiveness of every validator at pre-configured intervals.
Target of the test	An NPS server
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD – How often should the test be executed
2. HOST – The host for which the test is to be configured
3. PORT - The port at which the NPS server listens. The default is NULL.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every system health validator on the NPS server monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Client-communication failures:
Indicates the average number of Client-Communication failures per second from this health validator.</td><td>Failures/Sec</td><td rowspan="2">When a system health validator is not able to provide a health status to the Network Policy Server because of an error condition, it sends a Failure Category and code to the Network Policy Server.
If the error is on the client side, the system health validator sends either a Client Component Failure Category or a Client Communication Failure Category.
The value of these measures therefore indicate the rate at which client side failures occurred rendering a system health validator unable to determine the compliance status of a health statement.
In the Configuration Manager System Health Validator properties on the Network Policy Server, errors tagged with these failure categories match to SHA not responding to NAP client and SHA unable to contact required services, respectively.
Upon receipt of such failure categories from the system health validator, the NPS, by default, matches them to a non-compliant status.</td></tr>
<tr><td>Client-component failures:
Indicates the average number of Client-Component failures per second from this health validator.</td><td>Failures/Sec</td></tr>
</table>

	Compliances: Indicates the rate at which compliant decsisions were issued to the NPS by this system health validator.	Decisions/Sec	This condition occurs when the client's compliant status is successfully validated by the System Health Validator point because all the following apply: • The statement of health is not older than the setting **Date created must be after**. • The statement of health is within the configured **Validity period**. • The client site is valid. • The client has used up-to-date Configuration Manager NAP policies. • A failure did not occur on either the Configuration Manager client or the System Health Validator point. A high value is desired for this measure.
	Non-compliances: Indicates the rate at which non-compliant decsisions were issued to the NPS by this system health validator.	Decisions/Sec	This condition occurs when one of these situations apply: • The statement of health is older than the setting **Date created must be after**. • The statement of health is not within the configured **Validity period**. • The client does not have up-to-date Configuration Manager NAP policies. • The client has returned a non-compliant status because it does not have applicable software updates by the **Effective Date** as defined in the Configuration Manager NAP policies. A low value is desired for this measure.

	None failures: Indicates the rate at which none failures were reported by this system health validator.	Failures/Sec	
	Other failures: Indicates the rate at which other failures were reported by this system health validator.	Failures/Sec	
	Server-communication failures: Indicates the rate at which server-communication failures were reported by this system health validator.	Failures/Sec	When a system health validator is not able to provide a health status to the Network Policy Server because of an error condition, it sends a Failure Category and code to the Network Policy Server. If the error is on the server side, the system health validator sends either a Server Component Failure Category or a Server Communication Failure Category. The value of these measures therefore indicate the rate at which server side failures occurred rendering a system health validator unable to determine the compliance status of a health statement. In the Configuration Manager System Health Validator properties on the Network Policy Server, errors tagged with these failure categories match to **SHV not responding** and **SHV unable to contact required services**, respectively. Upon receipt of such failure categories from the system health validator, the NPS, by default, matches them to a non-compliant status.
	Server-component failures: Indicates the rate at which server-component failures were reported by this system health validator.	Failures/Sec	
	Last round-trip time: Indicates the interval (in hundredths of a second) between the most recent request to this system health validator and its response.	Secs	A low value is desired for this measure. A high value indicates that the system health validator is taking too long to validate health statements from the clients. Compare the value of this measure across system health validators to identify the slowest/least responsive validators.

12.1.6 NPS Remote Authentication Server Test

NPS performs centralized authentication, authorization, and accounting for wireless, authenticating switch, remote access dial-up and virtual private network (VPN) connections. When NPS is used as a RADIUS server, it provides a central authentication and authorization service for all access requests that are sent by RADIUS clients. NPS uses a Microsoft® Windows NT® Server 4.0 domain, an Active Directory® Domain Services (AD DS) domain, or the local Security Accounts Manager (SAM) user accounts database to authenticate user credentials for connection attempts.

The authenticating and authorization process is as follows:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the NPS server.
3. The NPS server evaluates the Access-Request message.
4. If required, the NPS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the NPS server.
5. The user credentials are checked and the dial-in properties of the user account are obtained by using a secure connection to a domain controller.
6. The connection attempt is authorized with both the dial-in properties of the user account and network policies.
7. If the connection attempt is both authenticated and authorized, the NPS server sends an Access-Accept message to the access server.
8. If the connection attempt is either not authenticated or not authorized, the NPS server sends an Access-Reject message to the access server.

If NPS challenges access requests frequently or rejects requests very often, administrators need to be instantly notified of this, so that they can look into these aberrations and uncover their reasons. Likewise, administrators should also rapidly capture any unusual delay in request authentication by NPS, so that they can swiftly determine and fix the reason for the delay. For this, administrators should periodically run the **NPS Remote Authentication Server** test. This test tracks the Access-Request messages sent by every access server configured to use NPS for authentication, and reports the rate at which these access requests are challenged/rejected by NPS. In addition, the test reveals the time taken by NPS to authenticate requests to every server, thus proactively alerting administrators to potential slowdowns in authentication. The rate at which access requests to a server are enqueued on NPS pending processing is also revealed, so that administrators are informed of bottlenecks in authentication.

Purpose	Tracks the Access-Request messages sent by every access server configured to use NPS for authentication, and reports the rate at which these access requests are challenged/rejected by NPS. In addition, the test reveals the time taken by NPS to authenticate requests to every server, thus proactively alerting administrators to potential slowdowns in authentication. The rate at which access requests to a server are enqueued on NPS pending processing is also revealed, so that administrators are informed of bottlenecks in authentication
Target of the test	An NPS server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for every access server that is configured to use NPS for authentication		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Access-Accepts: Indicates the rate at which RADIUS Access-Accept packets were received by this server from NPS.	Accepts/Sec	This is a good indicator of how frequently access requests from clients to a server are authenticated and authorized by NPS.
	Access-Challenges: Indicates the rate at which RADIUS Access-Challenge packets were sent by NPS to this server.	Challenges/Sec	A low value is desired for this measure. A high value indicates that NPS challenged many access requests, forcing the access server to send an updated Access-Request to NPS. In such cases, access clients are bound to experience delays in accessing the server.
	Access-Rejects: Indicates the rate at which RADIUS Access-Reject packets were sent by NPS to this server.	Rejects/Sec	Ideally, the value of this measure should be 0 or very low. A high value indicates too many or too frequent request rejections, which in turn may cause access clients to be denied access to the server.
	Access-Requests: Indicates the rate at which Access-Request packets were sent by this server to NPS.	Reqs/Sec	This is a good indicator of the load on NPS.
	Bad authenticators: Indicates the rate at which this server sent access requests containing an invalid Message Authenticator attribute to NPS.	Reqs/Sec	Ideally, the value of this measure should be 0.
	Packets dropped: Indicates the rate at which request packets sent by this server were silently discarded by NPS for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".	Packets/Sec	Ideally, the value of this measure should be 0.

	FullAccess-Decisions: Indicates the rate at which Full-access decisions were received from this server.	Decisions/Sec	NPS grants an access client full access if the client meets the defined health policies.
	Malformed packets: Indicates the rate at which NPS received malformed packets from this server.	Packets/Sec	Ideally, the value of this measure should be 0.
	Packets received: Indicates the rate at which requests packets were received from this server.	Packets/Sec	
	Probation-Decisions: Indicates the rate at which probation-decisions were received from this server.	Decisions/Sec	If NPS grants an access client full access but for a limited period only, the client is said to be on probation. This can happen if NPS finds that the client did not fulfill certain health policy requirements.
	Quarantine-Decisions: Indicates the rate at which quarantine decisions were sent by this server.	Decisions/Sec	When a remote access client dials in or connects via VPN to an access server, by default only the user's credentials (account name and password) are checked to determine whether access is granted. This means a computer that does not meet the network's policy requirements could still connect to the server and the network from a remote location. When quarantine control is deployed, after the user's credentials are authenticated the connection is "quarantined." In quarantine mode, the computer has an IP address and has limited access to some network resources (called quarantine resources) such as a DNS server and perhaps a file server or web server from which it can download files necessary to comply with the policies or where the user can get more information, but cannot access the rest of the network.
	Request timeouts: Indicates the rate at which requests to this server timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.

	Retransmissions: Indicates the rate at which requests were retransmitted to this server.	Reqs/Sec	Retransmits can increase the number of requests to NPS, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on NPS. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.
	Unknown type: Indicates the average number of unknown type (non-RADIUS) packets received by this server per second.	Packets/Sec	
	Last round-trip time: Indicates the interval (in hundredths of a second) between the most recent request to a remote NPS server and its response.	Secs	Ideally, the value of this measure should be very low. A high value indicates that that NPS is taking too long to authenticate requests.
	Pending requests: Indicates the rate of requests destined for this server that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on NPS or a high timeout setting. In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests.

12.1.7 NPS Remote Accounting Server Test

Network Policy Server (NPS) supports Remote Authentication Dial-In User Service (RADIUS) accounting, which you can use to track network usage for auditing and billing purposes. Accounting data can also be queried to assist with network access troubleshooting.

When a RADIUS client is configured to use RADIUS accounting, at the start of service delivery it generates an Accounting-Start message describing the type of service being delivered and the user it is being delivered to. The message is then sent to the RADIUS Accounting server, which sends back an acknowledgment to the RADIUS client. At the end of service delivery, the client generates an Accounting-Stop message describing the type of service that was delivered and optional statistics, such as elapsed time, input and output octets, or input and output packets. It then sends that data to the RADIUS accounting server, which sends back an acknowledgment to the RADIUS client.

The Accounting-Request message (whether for the Start or Stop message) is submitted to the RADIUS accounting server through the network. If the quality of this network connection is poor, then many request packets may be malformed, forcing the server to drop them. This in turn can delay or deny responses to accounting servers and acknowledgements to clients.

Also, if the accounting server is overloaded with requests, the server can choke slowing down accounting in the bargain.

To avoid such delays, administrators must track accounting requests and responses, proactively detect potential

slowdowns, accurately isolate what is causing it, and promptly fix it. This is where the **NPS Remote Accounting Server** test helps.

This test tracks the requests to and responses from each accounting server and reveals whether/not the servers are responding as quickly as the requests come in. You can also use this test to monitor the time each server takes to process requests, and thus identify the server that is experiencing a processing bottleneck. In addition, the test also captures the rate at which packets are dropped by the server and malformed/erroneous packets are received by the server, thus pointing to issues with the client or in the network connection between the server and the client. The load on each server is also revealed by monitoring the packets received by and the pending requests on the server from time to time. This way, the test pinpoints irregularities in load-balancing between servers, which in time can bottleneck request processing by the servers, resulting in slowdowns.

Purpose	Tracks the requests to and responses from each accounting server and reveals whether/not the servers are responding as quickly as the requests come in. You can also use this test to monitor the time each server takes to process requests, and thus identify the server that is experiencing a processing bottleneck. In addition, the test also captures the rate at which packets are dropped by the server and malformed/erroneous packets are received by the server, thus pointing to issues with the client or in the network connection between the server and the client. The load on each server is also revealed by monitoring the packets received by and the pending requests on the server from time to time. This way, the test pinpoints irregularities in load-balancing between servers, which in time can bottleneck request processing by the servers, resulting in slowdowns.		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for every access server that is configured to use NPS for authentication		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Accounting-Requests: Indicates the rate at which this server receives accounting requests.	Reqs/Sec	This is a good indicator of the workload on the server. Compare the value of this measure across servers to know which server is overloaded. If a single server appears to be handling a vast majority of the requests, you may want to consider sprucing up your load-balancing algorithm, so that the request load is uniformly balanced across servers. If required, you can even consider adding more servers.

	Accounting - Responses: Indicates the rate at which this server is responding to requests.	Reqs/Sec	If the value of this measure is much lower than the value of the *Accounting-Requests* measure, it could indicate that the server is not responding to requests quickly. You may want to investigate the reasons for the same.
	Bad authenticators: Indicates the rate at which this server received requests containing an invalid Message Authenticator attribute.	Reqs/Sec	Ideally, the value of this measure should be 0.
	Packets dropped: Indicates the rate at which this server were silently discarded the request packets it received for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".	Packets/Sec	Ideally, the value of this measure should be 0.
	Malformed packets: Indicates the rate at which this server received malformed packets.	Packets/Sec	Ideally, the value of this measure should be 0.
	Packets received: Indicates the rate at which requests packets were received by his server.	Packets/Sec	
	Request timeouts: Indicates the rate at which requests to this server timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.
	Retransmissions: Indicates the rate at which requests were retransmitted to this server.	Reqs/Sec	Retransmits can increase the number of requests to the server, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on the server. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.

	Unknown type: Indicates the average number of unknown type (non-RADIUS) packets received by this server per second.	Packets/Sec	
	Last round-trip time: Indicates the interval (in hundredths of a second) between the most recent request to this server and its response.	Secs	Ideally, the value of this measure should be very low. A high value indicates that that the accounting server is taking too long to perform accounting.
	Pending requests: Indicates the rate of requests destined for this server that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on the server or a high timeout setting (which could be causing many requests to be retransmitted to the server). In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests.

12.1.8 NPS Policy Engine Test

Every network policy must have at least one configured condition. NPS provides many conditions groups that allow you to clearly define the properties that the connection request received by NPS must have in order to match the policy. How quickly NPS matches requests with policies is a good measure of the efficiency of the NPS policy engine. Using the **NPS Policy Engine** test, administrators can measure just that! This test reports the time taken by NPS to process requests, the rate of pending requests on NPS, and the number of requests that matched configured policies. In the process, the test reveals processing bottlenecks on the NPS and how they impact policy matching.

Purpose	Reports the time taken by NPS to process requests, the rate of pending requests on NPS, and the number of requests that matched configured policies. In the process, the test reveals processing bottlenecks on the NPS and how they impact policy matching		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for the NPS monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Last round-trip time:** Indicates the interval (in hundredths of a second) between the most recent request to NPS and its response.	Secs	Ideally, the value of this measure should be very low. A high value indicates that NPS is taking too long to verify whether/not requests it receives match with the policy configuration.
	Matched remote access policy: Indicates the average number of remote access policies that have been matched with requests per second.	Number	
	Pending requests: Indicates the rate of requests destined for NPS that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on NPS or a high timeout setting (which could be causing many requests to be retransmitted to the NPS). In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests.

12.1.9 NPS Authentication Proxy Test

Network Policy Server (NPS) can be used as a RADIUS proxy to provide the routing of RADIUS messages between RADIUS clients access servers and RADIUS servers that perform user authentication, authorization, and accounting for the connection attempt. When used as a RADIUS proxy, NPS is a central switching or routing point through which RADIUS access and accounting messages flow.

Figure 12.3 shows NPS as a RADIUS proxy between RADIUS clients (access servers) and either RADIUS servers or another RADIUS proxy.

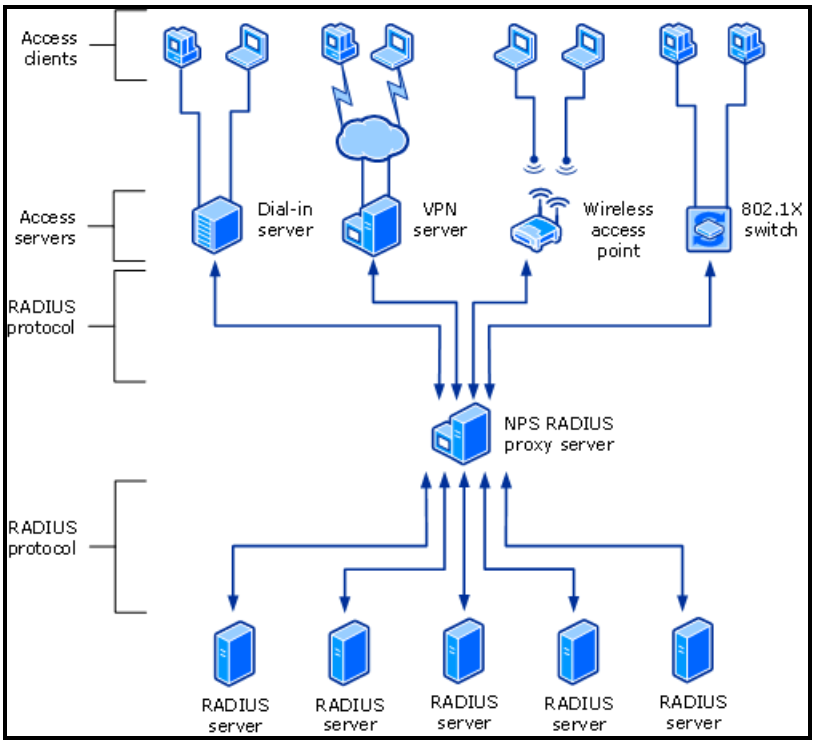

Figure 12.3: How NPS RADIUS Proxy works

When NPS is used as a RADIUS proxy between a RADIUS client and a RADIUS server, RADIUS messages for network access connection attempts are forwarded in the following way:

1. Access servers, such as dial-up network access servers, virtual private network (VPN) servers, and wireless access points, receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the NPS server that is being used as the NPS RADIUS proxy.
3. The NPS RADIUS proxy receives the Access-Request message and, based on the locally configured connection request policies, determines where to forward the Access-Request message.
4. The NPS RADIUS proxy forwards the Access-Request message to the appropriate RADIUS server.
5. The RADIUS server evaluates the Access-Request message.
6. If required, the RADIUS server sends an Access-Challenge message to the NPS RADIUS proxy, where it is forwarded to the access server. The access server processes the challenge with the access client and sends an updated Access-Request to the NPS RADIUS proxy, where it is forwarded to the RADIUS server.
7. The RADIUS server authenticates and authorizes the connection attempt.
8. If the connection attempt is both authenticated and authorized, the RADIUS server sends an Access-Accept message to the NPS RADIUS proxy, where it is forwarded to the access server.
9. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends an Access-Reject message to the NPS RADIUS proxy, where it is forwarded to the access server.

If the RADIUS server challenges or rejects connection requests frequently, the NPS RADIUS Proxy will transmit the same to the access servers. This in turn can cause many clients to be denied connections to access servers; some others may experience significant delays in connecting. Malformed request packets and those with invalid attributes/addresses can also be responsible for authentication delays/denials. Sometimes, a processing bottleneck on the NPS RADIUS Proxy server can also result in a slowdown in authentication. To enable clients to access remote services rapidly, administrators should keep a close watch on how the NPS RADIUS Proxy handles the requests and responses it receives, detect abnormalities rapidly, and quickly initiate measures to resolve them. The **NPS Authentication Proxy** test helps administrators do just that! This test keeps tabs on the access requests the NPS RADIUS Proxy receives from access servers and reports the rate at which the RADIUS server challenges / rejects these requests, thereby enabling administrators to instantly spot an abnormal number of challenges and rejections. Additionally, the test also reveals the rate at which erroneous request/response packets are received by the NPS RADIUS Proxy, thus providing administrators with effective pointers to what could be causing the high rate of challenges, rejects, or retransmisssions – is it because of malformed packets? packets with invalid authentication attributes? packets with invalid addresses? or non-RADIUS packets? The test also sheds light on the poor processing ability of the NPS RADIUS Proxy by reporting the number of pending requests on the proxy from time to time.

Purpose	Keeps tabs on the access requests the NPS RADIUS Proxy receives from access servers and reports the rate at which the RADIUS server challenges / rejects these requests, thereby enabling administrators to instantly spot an abnormal number of challenges and rejections. Additionally, the test also reveals the rate at which erroneous request/response packets are received by the NPS RADIUS Proxy, thus providing administrators with effective pointers to what could be causing the high rate of challenges, rejects, or retransmisssions – is it because of malformed packets? packets with invalid authentication attributes? packets with invalid addresses? or non-RADIUS packets? The test also sheds light on the poor processing ability of the NPS RADIUS Proxy by reporting the number of pending requests on the proxy from time to time.		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for the NPS RADIUS Proxy		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Access-Accepts: Indicates the rate at which RADIUS Access-Accept packets were received by the proxy.	Accepts/Sec	

	Access-Challenges: Indicates the rate at which RADIUS Access-Challenge packets were received by the NPS RADIUS proxy from the RADIUS server.	Challenges/Sec	A low value is desired for this measure. A high value indicates that the RADIUS server challenged many access requests, forcing the access server to send an updated Access-Request to it, through the NPS RADIUS proxy. In such cases, access clients are bound to experience delays in accessing the server.
	Access-Rejects: Indicates the rate at which RADIUS Access-Reject packets were sent by the RADIUS server to the NPS RADIUS Proxy.	Rejects/Sec	Ideally, the value of this measure should be 0 or very low. A high value indicates too many or too frequent request rejections, which in turn may cause access clients to be denied access to the server.
	Access-Requests: Indicates the rate at which Access-Request packets were received by the NPS RADIUS Proxy from the access servers.	Reqs/Sec	This is a good indicator of the load on NPS.
	Bad authenticators: Indicates the rate at which NPS RADIUS Proxy received access requests containing an invalid Message Authenticator attribute.	Reqs/Sec	Ideally, the value of this measure should be 0.
	Packets dropped: Indicates the rate at which request packets received by the NPS RADIUS proxy were silently discarded for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".	Packets/Sec	Ideally, the value of this measure should be 0.
	FullAccess-Decisions: Indicates the rate at which Full-access decisions were received the NPS RADIUS Proxy.	Decisions/Sec	The RADIUS server grants an access client full access if the client meets the defined health policies.
	Invalid addresses: Indicates the rate at which the NPS RADIUS Proxy received packets from unknown addresses.	Packets/Sec	Ideally, this value should be 0.

	Malformed packets: Indicates the rate at which the NPS RADIUS Proxy received malformed packets.	Packets/Sec	Ideally, the value of this measure should be 0.
	Packets received: Indicates the rate at which packets were received by the NPS RADIUS Proxy.	Packets/Sec	
	Probation-Decisions: Indicates the rate at which probation-decisions were received from the NPS RADIUS Proxy.	Decisions/Sec	If the RADIUS server grants an access client full access but for a limited period only, the client is said to be on probation. This can happen if the RADIUS server finds that the client did not fulfill certain health policy requirements.
	Quarantine-Decisions: Indicates the rate at which quarantine decisions were sent from the NPS RADIUS Proxy.	Decisions/Sec	When a remote access client dials in or connects via VPN to an access server, by default only the user's credentials (account name and password) are checked to determine whether access is granted. This means a computer that does not meet the network's policy requirements could still connect to the server and the network from a remote location. When quarantine control is deployed, after the user's credentials are authenticated the connection is "quarantined." In quarantine mode, the computer has an IP address and has limited access to some network resources (called quarantine resources) such as a DNS server and perhaps a file server or web server from which it can download files necessary to comply with the policies or where the user can get more information, but cannot access the rest of the network.
	Request timeouts: Indicates the rate at which requests to the NPS RADIUS proxy timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.

	Retransmissions: Indicates the rate at which requests were retransmitted to the NPS RADIUS Proxy.	Reqs/Sec	Retransmits can increase the number of requests to the proxy, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on NPS RADIUS Proxy. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.
	Unknown type: Indicates the average number of unknown type (non-RADIUS) packets received by this servethe NPS RADIUS proxy per second.	Packets/Sec	
	Pending requests: Indicates the rate of requests destined for the proxy that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on the NPS RADIUS Proxy or a high timeout setting. In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests. In the case of the former, you may want to consider adding more RADIUS servers, so that the NPS RADIUS Proxy is able to dynamically balance the load of connection requests across multiple RADIUS servers and thus speed up processing.

12.1.10 NPS Accounting Proxy Test

Once the access server completes the connection process with the access client, it sends an Accounting-Request message to the NPS RADIUS proxy. The NPS RADIUS proxy logs the accounting data and forwards the message to the RADIUS server. The RADIUS server then sends an Accounting-Response to the NPS RADIUS proxy, where it is forwarded to the access server.

If the access servers experience accounting delays, it could either be owing to a slowdown in the NPS RADIUS proxy that routes the responses or the poor processing ability of the RADIUS server that sends the responses to the proxy. Malformed packets, packets with invalid attributes/addresses, and non-RADIUS packets can aso contribute to the time lag at the proxy server end. Another common reason for the slowdown is a request overload on the NPS RADIUS proxy.

Using the **NPS Accounting Proxy** test, administrators can measure how effectively the proxy is handling accounting requests, detect slowdowns, and pinpoint the probable reasons for the same.

Purpose	Measures how well the proxy handles accounting requests, detects slowdowns, and pinpoints the probable reasons for the same.		
Target of the test	An NPS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The port at which the NPS server listens. The default is NULL.		
Outputs of the test	One set of results for the NPS RADIUS Proxy		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Accounting-Requests: Indicates the rate at which the NPS RADIUS proxy receives accounting requests.	Reqs/Sec	This is a good indicator of the workload on the server.
	Accounting - Responses: Indicates the rate at which the NPS RADIUS proxy is responding to requests.	Reqs/Sec	If the value of this measure is much lower than the value of the *Accounting-Requests* measure, it could indicate that the server is not responding to requests quickly. You may want to investigate the reasons for the same. This could either be caused by a processing bottleneck on the proxy server, a poor network connection between the proxy and the RADIUS server, or a slowdown on the RADIUS server.
	Bad authenticators: Indicates the rate at which the proxy received requests containing an invalid Message Authenticator attribute.	Reqs/Sec	Ideally, the value of this measure should be 0.
	Packets dropped: Indicates the rate at which the proxy silently discarded the request packets it received for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".	Packets/Sec	Ideally, the value of this measure should be 0.

	Invalid addresses: Indicates the rate at which the proxy received packets with invalid addresses.	Reqs/Sec	Ideally, the value of this measure should be 0.
	Malformed packets: Indicates the rate at which the proxy received malformed packets.	Packets/Sec	Ideally, the value of this measure should be 0.
	Packets received: Indicates the rate at which requests packets were received by the proxy.	Packets/Sec	
	Request timeouts: Indicates the rate at which requests to the proxy timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.
	Retransmissions: Indicates the rate at which requests were retransmitted to the proxy.	Reqs/Sec	Retransmits can increase the number of requests to the proxy server, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on the server. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.
	Unknown type: Indicates the average number of unknown type (non-RADIUS) packets received by the proxy per second.	Packets/Sec	

	Pending requests: Indicates the rate of requests destined for the proxy that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on the proxy, a high timeout setting on the proxy (which could be causing many requests to be retransmitted to the server), or the poor processing power of the RADIUS server. A flaky network connection between the proxy and the RADIUS server can also contribute to the processing delay and add to the count of pending requests.

Chapter 13

Monitoring the Microsoft RAS Server

Microsoft Remote Access Service (RAS) is a feature in the Windows Server family, including Windows Server 2003, Windows 2000 Server, and , NT4 Server. A Limited version of RAS is also included in Windows XP Professional. RAS allows remote dial-up clients to connect to a Local Area Network using analog phone lines or ISDN lines. A typical use would be by an ISP (Internet Service Provider) to allow users to dial in to their LAN, or by a corporate network administrator to allow their users to connect to the corporate LAN from remote sites. The remote clients connect to RAS using the TCP/IP protocol encapsulated in the Point-to-Point (PPP) protocol, which allows the remote client to access the LAN as if they were plugged directly into it.

Needless to say, even a brief non-availability of RAS can cause critical services to go out of the reach of remote clients. Continuous monitoring of the RAS server can alone ensure a higher uptime of the RAS service. Using the *Microsoft RAS* monitoring model (see Figure 13.1) presented by the eG Enterprise suite, administrators can closely observe RAS operations 24x7, be forewarned of probable issues, and quickly attend to the issues before any permanent damage occurs.

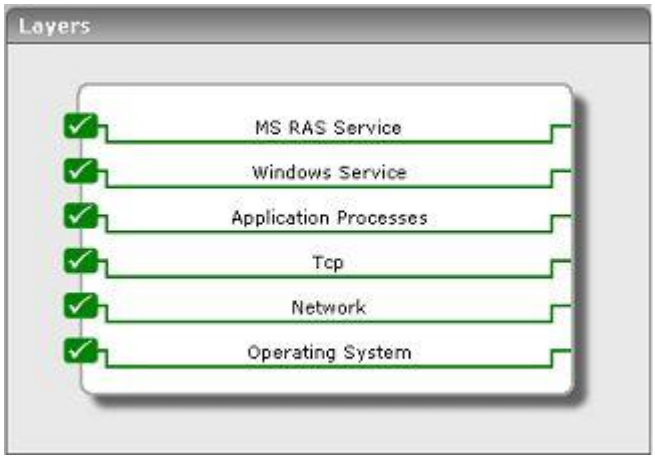

Figure 13.1: Layer model of the MS RAS server

The sections to come will deal with the tests mapped to the **MS RAS Service** layer only, as the remaining layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

13.1 The MS RAS Service Layer

Using the tests depicted by Figure 13.2, the **MS RAS Service** layer enables administrators to assess the effectiveness of the dial-up communication service provided by the RAS device.

Figure 13.2: The tests associated with the MSRAS_SERVICE layer

13.1.1 Microsoft RAS Port Test

The MsRasPort test reports the performance statistics pertaining to every port of the Remote Access Service (RAS) device on the computer.

Purpose	Reports the performance statistics pertaining to every port of the Remote Access Service (RAS) device on the computer		
Target of the test	A Microsoft RAS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The TCP port at which the RAS server listens. The default is NULL.		
Outputs of the test	One set of results for every RAS port		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes transmitted: Indicates the rate at which bytes were transmitted.	Bytes/Sec	
	Bytes received: Indicates the rate at which bytes were received.	Bytes/Sec	When viewed along with the Bytes_transmitted measure, this measure serves as a good indicator of the traffic on the network.

	Frames transmitted: Indicates the number of frames transmitted per second.	Frames/Sec	
	Frames received: Indicates the number of frames received per second.	Frames/Sec	
	Total errors: Indicates the number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors per second.	Number	
	Compression ratio for bytes sent: Indicates the compression ratio for the bytes being transmitted.	Percent	
	Compression ratio for bytes received: Indicates the compression ratio for the bytes being received.	Percent	
	Total connections: Indicates the number of remote access connections.	Number	
	CRC errors: Indicates the current number of CRC errors for this port.	Number	CRC errors occur when the frame received contains erroneous data.
	Timeout errors: Indicates the current number of timeout errors for this port.	Number	Timeout errors occur when an expected packet is not received in time.
	Serial overrun errors: Indicates the current number of serial overrun errors for this port t.	Reqs/Sec	Serial Overrun errors occur when the hardware cannot handle the rate at which data is received.
	Alignment errors: Indicates the current number of alignment errors for this port.	Number	Alignment errors occur when a received byte is different from the expected byte.

	Buffer overrun errors: Indicates the current number of buffer overrun errors for this port.	Number	Buffer Overrun errors occur when the software cannot handle the rate at which data is received.

13.1.2 Microsoft RAS Test

The MsRas test reports the performance statistics that are aggregated across all the ports of the Remote Access Service (RAS) device on the computer.

Purpose	Reports the performance statistics that are aggregated across all the ports of the Remote Access Service (RAS) device on the computer		
Target of the test	A Microsoft RAS server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The TCP port at which the RAS server listens. The default is NULL.		
Outputs of the test	One set of results for the RAS server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bytes transmitted: Indicates the rate at which bytes were transmitted.	Bytes/Sec	
	Bytes received: Indicates the rate at which bytes were received.	Bytes/Sec	When viewed along with the Bytes_transmitted measure, this measure serves as a good indicator of the traffic on the network.
	Frames transmitted: Indicates the number of frames transmitted per second.	Frames/Sec	
	Frames received: Indicates the number of frames received per second.	Frames/Sec	
	Total errors: Indicates the number of CRC, Timeout, Serial Overrun, Alignment, and Buffer Overrun errors per second.	Number	

	Compression ratio for bytes sent: Indicates the compression ratio for the bytes being transmitted.	Percent	
	Compression ratio for bytes received: Indicates the compression ratio for the bytes being received.	Percent	
	Total connections: Indicates the number of remote access connections.	Number	
	CRC errors: Indicates the current number of CRC errors for this port.	Number	CRC errors occur when the frame received contains erroneous data.
	Timeout errors: Indicates the current number of timeout errors for this port.	Number	Timeout errors occur when an expected packet is not received in time.
	Serial overrun errors: Indicates the current number of serial overrun errors for this port t.	Reqs/Sec	Serial Overrun errors occur when the hardware cannot handle the rate at which data is received.
	Alignment errors: Indicates the current number of alignment errors for this port.	Number	Alignment errors occur when a received byte is different from the expected byte.
	Buffer overrun errors: Indicates the current number of buffer overrun errors for this port.	Number	Buffer Overrun errors occur when the software cannot handle the rate at which data is received.

13.1.3 Windows Telephony Test

The MsTelephony test measures the performance of the telephone-communication activity on a computer running Windows 2000 or a higher operating system.

Purpose	Measures the performance of the telephone-communication activity on a computer running Windows 2000 or a higher operating system
Target of the test	A Microsoft RAS server
Agent deploying the	An internal agent

test			
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured 3. **PORT** - The TCP port at which the RAS server listens. The default is NULL.		
Outputs of the test	One set of results for the RAS server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Telephone lines: Indicates the number of telephone lines currently serviced by this computer.	Number	
	Telephone devices: Indicates the number of telephone devices (telephones or speaker phones) currently serviced by this computer.	Number	
	Active telephone lines: Indicates the number of telephone or integrated services digital network (ISDN) lines serviced by this computer that are currently in use by applications.	Number	
	Active telephone devices: Indicates the number of telephone devices (telephones or speaker phones) that are currently in use by applications.	Number	
	Outgoing calls: Indicates the rate at which outgoing calls are made by this computer.	Calls/Sec	
	Incoming calls: Indicates the rate at which incoming calls are answered by this computer.	Calls/Sec	

	Client applications using telephony services: Indicates the number of applications that are currently using telephony services.	Number	
	Current outgoing calls: Indicates the number of outgoing calls that are currently being serviced by this computer.	Number	
	Current incoming calls: Indicates the number of incoming calls that are currently being serviced by this computer.	Number	

Chapter 14

Monitoring Microsoft System Management Servers (SMS)

Microsoft Systems Management Server provides a comprehensive solution for change and configuration management for the Microsoft platform, enabling organizations to provide relevant software and updates to users quickly and cost-effectively.

In order to make sure that critical software updates are quickly and readily available to the users, the Microsoft SMS has to be monitored periodically for availability and optimal performance.

eG Enterprise provides administrators with an exclusive *Microsoft SMS* monitoring model that carefully examines the critical services and core functions of the Microsoft SMS, and proactively alerts them to performance aberrations that can adversely impact the user interaction with the server.

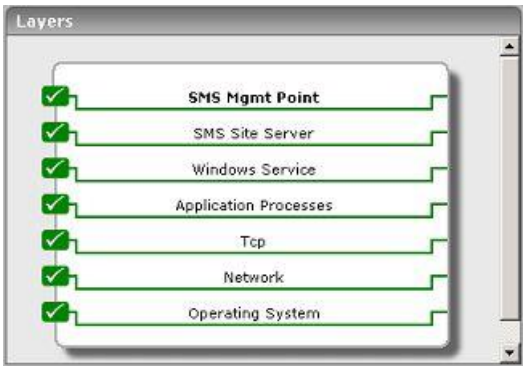

Figure 14.1: The layer model of Microsoft SMS

The sections to come discuss the top 2 layers of the hierarchical layer model depicted by Figure 14.1. The other layers have already been dealt with in the *Monitoring Unix and Windows Servers* document.

14.1 The SMS Site Server Layer

The tests mapped to this layer monitor the health of core components of the Microsoft SMS, such as:

- The Discovery Data Manager
- The Inventory Loader
- The SMS Memory Queue
- The SMS_STATUS_MANAGER
- The Software Inventory Processor

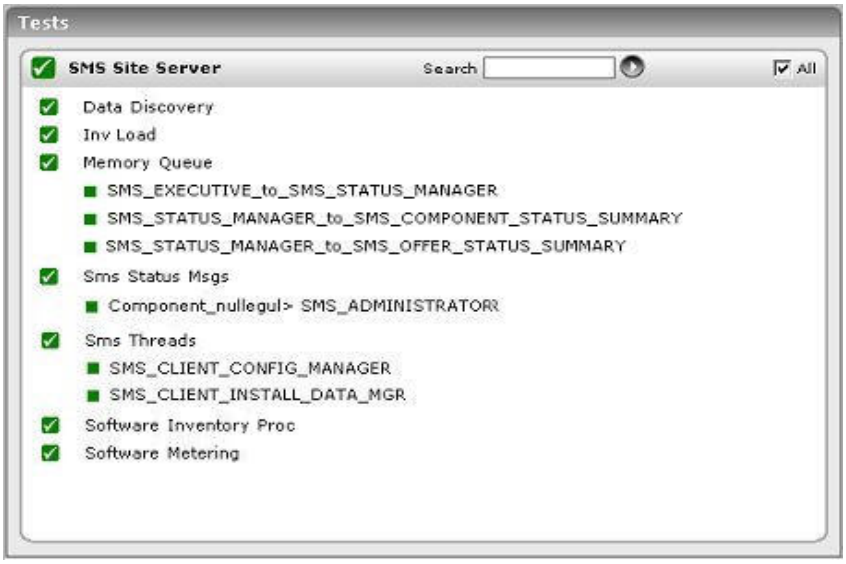

Figure 14.2: The tests associated with the SMS Site Server layer

Besides, the layer also reveals the state of threads executing on the Microsoft SMS, and the effectiveness of its Software Metering feature.

14.1.1 Data Discovery Test

This test monitors the Discovery data manager of SMS. This Data Manager discovers data about the SMS Clients (computers connected to the network and the SMS server).

Purpose	Monitors the Discovery data manager of SMS
Target of the test	Microsoft SMS
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured

Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bad data records processed: Indicates the number of bad (ill-formed or invalid) data records processed by the Discovery Data Manager.	Number	
	Data records waiting in the input queue: Indicates the number of SMS Discovery data records waiting in the Discovery Manager's input queue the last time the input queue was scanned minus the number of data records processed till then.	Number	When many data records are written to the input queue, this counter is too low until the Discovery Manager scans the input queue again. This means many data records have been processed in that period.
	Total data records processed: Indicates the number of Discovery Data records processed in the last test frequency.	Number	

14.1.2 Inv Load Test

This test reports metrics pertaining to the Inventory Data Loader of SMS, which loads the client configuration details pertaining to the system hardware.

Purpose	Reports metrics pertaining to the Inventory Data Loader of SMS		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Bad Management Information Files (MIFs) processed:** Indicates the number of bad (ill-formed or otherwise invalid) SMS hardware inventory records (in MIF - Management Information Format files) processed by Inventory Data Loader since it was last started.	Number	
	MIFs enqueued: Indicates the number of MIF files (containing SMS hardware inventory records) that were waiting in the Inventory Data Loader's input queue the last time Inventory Data Loader scanned the queue, minus the MIF files processed since then	Number	When many MIF files are being written to the input queue, this measure will be too low until Inventory Data Loader scans the input queue again.
	MIFs processed: Indicates the number of SMS hardware inventory records (in MIF files) processed by the Inventory Data Loader since it was last started.	Number	

14.1.3 Memory Queue Test

The MemoryQueue test monitors the health of the SMS memory queue. It is to this SMS Memory Queue thread that a component adds an object when waiting and another component picks the object for its function and removes it from the queue.

Purpose	Monitors the health of the SMS memory queue
Target of the test	Microsoft SMS
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured
Outputs of the test	One set of results for every memory queue thread on the monitored SMS server

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Objects dequeued: Indicates the number of objects that the destination component has removed from the queue.	Number	
	Objects enqueued: Indicates the number of objects that the source component has added to the queue	Number	

14.1.4 SMS Status Messages Test

The SmsStatusMsgs test tracks the status messages handled by the SMS_STATUS_MANAGER.

Purpose	Tracks the status messages handled by the SMS_STATUS_MANAGER		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every SMS component monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	High priority: Indicates the number of SMS status messages replicated to the parent site at high priority by the Status Manager.	Number	
	Low priority: Indicates the number of SMS status messages replicated to the parent site at low priority by the Status Manager.	Number	

	Normal priority: Indicates the number of SMS status messages replicated to the parent site at normal priority by the Status Manager.	Number	
	Report app evt log: Indicates the number of SMS status messages reported by the Status Manager to the Windows NT Application Event Log on the site server.	Number	
	Database writes: Indicates the number of SMS status messages queued by the Status Manager to be written to the SMS site database.	Number	This number equals the number of status messages actually written to the database, unless Status Manager cannot write to the database (because it is full, for example), in which case the number of queued messages (shown by this counter) will increase even though no messages are being written to the database. (Queued messages are stored as .SQL files in \SMS\Inboxes\Statmgr.box\Retry.) When the database becomes writable again, the queued messages will rapidly be written to it, and this counter will again reflect the actual number of messages written to the database.

14.1.5 SMS Threads Test

This test reports the state of the SMS threads.

Purpose	Reports the state of the SMS threads		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Running threads:** Indicates the number of running threads in the SMS Executive (SMSEXEC.EXE) service.	Number	When this measure is associated with a single thread instead of the entire service, its value is zero (the thread is not running) or one (the thread is running).

14.1.6 Software Inventory Proc Test

This test reports metrics pertaining to the Software Inventory Processor of SMS. The Software Inventory Processor processes the files produced by the Software Inventory Manager.

Purpose	Reports metrics pertaining to the Software Inventory Processor of SMS		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bad software inventory records processed: Indicates the number of bad (ill-formed or otherwise invalid) SMS software inventory records (SINVs) processed by the Software Inventory Processor since it was last started.	Number	
	Software inventory records waiting in input queue: Indicates the number of SMS software inventory records (SINVs) waiting in the Software Inventory Processor's input queue the last time Software Inventory Processor scanned the queue, minus the SINVs that have been processed since the queue was last scanned.	Number	When many SINVs are being written to the input queue, this counter is too low until Software Inventory Processor scans the input queue again.

	Total software inventory records processed: Indicates the number of SMS software inventory records (SINVs) processed by Software Inventory Processor since it was last started.	Number	

14.1.7 Software Metering Test

This test monitors the Software Metering feature, which allows one to monitor program usage on client computers. By using software metering, one can collect data about software usage in one's organization. Software metering data can be conveniently summarized to produce useful reports that can help one monitor licensing compliance and plan software purchases in one's organization. Software metering collects detailed information about the programs that you chose to monitor. This includes information about program usage, program users, the program start time, and the length of time it is used.

Purpose	Monitors the Software Metering feature, which allows one to monitor program usage on client computers		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Bad software metering files processed: Indicates the number of bad (ill-formed or otherwise invalid) SMS software metering usage files processed by Software Metering Processor since it was last started.	Number	

	Usage files waiting in the input queue: Indicates the number of SMS software metering usage files waiting in the Software Metering Processor's input queue, minus the number of files that have been processed since the queue was last scanned.	Number	
	Usage processing threads: Indicates the number of threads the Software Metering Processor is currently using to process incoming SMS software metering usage files.	Number	
	Total usage records processed: Indicates the number of software metering records processed by the SWM Processor.	Number	

14.2 The SMS Mgmt Point Layer

This layer tracks the health of the SMS Management Point components.

Figure 14.3: The tests associated with the SMS Mgmt Point layer

14.2.1 Management Point Data Loader Test

The MgmtPointDataLoader test reports metrics pertaining to the Management Point Data Loader object, which monitors the SMS interactions with the database.

Purpose	Reports metrics pertaining to the Management Point Data Loader object, which monitors the SMS interactions with the database		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Connections created: Indicates the number of database connections created by the management point.	Number	

	Connections create rate: Indicates the number of database connections created by the management point per second.	Conns/Sec	

14.2.2 MgmtPointHwInv Test

This test reports metrics pertaining to the Hardware Inventory Manager. The SMS hardware inventory feature automatically collects detailed information about the hardware characteristics of clients in an SMS hierarchy. By using this feature, you can collect a wide variety of information about client computers such as memory, operating system, peripherals, services, and processes that are running on the client computer.

Purpose	Reports metrics pertaining to the Hardware Inventory Manager		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Delta reports: Indicates the number of hardware inventory reports marked as Delta.	Number	
	Reports data generated: Indicates the size of generated reports.	MB	
	Reports processed: Indicates the number of reports processed, successfully or unsuccessfully.	Number	
	Reports process rate: Indicates the number of reports processed per second.	Reports/Sec	

14.2.3 Management Point Policy Manager Test

This test monitors the responses of the SMS Policy Manager to the policy requests of clients.

Purpose	Monitors the responses of the SMS Policy Manager to the policy requests of clients		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Request arrival rate: Indicates the rate at which Policy Assignment requests are arriving at the Policy Manager.	Requests/Sec	

14.2.4 Management Point Policy Test

This test reports the results of the client requests to the SMS Policy Manager. There are certain SMS policies which download in the client system. This is controlled by the SMS Policy Manager.

Purpose	Reports the results of the client requests to the SMS Policy Manager. There are certain SMS policies which download in the client system		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Cache hit rate:** Indicates the rate of requests to the Get Policy component that resulted in the policy being served from a cache.	Hits/Sec	
	Requests process rate: Indicates the rate of requests to the Get Policy component.	Requests/Sec	

14.2.5 Management Point Status Manager Test

This test reports metrics pertaining to the Status Manager of SMS. SMS generates status messages to report the activity of components on site systems and clients. A status message is a text string, generated by a component, describing a specific activity performed by the component. In addition, each status message contains important information such as which component generated the message, the exact time that the message was generated, and the severity of the message. Status messages are sent from clients and site systems to the site server and are stored in the SMS site database. You can then view status messages in the SMS Administrator console. Viewing status messages in the SMS Administrator console helps you monitor the activity of the various components, determine the health of SMS, and identify issues that might require your attention.

Purpose	Reports metrics pertaining to the Status Manager of SMS. SMS generates status messages to report the activity of components on site systems and clients		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Events processed: Indicates the number of events (i.e. status messages) processed, successfully or unsuccessfully.	Number	
	Events process rate: Indicates the number of events (i.e. status messages) processed per second.	Number	

14.2.6 Management Point Software Inventory Test

This test reports metrics pertaining to the reports generated by the Software Inventory manager of SMS. With the SMS Software Inventory Manager one can collect information about the applications listed in Add or Remove Programs in Control Panel. By using software inventory, one can collect a significantly larger amount of information about client's software.

Purpose	Reports metrics pertaining to the reports generated by the Software Inventory manager of SMS		
Target of the test	Microsoft SMS		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** – How often should the test be executed 2. **HOST** – The host for which the test is to be configured		
Outputs of the test	One set of results for every Microsoft SMS server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Delta reports: Indicates the number of Software Inventory reports marked as Delta.	Number	
	Reports data generated: Indicates the size of generated reports.	MB	
	Reports processed: Indicates the number of reports processed successfully or unsuccessfully.	Number	
	Reports process rate: Indicates the number of reports processed per second.	Reports/Sec	

Chapter 15

Externally Monitoring the Active Directory Server

The *Active Directory* server model discussed in Chapter 0 of this document, performs in-depth internal monitoring of the health of an Active Directory (AD) server. However, sometimes, administrators might be denied access to the AD servers to be monitored, and hence might be unable to install agents on them. Such administrators might still want to monitor the availability and responsiveness of the Active Directory. To cater to the needs of these administrators, eG Enterprise offers the *External AD* model (see Figure 15.1).

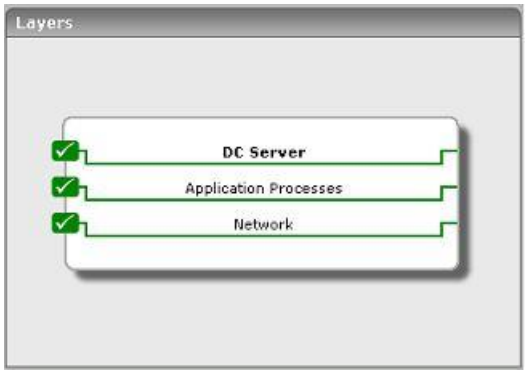

Figure 15.1: Layer model of the External AD server

Besides detecting the non-availability or slow responsiveness of an AD server, this model also runs port availability and network health checks, so as to ensure that all external performance parameters are functioning properly.

15.1 The Network Layer

The **Network** test (see Figure 15.2) associated with this layer performs network-level pings to assess the overall health of the network connection to the AD server.

Figure 15.2: The test associated with the Network layer

For details on the **Network** test, refer to the *Monitoring Unix and Windows Servers* document.

15.2 The Application Processes Layer

Using the TcpPortStatus test depicted by Figure 15.3, administrators can externally monitor the availability and responsiveness of the AD server port.

Figure 15.3: The tests associated with the Application Processes layer

Please refer to the *Monitoring Unix and Windows Servers* document for a discussion on the TcpPortStatus test.

15.3 The DC Server Layer

By emulating a user request to the AD server, the **ADServer** test associated with this layer (see Figure 15.4) helps determine the availability and responsiveness of the AD server.

Figure 15.4: The tests associated with the DC Server layer

For more details, refer to Chapter 0 of this document.

Chapter 16

Monitoring the AD Cluster Service

An active directory (AD) cluster service is a collection of physical AD servers that can act as a single logical server. Requests to a cluster are routed through a virtual cluster server that is assigned a cluster IP address and TCP port. Requests to this server can be handled by any of the individual nodes in the cluster at any given point in time, depending on which node is active at that time.

Since clusters are deployed in environments where 24*7 availability and responsiveness are critical, it is imperative that the performance of the clusters is monitored all the time.

To monitor an Active Directory cluster, an eG external agent is deployed, which emulates a user login to the cluster to determine the availability of the cluster and the speed with which the cluster responds to the emulated request. The emulated requests are directed at the virtual cluster server. Therefore, you need to manage the virtual cluster server as an *AD Cluster* using the eG administrative interface.

Note:

For more details on how eG Enterprise monitors clusters, refer to Chapter 0 of the *eG User Manual*.

The layer model of the *AD Cluster* has been depicted by Figure 16.1 below.

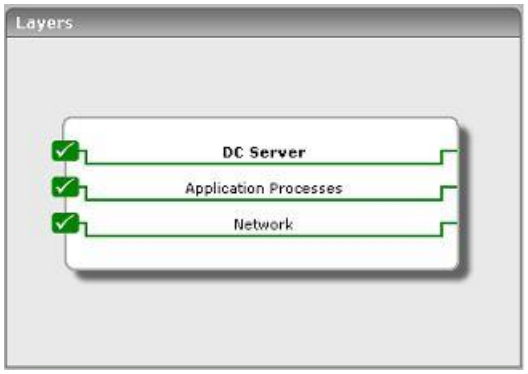

Figure 16.1: Layer model of the AD cluster service

The following section deals only with the first layer of Figure 16.1, as the other layers and the external metrics they report have been dealt with in the previous chapter.

16.1 The DC Server Layer

The **ADServer** test associated with this layer emulates a user login to the cluster to determine its availability and responsiveness. The test sends the emulated request to the virtual cluster server (i.e., the *AD Cluster*), which will promptly forward the request to that node in the cluster that currently owns the cluster server. If at least one node in the cluster is currently active, then the login request will succeed, indicating the good health of the cluster. On the other hand, if none of the nodes in the cluster are active, then the emulated request will fail, indicating the non-availability of the cluster.

Figure 16.2: The tests associated with the DC_SERVER layer

In Chapter 0 of this document, the **ADServer** test has been elaborately discussed. Refer to it for further details.

Chapter 17

Monitoring Windows Clusters

Microsoft Cluster Server (MSCS) is software designed to allow servers to work together as computer clust er, to provide failover and increased availability of applications, or parallel calculating power in case of high-performance computing (HPC) clusters (as in supercomputing).

To monitor Windows clusters, eG Enterprise provides a specialized *Microsoft Windows Cluster Node* monitoring model. Cluster monitoring enables you to do the following:

With the help of the tests mapped to this layer, you can determine the following:

- Know the clusters that are currently managed by the Windows Failover Cluster Manager;
- Know which nodes are part of a cluster;
- Determine the current state of each node;
- Rapidly detect the failure of the cluster service on the monitored node;
- Identify the services/applications that have been clustered, promptly detect service/application failures, and pinpoint the probable reasons for the same;
- Identify cluster networks that are currently down;
- Pinpoint cluster resources that are offline;
- Track the current capacity and usage of cluster disks and cluster shared volumes and proactively detect potential

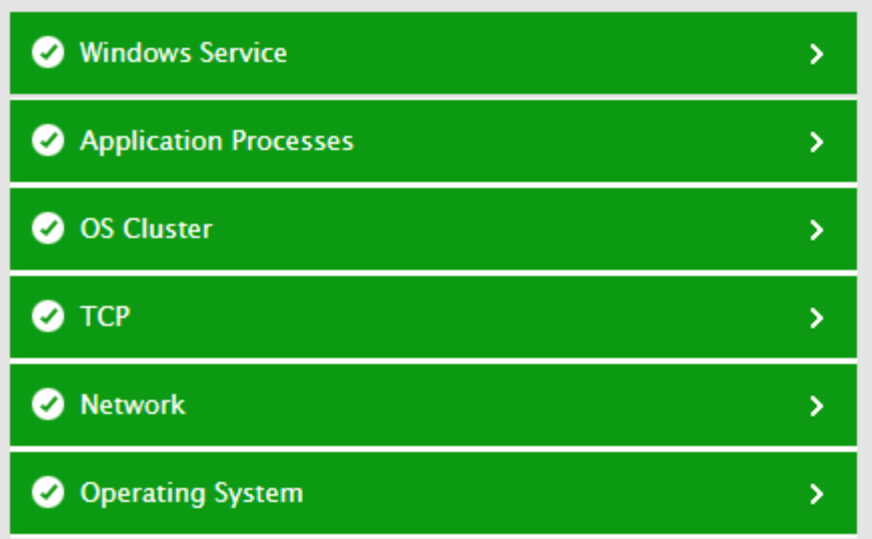

Figure 17.1: The layer model of the Microsoft Windows Cluster Node

The tests mapped to the **Windows Service** and **OS Cluster** layers, upon execution, provide information pertaining to the cluster status. This chapter will discuss the cluster-related tests mapped to the **Windows Service** layer alone. For details of tests mapped to the **OS Cluster** layer and all other layers of Figure 17.1, refer to the *Monitoring Unix and Windows Servers* document.

Note:

The tests mapped to the **Windows Service** and **OS Cluster** layers run only in the **agent-based** mode. This is why, you need to install an eG agent on at least one node in the cluster to enable these tests to report cluster-level metrics. For best results however, it is recommended that you install an eG agent on each node in the cluster; this way, even if one node goes down due to any reason, cluster health can continue to be monitored using the agents on the other nodes.

17.1 The Windows Service Layer

Using the tests mapped to this layer, you can determine the status of the Windows cluster service.

Figure 17.2: The tests mapped to the Windows Service layer

17.1.1.1 Cluster Groups Test

A resource group is a collection of resources, managed by the Cluster service as a single, logical unit. This logical unit is often referred to as a failover unit, because the entire group moves as a single unit between nodes. Resources and cluster elements are grouped logically according to the resources added to a resource group. When a Cluster service operation is performed on a resource group, the operation affects all individual resources contained in the group. Typically, a resource group is created that contains the individual resources required by the clustered program.

Cluster resources may include physical hardware devices, such as disk drives and network cards, and logical items such as IP addresses, network names, and application components.

The ClusterGroups test indicates the current status of the resource groups.

Purpose	Reports the current status of the resource groups		
Target of the test			
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows application. 4. **EXCLUDE CLUSTER OFFLINE GROUPS** - Provide a comma-separated list of cluster groups to be excluded from the monitoring scope of this test.		
Outputs of the test	One set of results for the server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Groups online: Indicates the number of resource groups in the cluster that are currently online.	Number	
	Groups offline: Indicates the number of groups in the cluster that are currently offline.	Number	This count includes only those groups in which all resources are offline.
	Groups partially online: Indicates the number of groups in the cluster that are partially online.	Number	This count includes only those groups in which some resources are online and some offline.

17.1.1.2 Windows Cluster Nodes Test

Every server in a cluster is referred to as a *Node*. Using the ClusterNodes test, you can identify how many nodes in the cluster are currently not available.

Purpose	Reports the current status of the nodes in the cluster		
Target of the test			
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows application.		
Outputs of the test	One set of results for the server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Is cluster node up?: Indicates the number of nodes that are currently up and running in the cluster.	Number	An online node or a node whose status is 'Up', is an active member of the cluster. It adheres to cluster database updates, contributes input into the quorum algorithm, maintains cluster network and storage heartbeats, and can own and run resource groups.
	Is cluster node down?: Indicates the number of nodes in the cluster that are currently offline.	Number	A node whose status is 'down' or 'offline' is considered to be an inactive member of the cluster. The node and its Cluster service might or might not be running.
	Is cluster node paused?: Indicates the number of nodes in the cluster that are paused.	Number	This refers to nodes that are active members of the cluster. The node adheres to cluster database updates, contributes input into the quorum algorithm, and maintains network and storage heartbeats, but it cannot accept resource groups. It can support only those resource groups that it currently owns. The paused state enables maintenance to be performed. Online and paused states are treated as equivalent states by the majority of the server cluster components.

17.1.1.3 Cluster Network Interfaces Test

The network adapter or adapters (also known as network interface cards or NICs) on each node in a cluster enables the nodes to communicate with each other and with clients. This test reveals whether/not there are any network interfaces in the cluster that are not functioning properly.

Purpose	Reveals whether/not there are any network interfaces in the cluster that are not functioning properly
Target of the	

<table>
<tr><td>test</td><td colspan="3"></td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the Windows application.
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the server being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Network interfaces that are up:
Indicates the number of network interfaces that are currently up and running in the cluster.</td><td>Number</td><td></td></tr>
<tr><td>Network interfaces that are down:
Indicates the number of network interfaces in the cluster that are currently not running.</td><td>Number</td><td>Each node in a failover cluster requires network connectivity with the other nodes. Problems with a network adapter or other network device (either physical problems or configuration problems) can interfere with connectivity.
Therefore, ideally, the value of this measure should be 0.</td></tr>
</table>

17.1.1.4 Windows Cluster Resources Test

A resource is a physical or logical entity that is capable of being managed by a cluster, brought online, taken offline, and moved between nodes. A resource can be owned only by a single node at any point in time.

This test reports the number of resources in the cluster and their current states.

Purpose	Reports the number of resources in the cluster and their current states

<table>
<tr><td>Target of the test</td><td colspan="3"></td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the Windows application.
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the server being monitored</td></tr>
<tr><td rowspan="3">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Online resources:
Indicates the number of resources that are currently online.</td><td>Number</td><td></td></tr>
<tr><td>Offline resources:
Indicates the number of resources that are currently offline.</td><td>Number</td><td></td></tr>
</table>

	Failed resources: Indicates the number of resources that have failed.	Number	Typically, a resource failure triggers a recovery action, which could be a resource restart or a transfer of the resource to another node. Typically, the Failover Manager and Resource Monitor work together to detect and recover from resource failures. Resource Monitors keep track of resource status by using the resource DLLs to periodically poll resources. Polling involves two steps, a cursory LooksAlive query and a longer, more definitive, IsAlive query. When Resource Monitor detects a resource failure, it notifies Failover Manager and continues to monitor the resource. Failover Manager maintains resources and resource group status. It also performs recovery when a resource fails and invokes Resource Monitors in response to user actions or failures. After a resource failure is detected, Failover Manager performs recovery actions that include restarting a resource and its dependent resources, or moving the entire resource group to another node. The recovery action that is taken is determined by resource and resource group properties, in addition to node availability. During failover, the resource group is treated as the unit of failover. This ensures that resource dependencies are correctly recovered. When a resource recovers from a failure, Resource Monitor notifies Failover Manager. Failover Manager then performs automatic failback of the resource group, based on the configuration of the resource group failback properties.

17.1.1.5 Windows Cluster Networks Test

A network (sometimes called an interconnect) performs one of the following roles in a cluster:

- A *private network* carries internal cluster communication. The Cluster service authenticates all internal communication, but administrators who are particularly concerned about security can restrict internal communication to physically secure networks.
- A *public network* provides client systems with access to cluster application services. IP Address resources are created on networks that provide clients with access to cluster services.
- A *mixed* (public-and-private) *network* carries internal cluster communication and connects client systems to cluster application services.

- A network that is not enabled for use by the cluster (that is, neither public nor private) carries traffic unrelated to cluster operation.

Regardless of the role that a network performs, its availability is critical to the smooth functioning of the cluster, as without the network, communication between cluster nodes and between clients and cluster nodes become impossible.

The ClusterNetworks test indicates whether the cluster networks are up or down.

<table>
<tr><td>Purpose</td><td colspan="3">Indicates whether the cluster networks are up or down</td></tr>
<tr><td>Target of the test</td><td colspan="3"></td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the Windows application.
4. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the server being monitored</td></tr>
<tr><td rowspan="2">Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td>Networks that are up:
Indicates the number of cluster networks that are currently available.</td><td>Number</td><td></td></tr>
</table>

	Networks that are down: Indicates the number of cluster networks that are currently unavailable.	Number	If there is only one cluster network available in a cluster and it goes down, the cluster nodes stop communicating with each other. When two nodes are unable to communicate, they are said to be partitioned. After two nodes become partitioned, the Cluster service automatically shuts down on one node to guarantee the consistency of application data and the cluster configuration. This can lead to the unavailability of all cluster resources. Therefore, it is recommended that you configure multiple networks as private or mixed to protect the cluster from a single network failure. For instance, if each node has at least two networks configured, and both are say, mixed networks, the Cluster service can tolerate network failures. In this scenario, the Cluster service can detect a public network adapter failure and fail over all resources that depend on that adapter (through its IP address) to a node where this network is available. This is accomplished because the private network is still functioning properly. If, on the other hand, an adapter on the private network fails, the Cluster service can use the public network for internal communication. This is accomplished because the public network is mixed, allowing both internal and client traffic.

Chapter 18

Monitoring Microsoft Sharepoint

Microsoft Sharepoint is a collection of products and software elements that include, Internet Explorer based collaboration functions, process management modules, search modules and a document-management platform. Sharepoint can be used to host web sites that access shared workspaces, information stores and documents, as well as host defined applications such as wikis and blogs. All users can manipulate proprietary controls called "web parts" or interact with pieces of content such as lists and document libraries.

If any of the services offered by Microsoft Sharepoint malfunction, it could deny users access to critical organizational data, thereby hampering their productivity and obstructing the achievement of business goals. It is therefore imperative that the Microsoft Sharepoint server is monitored 24x7 for performance deficiencies.

eG Enterprise offers two specialized monitoring models - one for each of the Sharepoint versions - *Microsoft Sharepoint 2007* and *Microsoft Sharepoint 2010*.

This chapter discusses both these models in great detail.

18.1 Monitoring Sharepoint 2007

The *Microsoft Sharepoint 2007* monitoring model continuously monitors the performance of the Sharepoint 2007 server, and proactively alerts administrators to issues.

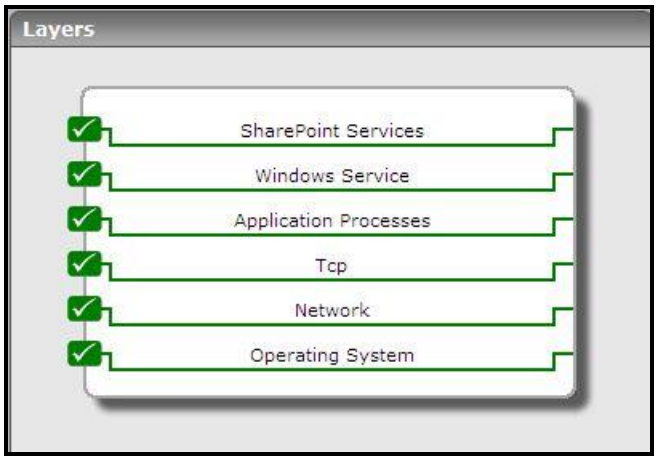

Figure 18.1: The layer model of Sharepoint

Each layer of Figure 18.1 is mapped to a wide variety of tests that report a number of metrics related to the health of the Sharepoint server in question. Using these metrics, the administrators can find quick and accurate answers for the following performance queries:

- Are there too many documents in the first and second queues of the archival plugin? Do these numbers indicate that the crawler is in a starved state?
- Were any error documents returned by the archival plugin?
- How well is the document converter functioning? Are too many conversion requests pending on the converter?
- How is the Excel calculation service performing? Is it responding to requests quickly? How effectively is the service using its cached charts? Are its workbook caches adequately sized?
- Are the Excel Web Access and Excel Web Services components experiencing any slowdowns in request processing?
- Is the content managed by Sharepoint adequately indexed? Are search queries been successfully executed or are too many queries failing?
- Is the gatherer service in a back-off state? If so, why?
- Are your site hit frequency rules very rigid? Are they creating too many delayed documents?
- Are too many threads waiting for documents?
- Are too many threads waiting for a response from the filter process? Is it owing to a network issue or is it because they are bound to a hungry-host?
- Was the gatherer unable to access any documents? If so, how many times?
- Are there too many unprocessed documents on the gatherer?
- Is the Sharepoint Publishing Cache well-tuned? Is the cache hit ratio high?

The sections to come discuss the tests associated with the **Sharepoint Services** layer only, as the remaining layers

have been dealt with elaborately in the *Monitoring Unix and Windows Servers* document.

18.1.1 The Sharepoint Services Layer

Using the tests mapped to this layer, administrators can periodically audit the service levels achieved by the components engaged in the searching and indexing of content managed by Sharepoint. These components include:

g. The Office Server Search Archival Plugin

h. The Office Server Search Schema Plugin

i. The Office Server Search Indexer Catalogs

j. The Office Server Search Gatherer

Similarly, the layer also sheds light on the core components of the Sharepoint Excel Services – namely, the Excel Calculation Service, the Excel Web Access, and the Excel Web Service.

In addition, the layer monitors the health of the object caches and the document converters on Sharepoint 2007.

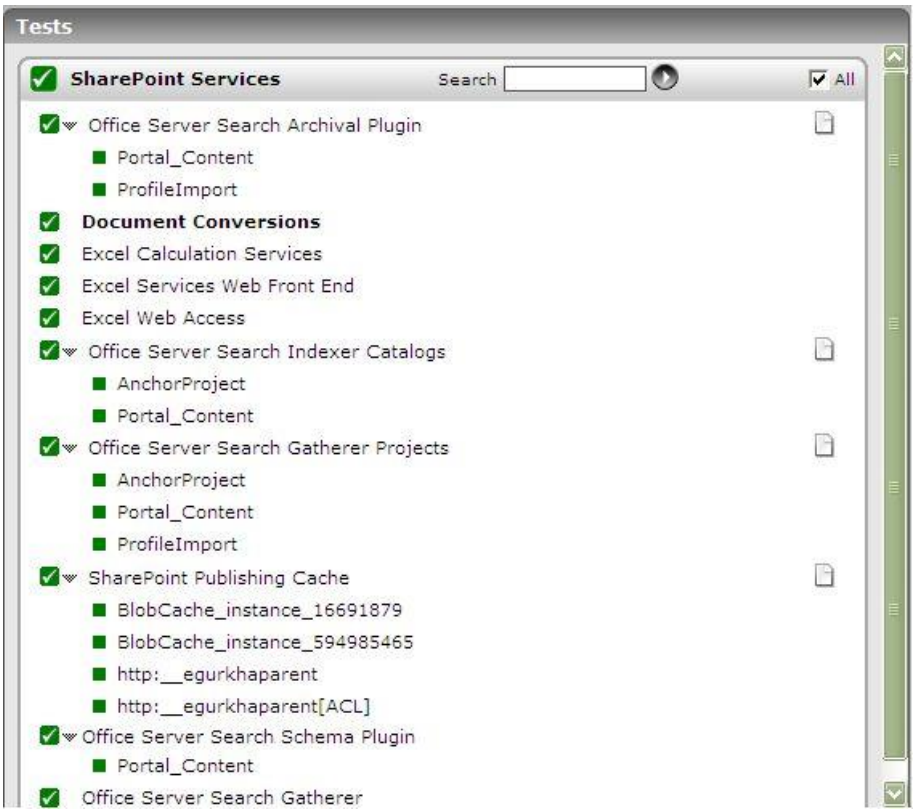

Figure 18.2: The tests mapped to the Sharepoint Services layer

18.1.1.1 Office Server Search Archival Plugin Test

The Search feature of the MOSS 2007 not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites. This is what makes MOSS 2007 that much more valuable to users.

The **Archival** and **Schema** plugins are internal components of the MOSS Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the MOSS search feature is functioning, identify irregularities early, and fine-tune the MOSS server to ensure peak performance of the search engine.

The **Office Server Search Archival Plugin** focuses on the archival plugin component, and helps assess its processing ability.

Purpose	Helps assess the processing ability of the archival plugin component		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 1. **HOST** - The host for which the test is to be configured 2. **PORT –** Refers to the port used by the Windows application.		
Outputs of the test	One set of results each for the *ProfileImport* and *Portal_Content* instances		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active documents in first queue: Indicates the number of documents that are actively using the first queue of the plugin.	Number	One of the more difficult tasks that a Search admin faces is figuring out how to build out the myriad of crawl schedules needed to keep the content on the Sharepoint server freshly indexed. When you are building out these schedules you will want to keep a close eye on the system and slowly add new schedules to minimize starving the crawl of resources while maxing out the utilization of the crawler. Starvation for Enterprise Search is defined as the crawlers inability to allocate another thread to retrieve the next document in the queue of work. This can be caused by resource (I/O) contention on the SQL machine, too many hosts concurrently participating in the crawl, "hungry" hosts that do not quickly relinquish a thread and finally back-ups (since crawls are paused during this time). The values of these measures typically help determine whether the crawler is in a starved state or not. If they are both consistently at 500 for the **Portal_Content** instance or 50 for the **ProfileImport** instance, then you are in a starved state and you are likely to be bottle-necked in SQL for I/O on the Crawl DB drive. Look into tuning SQL for better I/O.
	Active documents in second queue: Indicates the number of documents actively using the second queue of the plugin.	Number	

	Error documents: Indicates the number of documents which currently returned errors from the plugin.	Number	Ideally, this value should be low.
	Bulk insert sessions: Indicates the number of active bulk insert sessions to the database server.	Number	
	Active queue length: Indicates the number of documents currently available in the active queue.	Number	
	Blocked documents: Indicates the number of documents currently waiting for a queue.	Number	

18.1.1.2 Document Conversions Test

A document converter is a custom executable file that takes a document of one file type, and generates a copy of that file in another file type. For example, a document converter might take a Microsoft Office Excel file and use it to generate a Microsoft Office PowerPoint file. Using document converters, you can transform your content into different versions to suit your business needs.

Because document conversions can be resource intensive, Office Sharepoint Server 2007 relies on two services, DocConversionLoadBalancerService and DocConversionLauncherService, to manage the load balancing, prioritizing, and scheduling of the conversions. When a user initiates a document conversion, either through the user interface or object model, Office Sharepoint Server 2007 passes the document conversion request to these two services. It is the DocConversionLauncherService service that actually calls the document converter. When called, the document converter takes the original file and generates a converted copy. Office Sharepoint Server 2007 then takes the converted copy and performs certain post-processing actions on it. These actions include:

- Adding the metadata from the original file to the converted copy.
- Adding metadata that identifies the original file and document converter used to generate the converted copy.
- Notifying the specified people that the conversion has been performed.
- Placing the converted copy into the same document library as the original file.

This test monitors the document conversion process of the Sharepoint server and enables administrators to determine how well the converter is able to process document conversion requests.

Purpose	Monitors the document conversion process of the Sharepoint server and enables administrators to determine how well the coverter is able to process document conversion requests

Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows application.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Incoming E-mail messages processed: Indicates the rate at which e-mail messages have been received and processed by Sharepoint.	E-mails/Sec	
	Pending conversions: Indicates the number of document conversions that are currently pending.	Number	Ideally, the value of this measure should be low. A high value for the measure could indicate a processing bottleneck.

18.1.1.3 Excel Calculation Services Test

Excel Services is built on the Sharepoint products and technologies platform. There are three core components of Excel Services:

- Excel Calculation Service
- Excel Web Access
- Excel Web Service

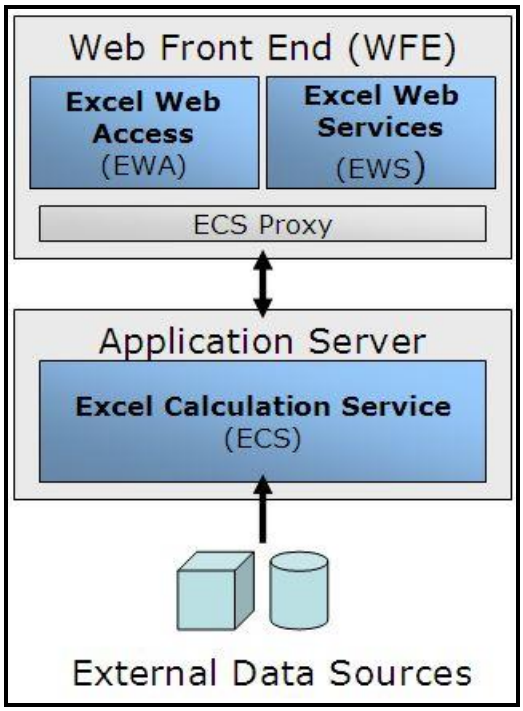

Figure 18.3: Excel services architecture

The role of Excel Calculation Service is to load workbooks, calculate them, call custom code (user-defined functions) and refresh external data. It also maintains the session state for interactivity. Excel Calculation Services maintains a session for the duration of interactions with the same workbook by a user or caller. A session is closed when the caller explicitly closes it or when the session times out on the server. Excel Services caches the opened Excel workbooks, calculation states, and external data query results, for improved performance when multiple users access the same set of workbooks.

In order to determine the quality of the user experience with the Excel Calculation Service, it is essential to know how smooth the user-service interaction is, how quickly the service is able to process the requests, and how effectively the service utilizes its caches. The **Excel Calculation Services** test closely monitors the aforesaid performance parameters, and accurately gauges the health of the service.

Purpose	Accurately gauges the health of the Excel Calculation Service		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the Windows application.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Requests with errors: Indicates the number of requests to the Excel Calculation Service that are returned with errors per second.	Requests/Sec	Ideally, the value of this measure should be low.
	Average number of sessions opened: Indicates the average number of sessions opened per second.	Sessions/Sec	a.
	Cached charts requested: Indicates the number of charts per second that were provided from a cached image.	Charts/Sec	b. A high value is generally desired for this measure, as it indicates the existence of a well-tuned cache. Such a cache goes a long way in reducing processing overheads.
	Active sessions: Indicates the number of currently active sessions on Excel Calculation Services.	Number	c. This value is a good indicator of the current workload on the service.
	Average processing time for a request: Indicates the average processing time for a request on Excel Calculation Services.	Secs	d. A high value for this measure or a gradual increase in this value could be indicative of a processing bottleneck on the service.
	Average session time: Indicates the average session time.	Secs	e.
	Current size of memory cache: Indicates the current size of unused items of the excel calculation service manager in bytes.	MB	f.
	Excel calculation service workbook cache size: Indicates the current size of the Excel Calculation Services workbook cache.	MB	g. A high value for this measure indicates that the cache is adequately sized. A poorly-sized cache can adversely impact service performance, especially when multiple users try to access the same set of workbooks.

	Rendered charts requested: Indicates the number of chart requests per second.	Charts/Sec	h.
	Requests received: Indicates the number of requests received per second on Excel Calculation Services.	Received/Sec	i.
	Active requests: Indicates the number of requests being actively processed on Excel Calculation Services.	Number	j.

18.1.1.4 Excel Services Web Front End Test

The core components of Excel Services - the Excel Web Access, Excel Services, and Excel Calculation Services components - can be divided into components on the Web front-end server and those that live on a back-end application server. The **Web front end** includes **Excel Web Access** and **Excel Web Services**.

Excel Web Services is the Excel Services component that provides programmatic access to its Web service. You can develop applications that call Excel Web Services to calculate, set, and extract values from workbooks, as well as refresh external data connections. Using Excel Web Services, you can incorporate server-side workbook logic into an application, automate the updating of Excel workbooks and create application-specific user interfaces around server-side Excel calculation.

Using the **Excel Services Web Front End** test, you can track the number and rate of requests to the Excel Web Services component.

Purpose	Tracks the number and rate of requests to the Excel Web Services component
Target of the test	A Sharepoint Server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.
Outputs of the test	One set of results for the Sharepoint server monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Active requests: Indicates the current number of requests to the Excel Web Services component.	Number	
	Requests rate: Indicates the rate at which requests were received by the Excel Web Services component.	Requests/Sec	k.

18.1.1.5 Excel Web Access Test

Excel Web Access is an Excel Services Web Part in Office Sharepoint Server 2007 that renders (in other words, creates the HTML for) live Excel workbooks on a Web page, and allows the user to interact with those workbooks and explore them. Excel Web Access is the visible Excel Services component for the user.

This test measures the responsiveness of the Excel Web Access component to user requests.

Purpose	Measures the responsiveness of the Excel Web Access component to user requests		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Average chart image request time: Indicates the average time taken between the request for a chart image and the issuance of the response to the web browser by Excel Web Access.	Secs	An unusually high value for this measure is a cause for concern, as it indicates a slowdown in the responsiveness of the Excel Web Access component.

	Chart image request: Indicates the Indicates the number of requests for chart images that are served by Excel Web Access per second.	Requests/Sec	l.
	Excel web access average request time: Indicates the excel web access average request time.	Secs	m.

18.1.1.6 Office Server Search Indexer Catalogs Test

The MOSS 2007 Search feature is implemented using two MOSS services:

- **Indexing:** Responsible for crawling content sources and building index files.
- **Searching:** Responsible for finding all information matching the search query by searching the index files.

All searching is performed against the index files; if these files do not contain what the user is looking for, there will not be a match. So, the index files are critical to the success of the search feature of MOSS. The search functionality can be described in its simplest form as a Web page where the user defines his or her search query. The index service works together with the searching service to let you search Office Sharepoint Server content.

This test monitors the search queries to every content index on the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by MOSS 2007.

Purpose	Monitors the search queries to every content index on the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by MOSS 2007		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Failed queries: Indicates the number of queries to the content index that currently failed.	Number	Ideally, this value should be 0.
	Succeeded queries: Indicates the number of queries to the content index that succeeded.	Number	n. A high number of successful queries serves as a good indicator of the efficiency of the index and query services provided by Sharepoint.
	Queries: Indicates the number of queries currently executing on the content index.	Number	o.
	Documents filtered: Indicates the number of documents currently filtered in the content index.	Number	p.
	Index size: Indicates the current size of the content index.	Number	q.

18.1.1.7 Office Server Search Gatherer Test

The MOSS 2007 Search feature is implemented using two MOSS services:

- **Indexing:** Responsible for crawling content sources and building index files.
- **Searching:** Responsible for finding all information matching the search query by searching the index files.

All searching is performed against the index files; if these files do not contain what the user is looking for, there will not be a match. So, the index files are critical to the success of the search feature of MOSS. The search functionality can be described in its simplest form as a Web page where the user defines his or her search query.

The index role can be configured to run on its own MOSS server, or run together with all the other roles, such as the Web service, Excel Services and Forms Services. It performs its indexing tasks following this general e:

1. Sharepoint stores all configuration settings for the indexing in its database.
2. When activated, the index will look in Sharepoint's databases to see what content sources to index, and what type of indexing to perform, such as a full or incremental indexing.
3. The index service will start a program called the **Gatherer**, which is a program that will try to open the content that should be indexed.
4. For each information type, the **Gatherer** will need an Index Filter, or **IFilter**, that knows how to read text inside this particular type of information. For example, to read a MS Word file, an IFilter for **.DOC** is needed.
5. The Gatherer will receive a stream of Unicode characters from the IFilter. It will now use a small program called a Word Breaker; its job is to convert the stream of Unicode characters into words.
6. However, some words are not interesting to store in the index, such as "the", "a", and numbers; the Gatherer

will now compare each word found against a list of Noise Words. This is a text file that contains all words that will be removed from the stream of words.

7. The remaining words are stored in an index file, together with a link to the source. If that word already exists, only the source will be added, so one word can point to multiple sources.
8. If the source was information stored in Sharepoint, or a file in the file system, the index will also store the security settings for this source. This will prevent a user from getting search results that he or she is not allowed to open.
9. Since the success of an indexing operation also depends upon how the **Gatherer** program functions, administrators need to keep their eyes open for irregularities in the functioning of the gatherer, so that such anomalies are detected instantly, and corrected before they can stall the indexing process.

This test monitors the gatherer, and reports issues in its performance (if any).

Purpose	Monitors the gatherer, and reports issues in its performance (if any)		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Documents filtered: Indicates the number of documents filtered per second.	Documents/Sec	If this rate is decreasing over time, you should perform some troubleshooting to find out why your server is not filtering documents. Look for memory issues, processor issues, network issues, or site hit frequency rules that slow the gatherer process.
	Filtering threads: Indicates the current number of filtering threads in the system.	Number	r.

	Threads accessing the network: Indicates the number of threads currently waiting for a response from the filter process.	Number	These threads have sent or are sending their request off to the remote data store and are either waiting for a response or consuming the response and filtering it. You can distinguish the difference between actually waiting on the network versus filtering the document by looking at a combination of CPU usage and Network usage counters. If this number is consistently high then you are either network bound or you are bound by a "hungry" host. If you are not meeting your crawl freshness goals, you can either change your crawl schedules to minimize overlapping crawls or look the remote repositories you are crawling to optimize them for more throughput.
	Active queue length: Indicates the number of documents currently waiting for robot threads.	Number	s. If the value of this measure is not 0, then all threads should be filtered.
	Admin clients: Indicates the number of currently connected administrative clients.	Number	t.

	Reason to back off: A code describing why the gatherer service went into back-off state.	Number	The values that this measure can take and the states they denote are available below: 0 - Up and Running. 1 - High system IO traffic. 2 - High notifications rate. 3 - Delayed recovery in progress. 4 - Due to user activity. 5 - Battery low. 6 - Memory low. 99 - Some internal reason. During a back-off period, indexing is suspended. To manually back off the gatherer service, pause the search service. If the search service itself generates the back-off, an event will be recorded and the search service will be paused automatically. There is no automatic restart, so you must manually start the search service in order to end a back-off state. Note that there is little reason to start the search service until you have solved the problem that caused the back-off in the first place.
	Threads waiting for plug-ins: Indicates the number of threads currently waiting for plug-ins to complete an operation	Number	These threads have the filtered documents and are processing it in one of several plug-ins. This is when the index and property store are created. If you have a consistently high number for this counter, check the metrics reported by the **Office Server Search Archival Plugin** test for problem pointers.
	Delayed documents: Indicates the number of documents that were currently delayed due to site hit frequency rules.	Number	If you have a plethora of rules and this number is steadily increasing over time, consider relaxing or simplifying your site hit frequency rules. A very high number may indicate a conflict in the rules that the gatherer cannot resolve or follow with efficiency.

	Idle threads: Indicates the number of threads that are currently waiting for documents.	Number	These threads are not currently doing any work and will eventually be terminated. If you consistently have a more than *Max Threads/Hosts* idle threads you can schedule an additional crawl. If this number is 0 then you are starved. Do not schedule another crawl in this time period and analyze the durations of your crawls during this time to see if they are meeting your freshness goals. If your goals are not being met you should reduce the number of crawls.
	Hearbeats: Indicates the number of heartbeats per second.	Hearbeats/Sec	A heartbeat occurs once every 10 seconds while the service is running. If the service is not running there will be no heartbeat.

18.1.1.8 Sharepoint Publishing Cache Test

Object caching Office Sharepoint Server 2007 supports caching of certain page items, such as navigation data and data accessed through cross-list queries. Caching page items reduces the requirement to retrieve field data from the database every time a page is rendered. The caching system also caches complete field data for a page, excluding data for any Web Part controls on the page.

Using the statistics provided by this test, you can fine-tune your cache size, so as to maximize cache hits and minimize object discards.

Purpose	Helps you fine-tune your cache size, so as to maximize cache hits and minimize object discards		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Publishing cache hit ratio: Indicates the ratio of hits to misses on the publishing cache.	Percent	A hit ratio greater than 90% and a low object discard rate are generally good signs that the current size is satisfactory. However, you should also measure user response time for key operations to adjust this setting. If you set the size too large, you might waste valuable memory for the other caches, such as the ASP.NET output cache if it is used. Certain Web Parts, such as the Content Query Web Part, stores their XSLT stylesheets in the output cache. If the object cache size is set too large, ASP.NET might flush output cache memory to make room for it. CPU usage might increase after the flushing. This is especially important for a system that is running on a 32-bit operating system because each worker process is limited to 2 GB application memory space. If you set the object cache size limit too large, the IIS worker process (w3wp) can run out of memory.
	Object discards: Indicates the total number of items that have been removed from the publishing cache since the last measurement period due to cache compaction.	Number	

18.1.1.9 Office Server Search Schema Plugin Test

The Search feature of the MOSS 2007 not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites. This is what makes MOSS 2007 that much more valuable to users.

The **Archival** and **Schema** plugins are internal components of the MOSS Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the MOSS search feature is functioning, identify irregularities early, and fine-tune the MOSS server to ensure peak performance of the search engine.

The **Office Server Search Schema Plugin** test focuses on the schema plugin component, and helps assess its processing ability.

Purpose	Focuses on the schema plugin component, and helps assess its processing ability
Target of the test	A Sharepoint Server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.
Outputs of the test	One set of results for the Sharepoint server monitored

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Aliases mapped: Indicates the total number of aliases which have been currently mapped to the schema.	Number	
	Duplicate aliases: Indicates the number of aliases that the schema currently ignored as they are duplicates.	Number	
	Refresh count: Indicates the number of aliases that have been refreshed from the database, currently.	Number	
	Error documents: Indicates the number of documents that have currently returned errors from the plug-in.	Number	Ideally, this value should be 0.

18.1.1.10 Office Server Search Gatherer Projects Test

As already mentioned, the indexing service will start a program called the **Gatherer**, which is a program that will try to open the content that should be indexed. Using an iFilter, the **Gatherer** reads the content as Unicode characters, converts the characters into words, identifies words that are worth indexing, and stores them in the content indexes.

For each content index, this test reports critical performance statistics revealing the content processing ability of the gatherer.

Purpose	Reports critical performance statistics revealing the health of the gatherer
Target of the test	A Sharepoint Server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.

Outputs of the test	One set of results for each content index on the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Documents added: Indicates the number of document additions per second.	Documents/Sec	
	Error: Indicates the number of filtered documents which returned an error per second.	Documents/Sec	A low value is typically desired for this measure.
	Retries: Indicates the total number of times that access to a document has been retried.	Number	A high value of this measure indicates that the gatherer is attempting to access a document numerous times, without success. You should check the gatherer logs and identify the problem document. Then ensure that it has the correct extension and that you have the correct IFilter for it.
	Incremental crawls: Indicates the number of incremental crawls currently in progress.	Number	
	Waiting documents: Indicates the current queue size of unprocessed documents in the gatherer.	Number	A high value of this measure could indicate a processing bottleneck on the gatherer. If this measure returns the value 0 on the other hand, it could indicate that the gatherer is idle.

18.2 Monitoring Sharepoint 2010/2013

Figure 18.4 depicts the *Microsoft Sharepoint 2010/2013* monitoring model.

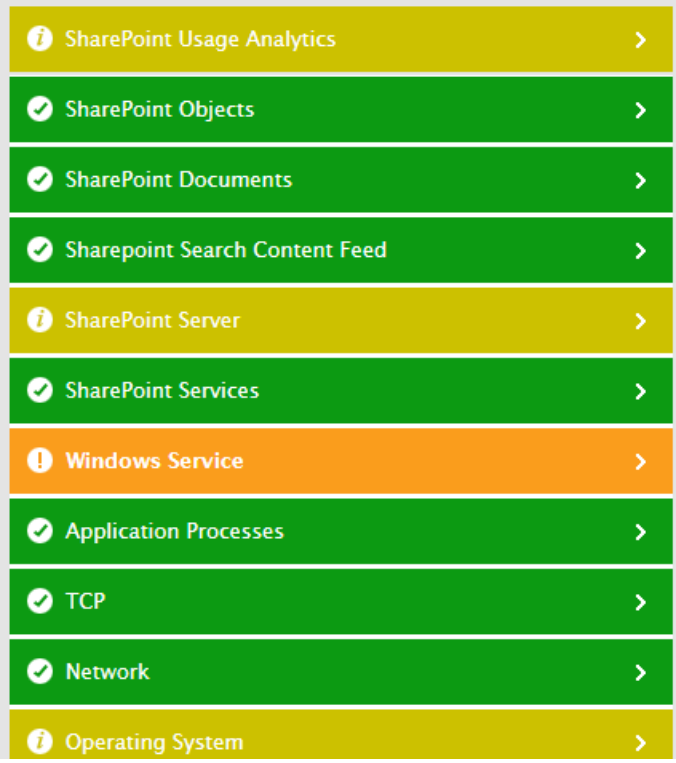

Figure 18.4: The layer model of Microsoft Sharepoint 2010/2013

Each layer of Figure 18.4 is mapped to a variety of tests that periodically check the health of the core components and services of the Sharepoint 2010/2013 server. Using the metrics reported by these tests, administrators can find quick and accurate answers for the following performance queries:

- Has the archival plugin marked too many documents for retry?
- Are too many documents in the archival plugin waiting for a queue?
- Have any errors occurred in index propagation?
- Is index reception error-free?
- Did any search query fail?
- Is query execution taking too long? If so, where is the query spending maximum time?
- Is the query CPU-intensive? If so, where is the query spending the maximum CPU time?
- Is any Sharepoint Foundation process overloaded? If so, which one is it?

- Is any Sharepoint Foundation process taking too long to execute requests? Which process is it?
- Which process is taking too much time to execute queries?
- Is the schema plugin able to process documents and properties quickly?
- Are there too many idle threads on the Sharepoint server?
- Is any thread waiting for a network response from the filter process?
- Have too many servers timed out?
- Was any slowdown noticed in document filtering? Is it due to site hit frequency rules? If so, how many documents were affected as a result?
- Is filtering failing for any document?

The sections that follow will only discuss the top 6 layers of Figure 18.4.

18.2.1 The Sharepoint Services Layer

The tests mapped to this layer shed light on the current status, overall health, and efficiency of the critical services offered by Sharepoint Foundation. This includes the Search archival and schema plugins, the search indexing mechanism, the search gatherer, and the critical Sharepoint Foundation processes.

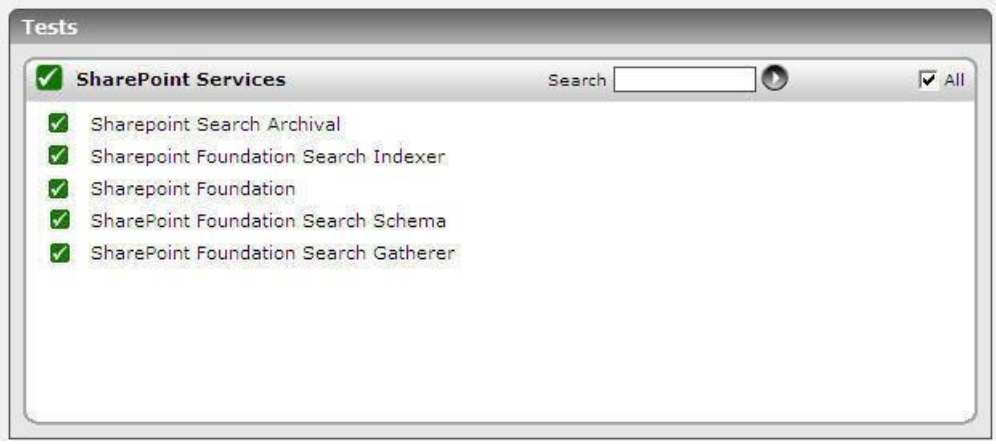

Figure 18.5: The tests mapped to the Sharepoint Services layer

18.2.1.1 Sharepoint Search Archival Test

The Search feature of the Microsoft Sharepoint server not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites.

The **Archival** and **Schema** plugins are internal components of the Microsoft Sharepoint server Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the Microsoft Sharepoint server search feature is functioning, identify irregularities early, and fine-tune the Microsoft Sharepoint server to ensure peak performance of the search engine.

This test monitors the performance of the Sharepoint Foundation Search Archival Plugin.

Purpose	Monitors the performance of the Sharepoint Foundation Search Archival Plugin		
Target of the test	A Sharepoint server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results each for the Sharepoint Server		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Upload queues available to filtering threads: Indicates the number of upload queues that are available to filtering threads in this plugin.	Number	
	Queues committing changes and completing uploads: Indicates the number of queues that are exclusively allotted for committing the changes and completing the uploads.	Number	
	Queues waiting to flush data to the property store: Indicates the number of queues that are waiting to flush data to the property store.	Number	A property store is a table of properties and their values that are used and maintained by the Search service. Each row in the table corresponds to a document in the full-text index.
	Queues being used by filtering threads: Indicates the number of queues that are being used by the filter threads in this plugin.	Number	

	Bulk insert sessions to the database server: Indicates the number of active bulk insert sessions to the database server.	Number	
	Documents processed: Indicates the number of documents that are processed in this plugin during the last measurement period.	Number	A high value is desired for this measure. If the value decreases steadily over a period of time, it indicates a performance bottleneck.
	Documents marked for retry by archival plugin: Indicates the number of documents that were marked for retry from this plugin during the last measurement period.	Number	Ideally the value of this measure should be low. A higher value may indicate a performance bottleneck.
	Documents waiting for a queue: Indicates the number of documents that were waiting for a queue during the last measurement period.	Number	Ideally the value of this measure should be low. A higher value may indicate a performance bottleneck.

18.2.1.2 Sharepoint Foundation Search Indexer Test

Using the Search feature of Sharepoint 2010, users can easily find the information they need in Sharepoint Foundation Sites.

The key components of Sharepoint's Search architecture are as follows:

- **Indexer**: Also referred to as the **Crawl Component** or the **Crawler**, the **Indexer** is solely responsible for building indexes. The indexers enumerate the source content and pass text information to the relevant index partition on the query server. The indexer also indexes any metadata to the search property database and updates the crawl status in the crawl database.
- **Crawl Database**: The **Crawl Database** tracks what needs to be crawled and what has been crawled.
- **Query Component**: Commonly referred to as the **Query Server**, this component will perform a search against an index created by the indexer. The query component will apply such things as security trimming, best bets, relevancy, removes duplicates, etc.
- **Index partition**: Indexes can be split into multiple partitions called **index partitions** to improve the amount of time it takes to perform a search by the query component. For every query component there will be a single index partition that is queried by the query component.
- **Index Partition Mirror**: Mirrors can be created for the index partitions. These mirrors again provide the ability to provide redundancy and better search result performance.
- **Property Database**: This database stores metadata and security information items in the index. The property database will be associated with one or more query components and is used as part of the query process. These properties will be populated as part of the crawling process which creates the index.
- **Search Admin Database**: The **Search Administration Database** is mostly responsible for managing information associated to the configuration and topology of the Sharepoint Search service. There will only be one instance of this database for each Search Application Service instance.

Figure 18.6 depicts how these components work together to implement the search feature of Sharepoint 2010.

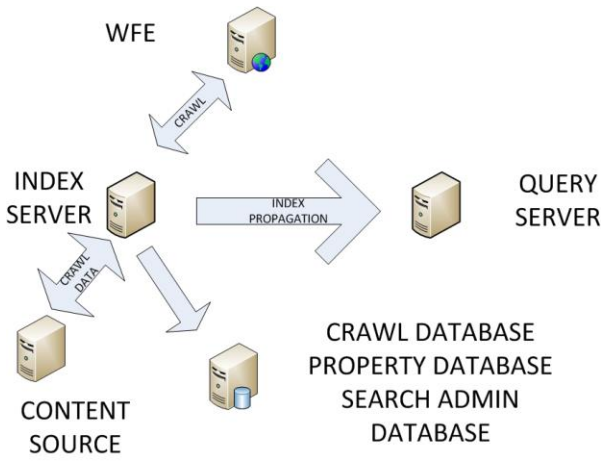

Figure 18.6: How Search works in Sharepoint 2010?

When a user enters a search query on a Web Front End (WFE) server, the query server processes the query. While processing, the query server retrieves the information that fulfills the query criteria from the index partition stored on its local file system, and also retrieves metadata information from the search property database. The index partition on the other hand, receives text information from the indexers that enumerate the source content. Once the desired query results are available, the query server packages the results, and delivers the results back to the requesting WFE server.

The success of Sharepoint Search feature therefore depends upon how quickly the query server processes the queries it receives, and how effective the index files built by the indexer are.

This test monitors the search queries to the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by Microsoft Sharepoint Server.

Purpose	Monitors the search queries to the Sharepoint server, promptly reports query failures, and thus reveals the overall efficiency of the Search feature offered by Microsoft Sharepoint Server		
Target of the test	A Sharepoint server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active connections to the indexer plugin: Indicates the number of currently active connections to this indexer plugin.	Number	
	Index size: Indicates the current size of the content index that is being managed by this indexer plugin.	Number	
	Tasks in queue of propagation task sender: Indicates the number of tasks that were in queue of the propogation task sender.	Number	
	Tasks in queue of index receiver: Indicates the number of tasks that were in queue of the index receiver.	Number	
	Tasks in queue of index propagator: Indicates the number of tasks that were in queue of the index propogator.	Number	

	Errors in Index propagation: Indicates the number of errors in index propagation during the last measurement period.	Number	Once the indexer builds the indexes, it propagates/pushes the index files from the index server to the query server. The indexer then waits for the query server to absorb the index, after which it acknowledges that the documents are successfully crawled. Ideally, no errors should occur in this process - i.e., the value of this measure should be ideally 0. The incidence of one or more errors can adversely impact the user experience with Sharepoint's Search mechanism.
	Errors in Index reception: Indicates the number of errors in index reception during the last measurement period.	Number	Ideally, no errors should occur in this process - i.e., the value of this measure should be ideally 0.
	Indexes received successfully: Indicates the number of indexes that were received successfully by this indexer plugin during the last measurement period.	Number	A high value is desired for this measure. A sudden/gradual decrease in the value of this measure may indicate a performance bottleneck in the Microsoft Server Search Indexer plugin.
	Indexes propagated successfully: Indicates the number of indexes that were propagated successfully by this indexer plugin during the last measurement period.	Number	A high value is desired for this measure. A sudden/gradual decrease in the value of this measure may indicate a performance bottleneck in the Microsoft Server Search Indexer plugin.
	Documents filtered: Indicates the number of documents that were filtered by this indexer plugin during the last measurement period.	Number	
	Documents in indexes that are being propagated: Indicates the number of documents in indexes which were being propagated by this indexer plugin during the last measurement period.	Number	

	Queries handled: Indicates the number of queries that were handled on the content index during the last measurement period.	Number	
	Successful queries: Indicates the number of queries that were processed successfully during the last measurement period.	Number	A high value is desired for this measure.
	Failed Queries: Indicates the number of queries that failed to process during the last measurement period.	Number	Ideally, the value of this measure should be zero.

	Avg latency of queries in the last minute: Indicates the average latency with which the queries were processed in the last minute.	Secs	Ideally, when an end user executes a query, results should be returned in less than one second. If this is not the case routinely, then end user experience with the Search feature is bound to suffer. The common reasons for poor query performance and their recommended solutions are as follows: • **One or more index partitions contain more than 10 million documents**: Add an additional index partition, and if possible, an additional index partition mirror. If all query servers already contain an active and a mirrored index partition, add more query servers. • **One or more query servers are memory bound and/or paging virtual memory on disk:** Add additional memory to the query server. Ensure that the query server has enough RAM to store 33% of each index partition (present on the query server) in memory. • **Query preformance suffers during the first few queries after the server is rebooted or during crawl processing and index propagation:** Ensure that the physical disk where the index partition is stored is capable of providing 2,000 IOPS for each index partition. • **Query latency is high though all query servers are adequately sized:** Ensure that the property database server has enough RAM available to store 33% of the property store tables in memory. Make sure that the property database server is not CPU or disk I/O bound. Additional property database servers or property databases can also be added based on need.

	Execution time to create a query restriction: Indicates the average execution time to create a query restriction.	Secs	Whenever query latency is very high - i.e., if the **Avg latency of queries in the last minute** measure reports a very high value - then, you can compare the values of these measures to understand where the query is spending too much time. You can thus identify the bottleneck areas and accordingly decide on the action to be taken to improve query performance.
	Execution time to resolve query: Indicates the average execution time to resolve a query.	Secs	
	Execution time to get row results of a query: Indicates the average execution time to get row results of a query.	Secs	
	Execution time spent in other parts of a query: Indicates the average time taken to create a query restriction.	Secs	
	CPU time to create a query restriction: Indicates the average CPU time that is required to create a query restriction.	Secs	If a query is found to be CPU-intensive, you can compare the values of these measures to determine where the query is consuming CPU excessively.
	CPU time to resolve a query: Indicates the average CPU time taken to resolve a query.	Secs	
	CPU time to get row results for a query: Indicates the average CPU time taken to get row results of a query.	Secs	
	CPU time spent in other parts of a query: Indicates the average CPU time taken to execute other parts of the query.	Secs	

18.2.1.3 Sharepoint Foundation Test

Microsoft Sharepoint Foundation is the essential solution for organizations that need a secure, manageable, web-based collaboration platform. It serves as the basis for Sharepoint server and offers out of the box elements such as

blogs, wikis, team workspaces, and document libraries, providing users with the ideal way to share information and collaborate within a customized website. In addition, it provides services such as Business Data Connectivity services to integrate external data, basic search services and workflow services.

This test auto-discovers the Sharepoint processes, and for each process, reports the workload on the process and how efficiently that process handles the load. This way, the test leads you to the processes that are very busy and provides pointers to what could be keeping them busy.

Purpose	Auto-discovers the Sharepoint processes, and for each process, reports the workload on the process and how efficiently that process handles the load		
Target of the test	A Sharepoint server 2010/2013		
Agent deploying the test	An internal/remote agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for each Sharepoint Foundation process		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active threads: Indicates the number of threads that are currently executing in Sharepoint code of this process.	Number	Many active threads is an indicator of a bottleneck.
	Incoming page requests: Indicates the number of incoming requests to access a particular page in the last second.	Number	This measure is a good indicator of the workload on this process.
	Requests being processed currently: Indicates the requests that are currently processed by this Sharepoint process.	Reqs	

	Avg execution time of requests processed: Indicates the average time taken by this process for executing the requests that are processed in the last second.	Secs	Ideally, this value should be low. If the value of this measure increases steadily, then it indicates a performance bottleneck.
	Requests rejected: Indicates the number of page requests that were rejected by this process during the last second.	Number	Ideally, the value of this measure should be zero.
	Requests responded to by the Sharepoint server: Indicates the number of page requests that were responded by this Sharepoint process during the last second.	Number	
	Throttled page requests: Indicates the number of page requests that have been throttled by this process during the last measurement period.	Number	
	Executing SQL queries: Indicates the number of SQL queries that are currently executing on this Sharepoint server.	Number	
	Query execution time: Indicates the average time taken by this Sharepoint server to execute the SQL queries.	Secs	If the time taken to execute a query is *high*, it indicates that the query is unoptimal or there may be a database slowdown.
	Native heaps in use: Indicates the number of native heaps that are currently in use by this Sharepoint process.	Number	

	Native heaps allocated by process: Indicates the number of native heaps that are allocated by this Sharepoint process.	Number	
	Global heap size: Indicates the size of the global heaps that are used by this Sharepoint process for cache related activity.	MB	
	Size of all per thread native heaps: Indicates the size of the native heaps that are used by all the threads that are being executed by this Sharepoint process.	MB	

18.2.1.4 Sharepoint Foundation Search Schema Test

The Search feature of the Microsoft Sharepoint server not only makes it possible to search through content, documents, and people within the Sharepoint sites, but also through external sources such as Windows file shares, public Microsoft Exchange server folders, and standard web sites.

The **Archival** and **Schema** plugins are internal components of the Microsoft Sharepoint server Search engine, typically responsible for processing the metadata of indexed documents. By monitoring these components, administrators can efficiently evaluate how well the Microsoft Sharepoint server search feature is functioning, identify irregularities early, and fine-tune the Microsoft Sharepoint server to ensure peak performance of the search engine.

This test monitors the performance of the Sharepoint Foundation Search Schema and Alias Mapping Plugin, and enables an informed assessment of its processing ability.

Purpose	Monitors the performance of the Sharepoint Foundation Search Schema and Alias Mapping Plugin, and enables an informed assessment of its processing ability.
Target of the test	A Sharepoint server 2010/2013
Agent deploying the test	An internal/remote agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.
Outputs of the test	One set of results each for the *ProfileImport* and *Portal_Content* instances

Measurements made by the test	Measurement	Measurement Unit	Interpretation
	Documents processed by schema plugin: Indicates the number of documents that are processed by this schema plugin during the last measurement period.	Number	
	Properties processed by schema plugin: Indicates the number of properties that are processed by this schema plugin during the last measurement period.	Number	
	Aliases loaded: Indicates the number of aliases that have been currently loaded to this schema plugin.	Number	
	Aliases have been mapped: Indicates the total number of aliases that have been currently mapped to this schema plugin during the last measurement period.	Number	
	Aliases ignored as they are duplicates: Indicates the number of aliases that the schema currently ignored as they are duplicates during the last measurement period.	Number	
	Aliases refreshed from the database: Indicates the number of aliases that have been refreshed from the database during the last measurement period.	Number	

18.2.1.5 Sharepoint Foundation Search Gatherer Test

The search functionality can be described in its simplest form as a Web page where the user defines his or her search query. The index role can be configured to run on its own Microsoft Sharepoint server, or run together with all the other roles, such as the Web service, Excel Services and Forms Services. It performs its indexing tasks following this general workflow:

1. Sharepoint stores all configuration settings for the indexing in its database.
2. When activated, the index will look in Sharepoint's databases to see what content sources to index, and what type of indexing to perform, such as a full or incremental indexing.
3. The index service will start a program called the Gatherer, which is a program that will try to open the content that should be indexed.
4. For each information type, the Gatherer will need an Index Filter, or IFilter, that knows how to read text inside this particular type of information. For example, to read a MS Word file, an IFilter for .DOC is needed.
5. The Gatherer will receive a stream of Unicode characters from the IFilter. It will now use a small program called a Word Breaker; its job is to convert the stream of Unicode characters into words.
6. However, some words are not interesting to store in the index, such as "the", "a", and numbers; the Gatherer will now compare each word found against a list of Noise Words. This is a text file that contains all words that will be removed from the stream of words.
7. The remaining words are stored in an index file, together with a link to the source. If that word already exists, only the source will be added, so one word can point to multiple sources.
8. If the source was information stored in Sharepoint, or a file in the file system, the index will also store the security settings for this source. This will prevent a user from getting search results that he or she is not allowed to open.

Since the success of an indexing operation also depends upon how the Gatherer program functions, administrators need to keep their eyes open for irregularities in the functioning of the gatherer, so that such anomalies are detected instantly, and corrected before they can stall the indexing process.

This test monitors the performance of the Sharepoint Foundation Search Gatherer, and reports issues in its performance (if any).

Purpose	Monitors the performance of the Sharepoint Foundation Search Gatherer, and reports issues in its performance (if any)		
Target of the test	A Sharepoint server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results each for the *ProfileImport* and *Portal_Content* instances		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Filtering threads in the system: Indicates the current number of filtering threads in the system.	Number	
	Threads waiting for documents: Indicates the number of threads that are currently waiting for documents.	Number	These threads are not currently doing any work and will eventually be terminated. If you consistently have more than *Max Threads/Hosts* idle threads you can schedule an additional crawl. If this number is 0 then you are starved. Do not schedule another crawl in this time period and analyze the durations of your crawls during this time to see if they are meeting your freshness goals. If your goals are not being met you should reduce the number of crawls.
	Threads waiting for network response from the filter process: Indicates the number of threads that were waiting for a response from the filter process.	Number	If you figure out that there is no activity that is taking place as far as this measure is concerned, and if the value of this measure is equal to the *Filtering threads in system* measure, it indicates a network issue or the unavailability of the server that is crawling into.
	Threads committing transactions: Indicates the number of threads that are committing transactions.	Number	
	Threads waiting for plug-ins to complete an operation: Indicates the number of threads currently waiting for plug-ins to complete an operation.	Number	These threads have the filtered documents and are processing it in one of several plug-ins. This is when the index and property store are created.
	Threads loading transactions from persisted crawl queue: Indicates the number of transactions that are loaded from the persisted crawl queue.	Number	
	Threads processing links: Indicates the number of threads that are processing links.	Number	

	Filtering processes in the system: Indicates the number of filtering processes that are active in the system.	Number	
	Filter objects in the system: Indicates the number of filter objects in the system.	Number	
	Documents waiting for robot threads: Indicates the number of documents that are waiting for robot threads.	Number	If the value of this measure is *0*, then it implies that all the threads are filtering threads.
	Currently connected admin clients: Indicates the number of currently connected admin clients.	Number	
	Amount of resources allowed for the Gatherer service: Indicates the amount of resources that the Gatherer service is allowed to use.	Number	
	Servers recently accessed by the system: Indicates the number of servers that were recently accessed by the system.	Number	
	Servers currently unavailable: Indicates the number of servers that are currently unavailable to the system.	Number	A server becomes unavailable if the requests made to the server is timed out.
	Available cached stemmer instances: Indicates the number of cached stemmer instances in the system.	Number	Stemmers are nothing but components shared by the Search and Indexing engines that generate inflected forms for a word. Too many stemmer instances that are cached may indicate a resource usage problem.

	System I/O rate: Indicates the rate at which the system IO disk traffic is detected during back off period.	KB/Sec	During a back-off period, indexing is suspended. To manually back off the gatherer service, pause the search service. If the search service itself generates the back-off, an event will be recorded and the search service will be paused automatically. There is no automatic restart, so you must manually start the search service in order to end a back-off state. Note that there is little reason to start the search service until you have solved the problem that caused the back-off in the first place.
	Timeouts: Indicates the number of timeouts detected by the system during the last measurement period.	Number	Ideally, this value should be zero.
	Documents filtered: Indicates the rate at which the documents are filtered in the system.	KB/Sec	If this rate is decreasing over time, you should perform some troubleshooting to find out why your server is not filtering documents. Look for memory issues, processor issues, network issues, or site hit frequency rules that slow the gatherer process.
	Documents successfully filtered: Indicates the rate at which the documents are filtered successfully in the system.	KB/Sec	
	Documents delayed due to site hit frequency rules: Indicates the number of documents that were currently delayed due to site hit frequency rules.	Number	If you have a plethora of rules and this number is steadily increasing over time, consider relaxing or simplifying your site hit frequency rules. A very high number may indicate a conflict in the rules that the gatherer cannot resolve or follow with efficiency.
	Document entries currently in memory: Indicates the number of document entries that are currently available in the memory of the system.	Number	
	Documents filtered: Indicates the total number of documents filtered in the system during the last measurement period.	Number	

	Documents successfully filtered: Indicates the total number of documents that are successfully filtered in the system during the last measurement period.	Number	If the value of this measure is less than the value of the *Documents filtered* measure, use the gatherer logs to figure out the cause for the documents that are attempting to be filtered but are failing.

18.2.1.6 Distributed Cache Service Test

SharePoint uses the Distributed Cache to store data for very fast retrieval across all entities. The Distributed Cache service provides in-memory caching services to several features in SharePoint Server 2013. Some of the features that use the Distributed Cache service include:

- Newsfeeds
- Authentication
- pOneNote client access
- Security Trimming
- Page load performance

Besides services, several caches that exist in Sharepoint 2013 depend upon the Distributed Cache service for their proper functioning.

Any server in the farm running the Distributed Cache service is known as a **cache host**. A **cache cluster** is the group of all cache hosts in a SharePoint Server 2013 farm. A cache host joins a cache cluster when a new application server running the Distributed Cache service is added to the farm. When using a cache cluster, the Distributed Cache spans all application servers and creates one cache in the server farm. The total cache size is the sum of the memory allocated to the Distributed Cache service on each of the cache hosts.

If the distributed cache is not able to service requests efficiently, it is bound to significantly impact the performance of the dependent services/caches. Furthermore, it will add significantly to the processing overheads of Sharepoint, as poor cache usage translates into increased database accesses. If this is to be prevented, administrators should keep a close watch on the distributed cache's ability to service requests, rapidly detect poor cache usage patterns, and accurately pinpoint the reason for the same – is it because adequate objects are not cached in the distributed cache? If so, why? Is it owing to insufficient cache size? Will allocating more memory to the cache help or should more servers be added to the cache cluster? The **Distributed Cache Service** test helps answer all these questions! This test continuously monitors the requests to the cache, reports the count of requests serviced and rejected by the cache, and thus enables administrators to ascertain how well the cache is utilized. In the event of poor cache usage, close scrutiny of these test results will provide administrators with useful pointers to what is impeding cache usage and whether/not right-sizing the cache will help clear the bottleneck.

Purpose	Continuously monitors the requests to the cache, reports the count of requests serviced and rejected by the cache, and thus enables administrators to ascertain how well the cache is utilized
Target of the test	A Sharepoint Server 2013
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results each for the Sharepoint server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Cache data transferred rate: Indicates the number of cached entries transferred per second.	Number	
	Cache hit count: Indicates the number of requests serviced by the cache during the last measurement period.	Number	A high value is desired for this measure. A sudden/steady dip in this value indicates that the cache is unable to process requests, thereby increasing direct database accesses.

	Cache hit ratio: Indicates the percentage of requests that were serviced by the cache.	Percent	A high value is desired for this measure. A sudden/steady drop in this value is indicative of poor cache usage, which in turn can cause direct database accesses to increase and strain the database. One of the common reasons for a low cache hit ratio is insufficient memory allocation to the cache. In the absence of adequate memory resources, the cache may not be able to hold many frequently-accessed objects within, and may hence not be able to service many requests. Under such circumstances, you may want to consider allocating more memory to the cache. Here are a few recommendations from Microsoft with regard to how to size the distributed cache: • The Distributed Cache service actually uses twice the allocated amount of RAM, using the extra for housekeeping. In a small farm with fewer than 10,000 users, Microsoft recommends allocating 1GB of RAM for the Distributed Cache. This can be either a dedicated server or collocated with other SharePoint services, such as the Web Application Service. Beyond this the recommendation is using dedicated servers for the cache. A medium farm with fewer than 100,000 users should look to allocate around 2.5GB for the cache, and a large farm with up to 500,000 users should set aside around 12GB of RAM allocated for the cache. • It is a very strong recommendation that you should not allocate more than 16GB to any one Cache Host. This may cause the Cache Service to timeout during housekeeping operations and become unresponsive for several seconds at a time. If you need a cache size of greater than 16GB, it is better to use multiple servers in a Cache Cluster. You can have up to a maximum of 16 hosts in a Cache Cluster.

	Cache miss count: Indicates the number of requests that were not serviced by the cache since the last measurement period.	Number	Ideally, the value of this measure should be low. A sudden/steady increase in this value is indicative of poor cache usage, which in turn can cause direct database accesses to increase and strain the database.
	Cache read requests rate: Indicates the number of read requests to the cache per second, during the last measurement period.	Number	A high value for these measures is often indicative of heavy load on the distributed cache. In such a situation, for better cache performance, it is recommended that you opt for the **dedicated mode** of cache deployment. In this mode, all services other than the Distributed Cache service are stopped on the application server that runs the Distributed Cache service, thus ensuring that all critical resources on the server are at at the disposal of the distributed cache. This in turn, will help the cache handle the load efficiently!
	Cache write requests rate: Indicates the number of write requests to the cache per second, during the last measurement period.	Number	
	Total cache read requests: Indicates the total number of read requests received by the cache since the last measurement period.	Number	
	Total cache write requests: Indicates the total number of write requests received by the cache since the last measurement period.	Number	

18.2.2 The SharePoint Server Layer

The tests mapped to this layer report the composition and state of the farm in which the target SharePoint server operates, captures health analyzer alerts related to server status, reads ULS logs and reports problem conditions logged therein, and also tracks timer jobs that run on the server.

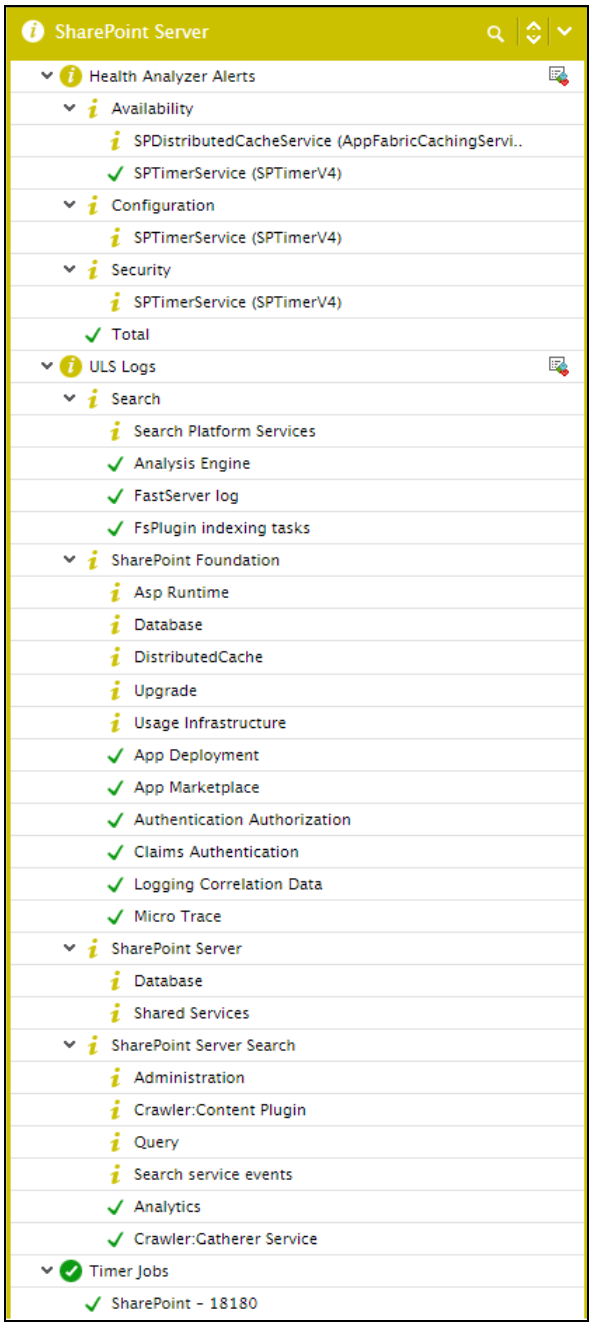

Figure 18.7: The tests mapped to the SharePoint Server layer

18.2.2.1 Timer Jobs Test

A timer job runs in a specific Windows service for SharePoint 2013. Timer jobs perform infrastructure tasks for the Timer service, such as clearing the timer job history and recycling the Timer service. Timer jobs also perform tasks for web applications, such as sending email alerts. A timer job contains a definition of the service to run and specifies how frequently the service is started. The SharePoint Timer service (SPTimerv4) runs timer jobs.

By tracking timer jobs run for web applications, administrators can quickly detect job failures. This is exactly what the **Timer Jobs** test does! For each web application, this test reports the count of successful and failed timer jobs.

Administrators are proactively alerted if even a single timer job fails! Deeps diagnostics reported by the test also provides details about the failed jobs, thereby enabling you to troubleshoot better.

Purpose	Reports the count of successful and failed timer jobs		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **DOMAIN**, **DOMAIN USER**, **PASSWORD**, and **CONFIRM PASSWORD** – This test requires the credentials of a domain user to run and collect farm-related metrics from the target SharePoint server. Therefore, specify the domain to which that user belongs in the **DOMAIN** text box, and then, enter the credentials of the user in the **DOMAIN USER** and **PASSWORD** text boxes. To confirm the password, retype it in the **CONFIRM PASSWORD** text box. 5. **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.		
Outputs of the test	One set of results for each web application on the SharePoint server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Successful jobs: Indicates the number of timer jobs for this web application that were successful.	Number	
	Failed jobs: Indicates the number of jobs run for this web application that failed.	Number	Ideally, the value of this measure should be 0. If a non-zero value is reported, use the detailed diagnosis of this measure to know which jobs failed, on which server, and when the failure occurred.

18.2.2.2 ULS Logs Test

The Unified Logging Service (ULS) writes SharePoint Foundation events to the SharePoint Trace Log, and stores them in the file system. ULS logging, when implemented effectively, can provide very useful information for developers, server administrators, and support personnel alike. The ULS logs can collect data at varying levels depending on the logging settings. Typically, every ULS log record indicates the diagnostic area and the specific category under the diagnostic area that has been traced.

Using the **ULS Logs** test, you can capture the number and nature of messages of various types and levels that are logged in the ULS logs. These statistics are grouped by area and category, so that you can instantly isolate the problem-prone categories and the diagnostic areas they belong to. This way, you will be enabled to investigate issues more effeciently and resolve them quickly.

For this test to run and report metrics, you need to enable the collection of health data on the SharePoint server. To know how to achieve this, refer to Section 18.2.2.2.1 of this document.

Purpose	Using the **ULS Logs** test, you can capture the number and nature of messages of various types and levels that are logged in the ULS logs
Target of the test	A Sharepoint Server
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **DOMAIN**, **DOMAIN USER**, **PASSWORD**, and **CONFIRM PASSWORD** – This test requires the credentials of a domain user to run and collect farm-related metrics from the target SharePoint server. Therefore, specify the domain to which that user belongs in the **DOMAIN** text box, and then, enter the credentials of the user in the **DOMAIN USER** and **PASSWORD** text boxes. To confirm the password, retype it in the **CONFIRM PASSWORD** text box. 5. **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.

Outputs of the test	One set of results for each category in each diagnostic area First-level descriptor: Area Second-level descriptor: Category		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Medium severity messages: Indicates the number of messages of a medium severity that are currently logged in the ULS log for this category of this area.	Number	*Medium severity messages* represent all messages except Verbose and VerboseEx messages. Such messages record all high-level information about operations that were performed. These messages provide enough detail to construct the data flow and sequence of operations. Administrators or support professionals could use such messages to troubleshoot issues. Use the detailed diagnosis of this measure to view the complete description of the medium severity messages logged in the ULS log for a specific category of an area.
	High severity messages: Indicates the number of messages of a high severity that are currently logged in the ULS log for this category of this area.	Number	*High severity messages* record all events that are unexpected but which do not stop the processing of a solution. Use the detailed diagnosis of this measure to view the complete description of the high severity messages logged in the ULS log for a specific category of an area.
	Monitorable messages: Indicates the number of monitorable messages that are currently logged in the ULS log for this category of this area.	Number	*Monitorable messages* capture all unrecoverable events that limit the functionality of the solution but do not stop the application. Use the detailed diagnosis of this measure to view the complete description of the monitorable messages logged in the ULS log for a specific category of an area.
	Warning messages: Indicates the number of warning messages that are currently logged in the ULS log for this category of this diagnostic area.	Number	This message type indicates a potential problem or issue that might require attention. You should review and track warning messages for patterns over time. Use the detailed diagnosis of this measure to view the complete description of the warning messages logged in the ULS log for a specific category of an area.

	Error messages: Indicates the number of error messages that are currently logged in the ULS log for this category of this area.	Number	A message of this type indicates an urgent condition. You should investigate all error events. Use the detailed diagnosis of this measure to view the complete description of the error messages logged in the ULS log for a specific category of an area.
	Critical messages: Indicates the number of critical messages that are currently logged in the ULS log for this category of this area.	Number	This message type indicates a serious error that has caused a major failure in the solution. Use the detailed diagnosis of this measure to view the complete description of the critical messages logged in the ULS log for a specific category of an area.
	Unexpected messages: Indicates the number of unexpected messages logged in the ULS log for this category of this area.	Number	Unexpected messages record events that cause solutions to stop processing. Use the detailed diagnosis of this measure to view the complete description of the unexpected messages logged in the ULS log for a specific category of an area.

Use the detailed diagnosis of the *Medium severity messages* measure to view the complete description of the medium severity messages logged in the ULS log for a specific category of an area.

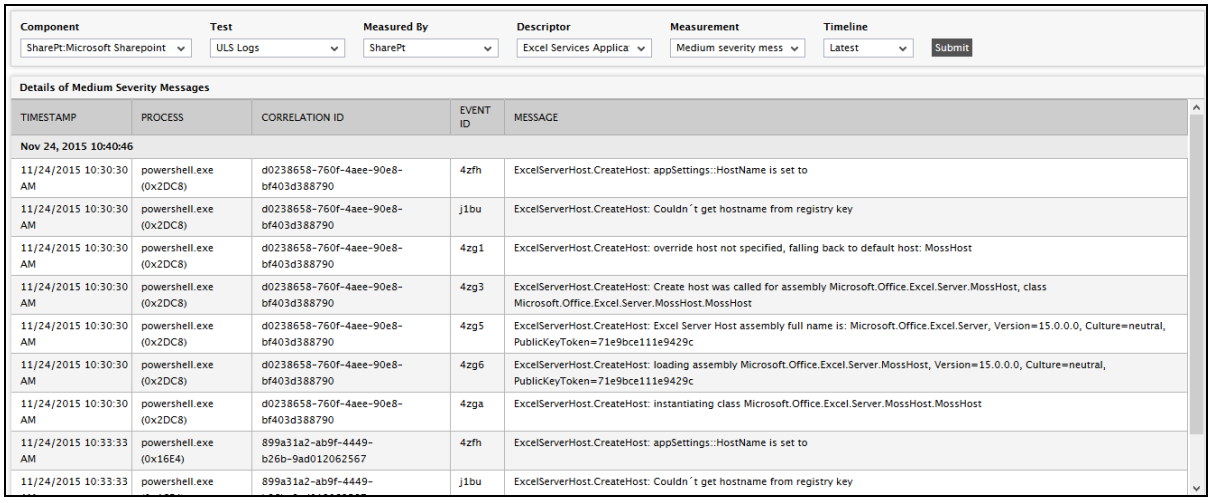

Figure 18.8: The detailed diagnosis of the Medium severity messages measure

Use the detailed diagnosis of the *High severity messages* measure to view the complete description of the high severity messages logged in the ULS log for a specific category of an area.

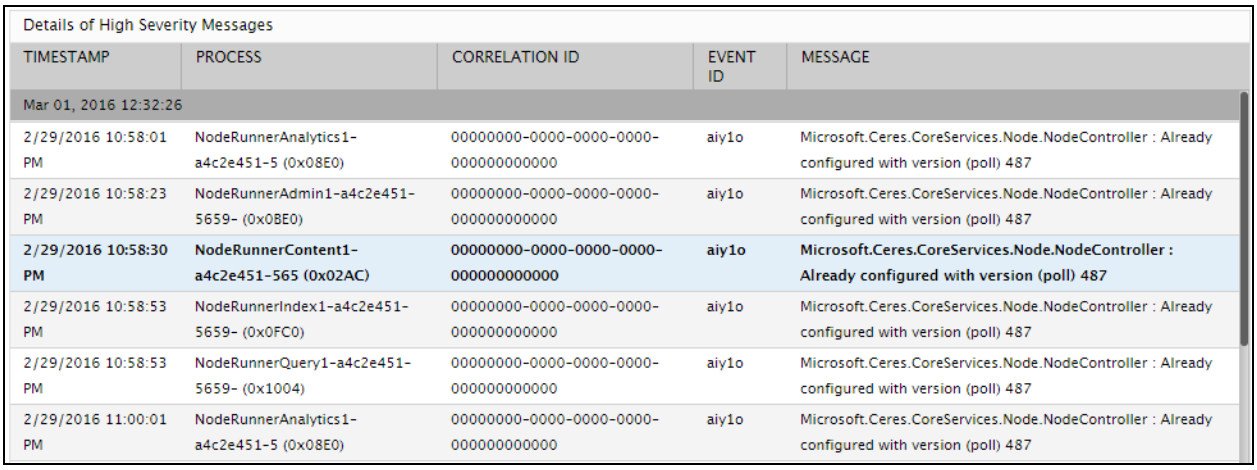

Details of High Severity Messages

TIMESTAMP	PROCESS	CORRELATION ID	EVENT ID	MESSAGE
Mar 01, 2016 12:32:26				
2/29/2016 10:58:01 PM	NodeRunnerAnalytics1-a4c2e451-5 (0x08E0)	00000000-0000-0000-0000-000000000000	aiy1o	Microsoft.Ceres.CoreServices.Node.NodeController : Already configured with version (poll) 487
2/29/2016 10:58:23 PM	NodeRunnerAdmin1-a4c2e451-5659- (0x0BE0)	00000000-0000-0000-0000-000000000000	aiy1o	Microsoft.Ceres.CoreServices.Node.NodeController : Already configured with version (poll) 487
2/29/2016 10:58:30 PM	**NodeRunnerContent1-a4c2e451-565 (0x02AC)**	**00000000-0000-0000-0000-000000000000**	**aiy1o**	**Microsoft.Ceres.CoreServices.Node.NodeController : Already configured with version (poll) 487**
2/29/2016 10:58:53 PM	NodeRunnerIndex1-a4c2e451-5659- (0x0FC0)	00000000-0000-0000-0000-000000000000	aiy1o	Microsoft.Ceres.CoreServices.Node.NodeController : Already configured with version (poll) 487
2/29/2016 10:58:53 PM	NodeRunnerQuery1-a4c2e451-5659- (0x1004)	00000000-0000-0000-0000-000000000000	aiy1o	Microsoft.Ceres.CoreServices.Node.NodeController : Already configured with version (poll) 487
2/29/2016 11:00:01 PM	NodeRunnerAnalytics1-a4c2e451-5 (0x08E0)	00000000-0000-0000-0000-000000000000	aiy1o	Microsoft.Ceres.CoreServices.Node.NodeController : Already configured with version (poll) 487

Figure 18.9: The detailed diagnosis of the High severity messages measure

The detailed diagnosis of all the measures reported by the **ULS Logs** test also point to a correlation ID. Correlation Ids are GUIDs assigned to events which transpire during the lifecycle of a resource request. An administrator can then use the correlation Id to locate and isolate the request in the ULS log. Correlation Ids also span machine boundaries, so in the event a conversion crosses a machine boundary, such as a Web front-end calling a Web service on an application server, etc., a unique Correlation Id is assigned to the conversation enabling a complete view of the request and what transpired during the operation. This way, administrators can dig deeper and troubleshoot issues more effectively.

18.2.2.2.1 Configuring the eG Agent to Monitor ULS Logs

The **ULS Logs** test can run and report metrics only if **health data collection** is enabled on the target SharePoint server. To achieve this, follow the steps below:

1. In the SharePoint management console, select the **Monitoring** node under **Central Administration.** Then, click on the **Configure usage and health data collection** option under **Reporting** (see Figure 18.10).

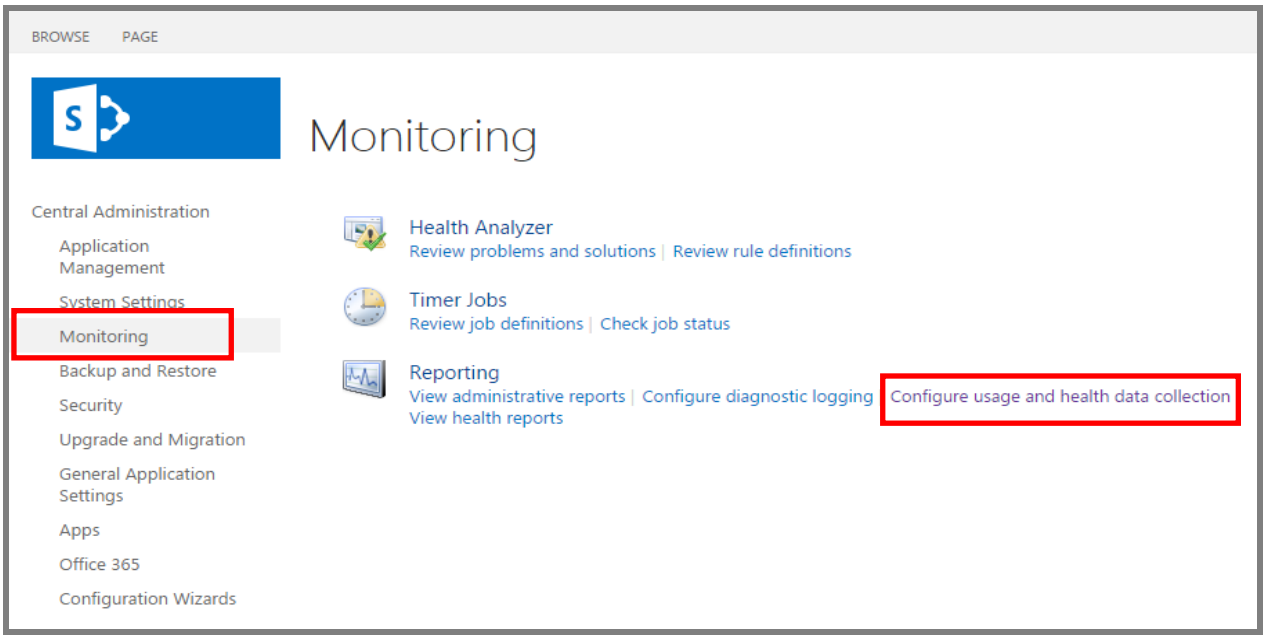

Figure 18.10: Selecting the Configure usage and health data collection option

2. Figure 18.11 will then appear.

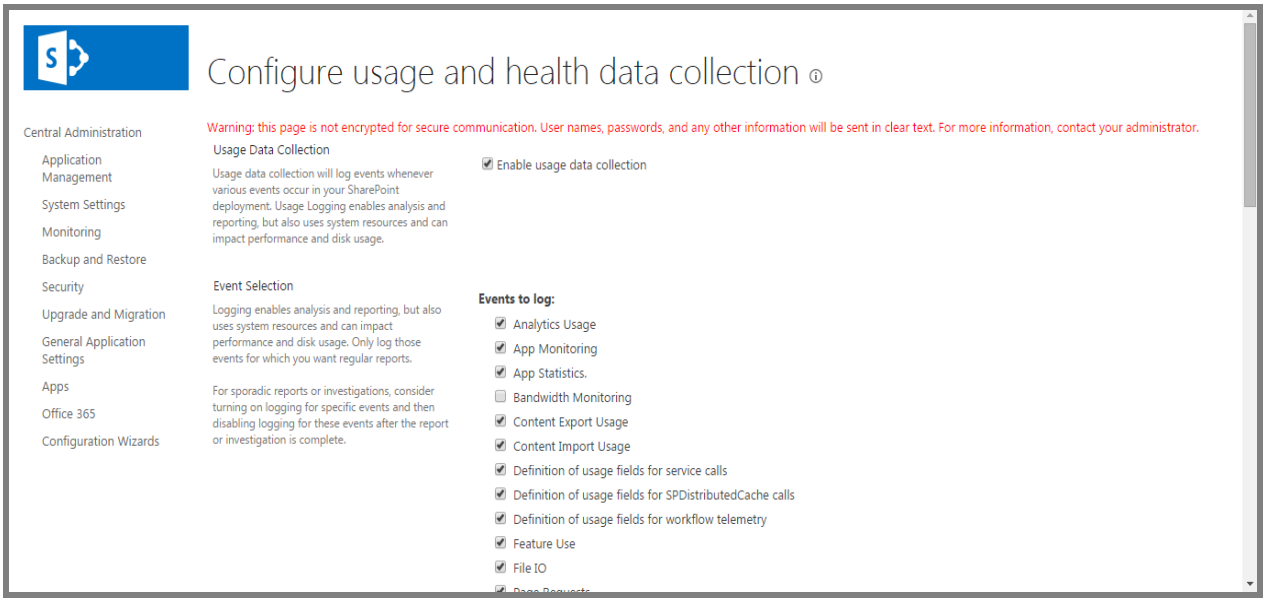

Figure 18.11: Scrolling down the 'Configure usage and health data collection' page

3. Scroll down Figure 18.12 and then select the **Enable health data collection** check box that becomes visible.

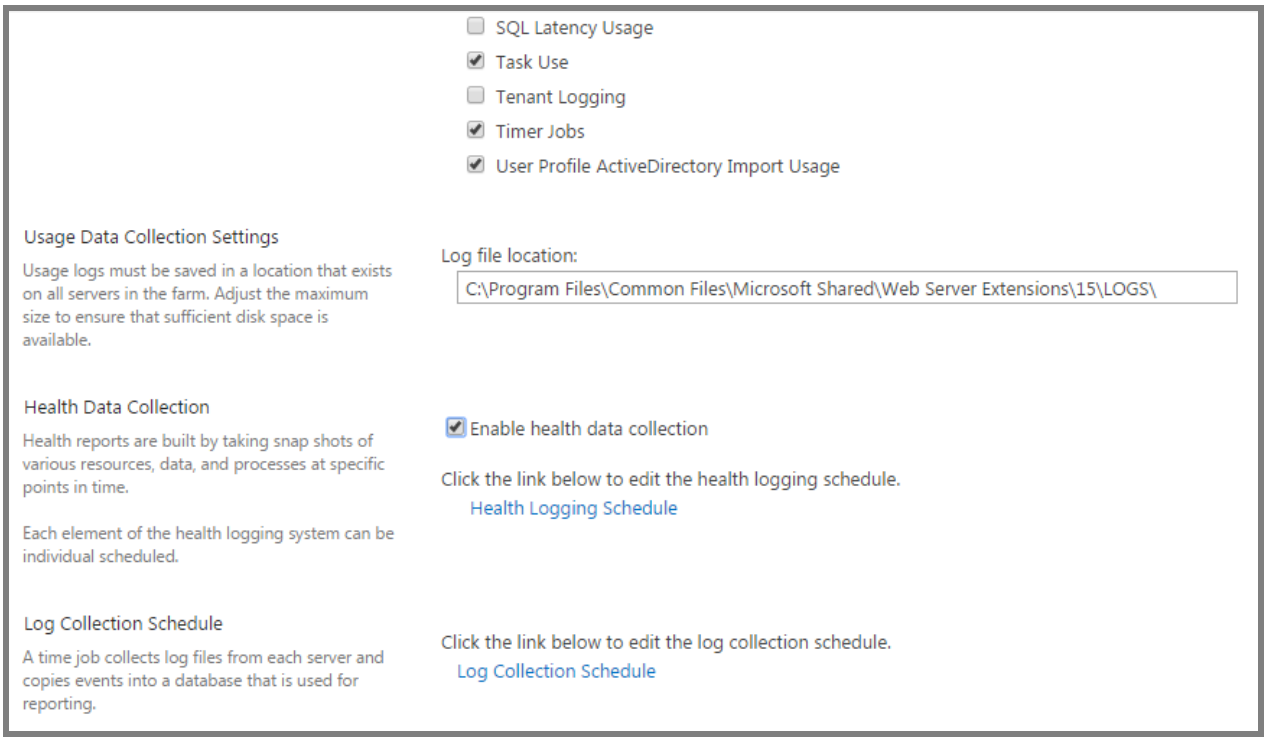

Figure 18.12: Enabling health data collection

18.2.2.3 Health Analyzer Alerts

SharePoint includes a tool named SharePoint Health Analyzer that enables you to diagnose and resolve configuration, performance, and usage problems. SharePoint Health Analyzer runs predefined health rules against servers in the farm. A health rule runs a test and returns an alert that tells you the outcome of the test. Every alert will indicate its severity – i.e., whether it is an *Error, Warning, Information, or a Rule execution failure*. Also, depending upon their nature, alerts are also automatically grouped into any of the standard categories, namely – *Security, Performance, Configuration, or Availability* – or can be part of any user-configured category. Every alert will also indicate the category to which it belongs.

The **Health Analyzer Alerts** test captures these alerts, ascertains their severity and category, and reports the count of alerts that were raised per severity for every category of alerts. You can also use the detailed diagnostics provided by this test to view the complete alert messages, the health rules that generated the alerts, and the server and services that were impacted by the problems for which the alerts were generated. This will not only lead you to the problem areas, but will also shed light on the probable problem cause, so that you can resolve the issue quickly.

Purpose	Captures Health Analyzer alerts, ascertains their severity and category, and reports the count of alerts that were raised per severity for every category of alerts
Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the HOST.
4. DOMAIN, DOMAIN USER, PASSWORD, and CONFIRM PASSWORD – This test requires the credentials of a farm administrator to run and collect metrics from the target SharePoint server. Therefore, specify the domain to which that farm administrator belongs in the DOMAIN text box, and then, enter the credentials of the farm administrator in the DOMAIN USER and PASSWORD text boxes. To confirm the password, retype it in the CONFIRM PASSWORD text box.
5. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each alert category</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Error messages:
Indicates the number of error alerts of this category that were currently generated by the Health Analyzer.</td><td>Number</td><td>Ideally, the value of this measure should be 0. If a non-zero value is reported, then use the detailed diagnosis of this measure to understand the errors that occurred and the servers and services impacted by the same.</td></tr>
<tr><td></td><td>Warning messages:
Indicates the number of warning alerts of this category that were currently generated by the Health Analyzer .</td><td>Number</td><td>Ideally, the value of this measure should be 0. If a non-zero value is reported, then use the detailed diagnosis of this measure to view the warning messages and the servers and services that may be impacted by the warnings.</td></tr>
<tr><td></td><td>Information messages:
Indicates the number of information alerts of this category that were currently generated by the Health Analyzer.</td><td>Number</td><td>Use the detailed diagnosis of this measure to view the information messages and the servers and services that they pertain to.</td></tr>
</table>

	Rule execution failure messages: Indicates the number of rule execution failure messages of this category that were currently generated by the Health Analyzer.	Number	Ideally, the value of this measure should be 0. If a non-zero value is reported, then use the detailed diagnosis of this measure to view the descriptive messages and to determine which servers and services failed because of rule execution failure.

Use the detailed diagnosis of the *Error messages* measure to understand what errors occurred, when, and which servers and services impacted by the same.

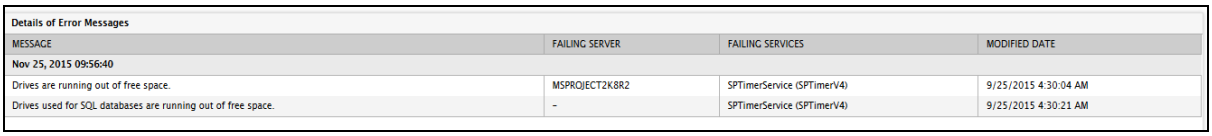

Details of Error Messages

MESSAGE	FAILING SERVER	FAILING SERVICES	MODIFIED DATE
Nov 25, 2015 09:56:40			
Drives are running out of free space.	MSPROJECT2K8R2	SPTimerService (SPTimerV4)	9/25/2015 4:30:04 AM
Drives used for SQL databases are running out of free space.	-	SPTimerService (SPTimerV4)	9/25/2015 4:30:21 AM

Figure 18.13: The detailed diagnosis of the Error messages measure

Use the detailed diagnosis of the *Warning messages* measure to view the warning messages and the servers and services that will potentially be impacted if the warnings are ignored.

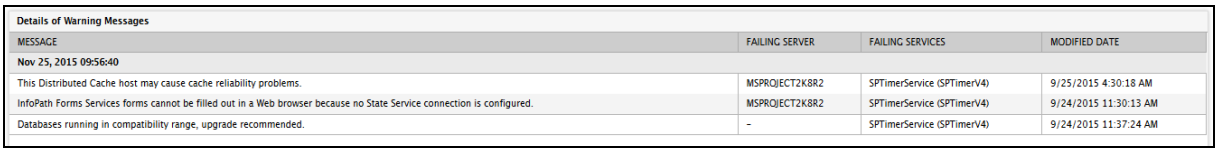

Details of Warning Messages

MESSAGE	FAILING SERVER	FAILING SERVICES	MODIFIED DATE
Nov 25, 2015 09:56:40			
This Distributed Cache host may cause cache reliability problems.	MSPROJECT2K8R2	SPTimerService (SPTimerV4)	9/25/2015 4:30:18 AM
InfoPath Forms Services forms cannot be filled out in a Web browser because no State Service connection is configured.	MSPROJECT2K8R2	SPTimerService (SPTimerV4)	9/24/2015 11:30:13 AM
Databases running in compatibility range, upgrade recommended.	-	SPTimerService (SPTimerV4)	9/24/2015 11:37:24 AM

Figure 18.14: The detailed diagnosis of the Warning messages measure

Use the detailed diagnosis of the *Information messages* measure to view the information messages and the servers and services they pertain to.

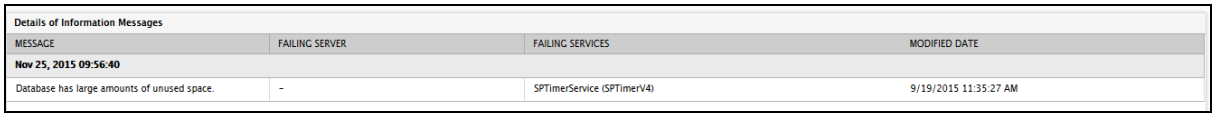

Details of Information Messages

MESSAGE	FAILING SERVER	FAILING SERVICES	MODIFIED DATE
Nov 25, 2015 09:56:40			
Database has large amounts of unused space.	-	SPTimerService (SPTimerV4)	9/19/2015 11:35:27 AM

Figure 18.15: The detailed diagnosis of the Information messages measure

Use the detailed diagnosis of the *Rule execution failure messages* measure to view the descriptive execution failure messages and to determine which servers and services failed because of rule execution failure.

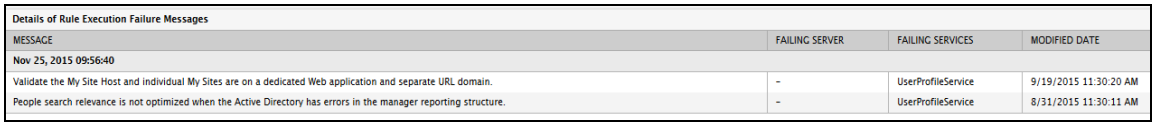

Details of Rule Execution Failure Messages

MESSAGE	FAILING SERVER	FAILING SERVICES	MODIFIED DATE
Nov 25, 2015 09:56:40			
Validate the My Site Host and individual My Sites are on a dedicated Web application and separate URL domain.	-	UserProfileService	9/19/2015 11:30:20 AM
People search relevance is not optimized when the Active Directory has errors in the manager reporting structure.	-	UserProfileService	8/31/2015 11:30:11 AM

Figure 18.16: The detailed diagnosis of the Rule execution failure messages measure

18.2.2.4 Backup and Restores Test

A backup is a copy of data that is used to restore and recover that data after a system failure. If a backup job fails, then all the data that could not be backed up cannot be recovered at the time of system failure, thus resulting in significant data loss. This is why, it is imperative that administrators be instantly alerted if any backup or restore job fails. This is exactly what the **Backup and Restores** test does! This test monitors each configured backup directory (local and/or remote), tracks the backups job and restores from every directory, and reports the count of backup and restore jobs that succeeded and/or failed on that directory. This way, the test notifies administrators as soon as a backup or restore job fails and also points them to the exact directory where the failure occurred. Detailed diagnostics provided by this measure also lead you to what exactly caused the backup or restore activity to fail, thereby enabling you to resolve isssues quickly and ensure smooth operations.

Purpose	Monitors each configured backup directory (local and/or remote), tracks the backups job and restores from every directory, and reports the count of backup and restore jobs that succeeded and/or failed on that directory.
Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **DOMAIN**, **DOMAIN USER**, **PASSWORD**, and **CONFIRM PASSWORD** – This test requires the credentials of a **farm administrator** to run and collect metrics from the target SharePoint server. Therefore, specify the domain to which that farm administrator belongs in the **DOMAIN** text box, and then, enter the credentials of the farm administrator in the **DOMAIN USER** and **PASSWORD** text boxes. To confirm the password, retype it in the **CONFIRM PASSWORD** text box. 5. **DIR PATH** – Provide the full path to the backup and/or restore directory to be monitored. Multiple paths can be provided as a comma-separated list. For example – *C:|BackupSite,C:|RestoreSite,D:|BackupDir*. Your specification can include both local and remote directories. For example - *C:|BackupSite,C:|RestoreSite,||192.168.9.70|backup|250216*. However, bear the following points in mind when including remote directory paths in your specifications: • While specifying the path of a remote directory, make sure that your specification begins with \\ (two forward slashes) followed by the IP/hostname of the remote server in which the directory resides. This should be followed by the full path of the remote directory to be monitored. For example - *||192.168.9.70|backup|250216*. • Your **DIR PATH** specification can include the path to multiple remote directories. Each of these directories can be in a different remote server. However, all these remote servers should operate in the same domain. • A single user in the remote domain should have access to all the remote directories configured against **DIR PATH**. 6. **REMOTE SERVER DOMAIN** – **This parameter is applicable only if the DIR PATH specification includes one/more remote directories.** In this case, against **REMOTE SERVER DOMAIN**, specify the domain to which the servers hosting the remote directories belong. If your **DIR PATH** specification does not include any remote directories, set **REMOTE SERVER DOMAIN** to *none*. **Note:** • Only a single domain name can be specified against **REMOTE SERVER DOMAIN**. • For proper results, all the servers that host the remote directories configured against **DIR PATH** should belong to the **REMOTE SERVER DOMAIN** you specify. 7. **REMOTE SERVER USER NAME** and **REMOTE SERVER PASSWORD** – **These parameters are applicable only if the DIR PATH specification includes one/more remote directories.** In such a case, against these parameters, specify the credentials of a user who fulfills the following conditions: • Should be a valid user in the **REMOTE SERVER DOMAIN** that you have configured;

<table>
<tr><td></td><td colspan="3">
Should be a user who has at least read-only access to all the remote directories configured for monitoring against the DIR PATH parameter.
However, if the DIR PATH specification does not include any remote directories, then you can set both REMOTE SERVER USER NAME and REMOTE SERVER PASSWORD to none.
8. CONFIRM PASSWORD – Confirm the REMOTE SERVER PASSWORD by retyping it here. This parameter again is applicable only if the DIR PATH specification includes one/more remote directories.
9. 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.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each DIR PATH configured for monitoring</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Failed backups:
Indicates the number of backups to this directory that failed currently.</td><td>Number</td><td>Ideally, the value of this measure should be 0.
If this measure reports a non-zero value, then use the detailed diagnosis of this measure to figure out which backup jobs failed, when the failure occurred, what error caused the failure, what backup method was employed, who initiated the backup, and many other details regarding the backup jobs that failed. Using these details, you can troubleshoot the failure easily.</td></tr>
<tr><td></td><td>Successful backups:
Indicates the number of backups to this directory that succeeded presently.</td><td>Number</td><td>Ideally, the value of this measure should be high.
Use the detailed diagnosis of this measure to know which backup jobs succeeded.</td></tr>
</table>

	Failed restores: Indicates the number of restores from this directory that failed currently.	Number	Ideally, the value of this measure should be 0. If this measure reports a non-zero value, then use the detailed diagnosis of this measure to figure out which restore jobs failed, when the failure occurred, what error caused the failure, what restore method was employed, who initiated the backup, and many other details regarding the failed restore jobs. Using these details, you can troubleshoot the failure easily.
	Successful restores: Indicates the number of restores from this directory that succeeded presently.	Number	Ideally, the value of this measure should be high. Use the detailed diagnosis of this measure to know which restore jobs succeeded.

Use the detailed diagnosis of the *Successful backups* measure to know which backup jobs succeeded.

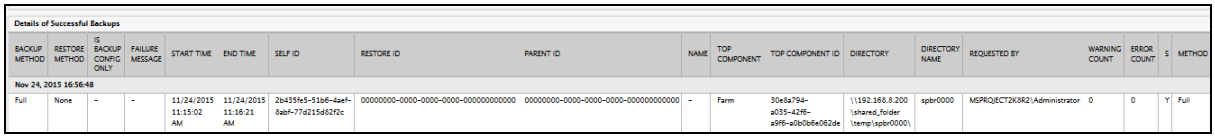

Figure 18.17: The detailed diagnosis of the Successful backups measure

18.2.3 Sharepoint Search Content Feed Layer

The key components of the Sharepoint content feeding chain are:

- Crawl Database
- Crawl Component
- Content Processing Component
- Index Component

When search queries execute slowly, administrators need to figure out where in the feeding chain the slowdown originated. The tests mapped to this layer run checks on all the aforesaid components, so that administrators can accurately isolate the probable cause of this slowdown.

18.2.3.1 Search Gatherer Threads Test

Search in SharePoint 2013 enables users to find relevant information more quickly and easily than ever before and makes it easy for Search administrators to customize the search experience.

The search architecture consists of the following areas:

- Crawl and content processing

- Index
- Query processing
- Search administration
- Analytics

Figure 18.18 depicts how these components work together to implement the search functionality in Sharepoint 2013.

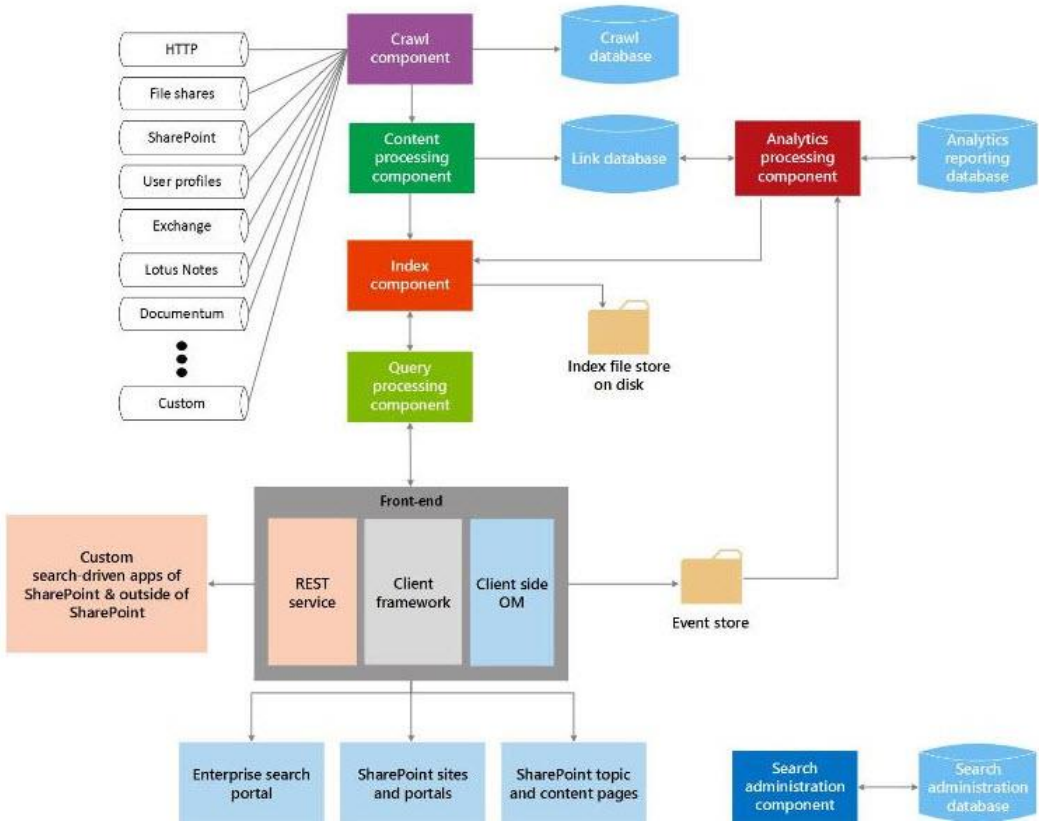

Figure 18.18: How search works in Sharepoint 2013?

From Figure 18.18, it is clear that the crawl component lays the foundation for the search mechanism! The crawl component crawls content sources to collect crawled properties and metadata from crawled items and sends this information to the content processing component. This means that if the crawl component is unable to crawl the content hosts, it could impact the speed of every dependent operation – be it content processing, indexing, query processing etc. – thereby crippling the entire search engine! Hence, for search in Sharepoint 2013 to be quick and efficient, administrators should primarily keep an eye on the crawl component, swiftly isolate painpoints in crawling, and clear them rapidly. To achieve this, administrators can use the **Search Gatherer Threads** test. This test monitors the crawling process and reveals how well the crawling worker threads are doing their jobs. While at it, the test proactively notifies administrators of a potential slowdown (if any) n crawling and pinpoints what is causing the slowdown – a hungry content host? or improperly configured crawls? .

Purpose	Monitors the crawling process and reveals how well the crawling worker threads are doing their

	jobs. While at it, the test proactively notifies administrators of a potential slowdown (if any) n crawling and pinpoints what is causing the slowdown – a hungry content host? or improperly configured crawls?		
Target of the test	A Sharepoint Server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Threads accessing the network: Indicates the number of threads that are waiting on the content host to return the requested content.	Number	If this number is consistently high then you are either network bound or you are bound by a "hungry" host. If you are not meeting your crawl freshness goals, you can either change your crawl schedules to minimize overlapping crawls or look at the remote repositories you are crawling to optimize them for more throughput.
	Filtering threads: Indicates the current number of filtering threads in the system.	Number	If the value of the *Threads accessing the network* measure is close to that of the *Filtering threads* measure, it is an indication that a bottleneck exists at the content source/host. When this happens, you may also want to check whether processor usage on the crawl component servers is low. Likewise, look for disk latency issues on the crawl database. If all the above exist, it is a clear indicator that the content host/source is where the bottleneck lies!
	Idle threads: Indicates the number of threads that are currently waiting for documents.	Number	These threads are not currently doing any work and will eventually be terminated. If you consistently have a more than *Max Threads/Hosts* idle threads you can schedule an additional crawl. If this number is 0 then you are starved. Do not schedule another crawl in this time period and analyze the durations of your crawls during this time to see if they are meeting your freshness goals. If your goals are not being met you should reduce the number of crawls.

18.2.3.2 Search Gatherer Transactions Test

Crawls, when scheduled to occur too frequently, can significantly impact the processing ability of the content processing component, the level of I/O activity on the crawl database, and ultimately, the search throughput! Likewise, a resource-starved content processing component and/or a crawl database can also considerably slowdown Sharepoint search, as they may not be able to handle the workload generated by the crawler! This is why, when end-users complain of slow searching by Sharepoint, administrators need to be able to quickly figure out where the bottleneck is and how to clear it – should the crawl schedules be changed so that less crawls occur? Or should the processing power of the content processor and crawl database change in tandem with the frequency of crawls? This is where the **Search Gatherer Transactions** test helps!

This test monitors the transactions on the crawl component and reports the count of transactions that are waiting for processing by the content processor and those that have completed processing. In the process, the test turns the spotlight on a potential processing slowdown and accurately pinpoints what is causing it – is it owing to too many crawls? Or is it because the content processor and/or the crawl database are incorrectly sized? Based on the results of this test, administrators can clearly understand what needs to be fine-tuned and how.

Purpose	Monitors the transactions on the crawl component and reports the count of transactions that are waiting for processing by the content processor and those that have completed processing. In the process, the test turns the spotlight on a potential processing slowdown and accurately pinpoints what is causing it – is it owing to too many crawls? Or is it because the content processor and/or the crawl database are incorrectly sized? Based on the results of this test, administrators can clearly understand what needs to be fine-tuned and how		
Target of the test	A Sharepoint Server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Waiting transactions: Indicates the number of transactions that are currently waiting to be processed by the content processing component.	Number	Ideally, this value should be low (less than a few thousand). If so, it implies that content processing is keeping up with content crawling. On the other hand, if the value of this measure is high and/or consistently rising, then it means that the crawl component is pushing more data for processing than what the content processing component can handle. This will slow down content processing and eventually affect Sharepoint search! Under such circumstances, you can do either of the following: • Provide more processing power to the content processing component, so that it is able to handle the load imposed by the crawl component. You can also add more content processing components to uniformly distribute the processing load. • Reconfigure the crawl component to run crawls less frequently, so that the crawl component does not overload the content processing component
	Transactions in progress: Indicates the number of transactions that are currently being processed by the crawl component.	Number	This is a good indicator of the current load on the crawl component.
	Completed transactions: Indicates the number of transactions that are completed	Number	If this value is very high (say, greater than a few hundred), it means that too many transactions are getting completed and are written to the crawl database, causing disk activity on the database to increase. At this juncture, check the crawl database for disk latency. If the disk latency and disk queue length are also high, you can conclude that the crawl database is where the bottleneck is.

18.2.3.3 Search Submission Test

Like problems in the content acquisition process, snags in the content processing routine can also delay searching. Content processing in Sharepoint is performed by the content processing component (CPP) and the index component. Once crawling is complete, the Content plug-in on the crawl component first routes the content to the **Content Submission Service** (CSS) of the content processing component. An instance of the CSS runs alongside each instance of a content processing component. Once the content plug-in on the crawl component establishes a session with the CSS, the CSS load-balances the incoming content by uniformly distributing the content to the content processing components (CPC). Upon receipt of documents from the CSS, the content processing component processes the documents and then sends them to the indexer for indexing.

If a crawler session is unexpectedly terminated by CSS, then some crawled content may not even reach the CSS, and will hence not be processed or indexed; this will eventually impact the search service! Moreover, if CSS is not able to push its document load to the content processing component fast enough, documents may get timed out from the CSS itself, and will hence be omitted from the search index; this again will result in a poor search experience. Likewise, if the content processing component suffers a slowdown, document processing and indexing will be significantly delayed, which in turn can affect querying. If such problems are to be avoided, administrators should closely monitor the availability and processing ability of the CSS and the CPC, and rapidly isolate bottlenecks. This is where the **Search Submission** test helps. This test periodically checks the sessions to CSS, monitors how quickly the CSS load-balances the content and transmits it to the CPC, and measures the processing capacity of the CPC. When users complain of their search queries being slow, then this test will shed light on the probable cause of the delay – is it owing to sudden/sporadic breaks in the crawler sessions to CSS? Is it because of a load-balancing bottleneck experienced by the CSS? Or is it due to a processing slowdown at the CPC? Based on the findings reported by this test, administrators can initiate the appropriate remedial measures.

Purpose	Periodically checks the sessions to CSS, monitors how quickly the CSS load-balances the content and transmits it to the CPC, and measures the processing capacity of the CPC		
Target of the test	A Sharepoint Server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Aborted sessions: Indicates the number of sessions that aborted since the start of the component.	Number	Ideally, the value of this measure should be 0. A high value is a cause for concern as it indicates frequent breaks in the crawler sessions on the CSS. Too many broken sessions can seriously impede the transfer of crawled content from the crawler to the CSS, resulting in incomplete transfers! This warrants an investigation into the reason for the frequent session failures.
	Active sessions: Indicates the number of crawler sessions that are currently active on the CSS.	Number	This is a good indicator of the current load on the CSS.
	Available calbacks: Indicates the current number of callbacks ready for consumption, but not yet consumed by the client.	Number	Once the content processing component processes the content it receives and writes it to the index, it sends out a 'call back' to the content plug-in on the crawler indicating the processing status of that content. A high value for this measure indicates that while the CPC has been able to generate callbacks, many of these callbacks have not yet been consumed by – i.e., have not yet reached – the crawler. This hints at an error in network communication between the crawler and the CPC.
	Total callbacks: Indicates the total number of callbacks produced by the submission service since the start of the component.	Number	You may want to compare the value of the *Available callbacks* measure with that of this measure to understand what fraction of callbacks is still to be consumed by the crawl component.
	Client polls: Indicates the total number of client polls since the start of the component.	Number	Each time a client refreshes the session to check for callbacks this measure will be incremented.
	Client submits: Indicates the total number of submits performed by clients since the start of the component.	Number	

	Skipped documents: Indicates the total number of documents skipped in the submission service before being delivered to the content processing component.	Number	A non-zero value is desired for this measure. A high value is disconcerting as it indicates that too many crawled documents are not reaching the CPC for processing as the CSS disregards them. Further investigation into the reasons is necessitated.
	Timed out documents: Indicates the total number of documents that timed out in the submission service.	Number	A low value is desired for this measure. A high value implies that the search index may not include many crawled documents as they have been timed out of the submission queue itself. This in turn may result in ineffective search queries. You may hence want to reset the timeout value for documents in the submission service.
	Flows used for feeding: Indicates the current number of flows used for feeding.	Number	The CPC uses Flows and Operators to process the content. Flows define how to process content, queries and results and each flow processes one item at a time. The number of current flows is hence an indicator of the number of documents that are being processed by the CPC.
	Pending items: Indicates the current number of items delivered to the content processing component but where no callback has yet been received.	Number	A high value or a consistent rise in the value for this measure could indicate a bottleneck in content processing.

18.2.3.4 Search Flow Test

Content processing in Sharepoint is performed by the content processing component (CPC) and the index component.

The Content Processing Component (CPC) uses Flows and Operators to process the content (see Figure 18.19). Flows define how to process content, queries and results and each flow processes one item at a time. Flows consist of operators and connections organized as graphs. This is where activities like language detection, word breaking, security descriptors, content enrichment (web service callout), entity and metadata extraction, deep link extraction and many others take place. The flow has branches that handle different operations, like inserts, deletes and partial updates.

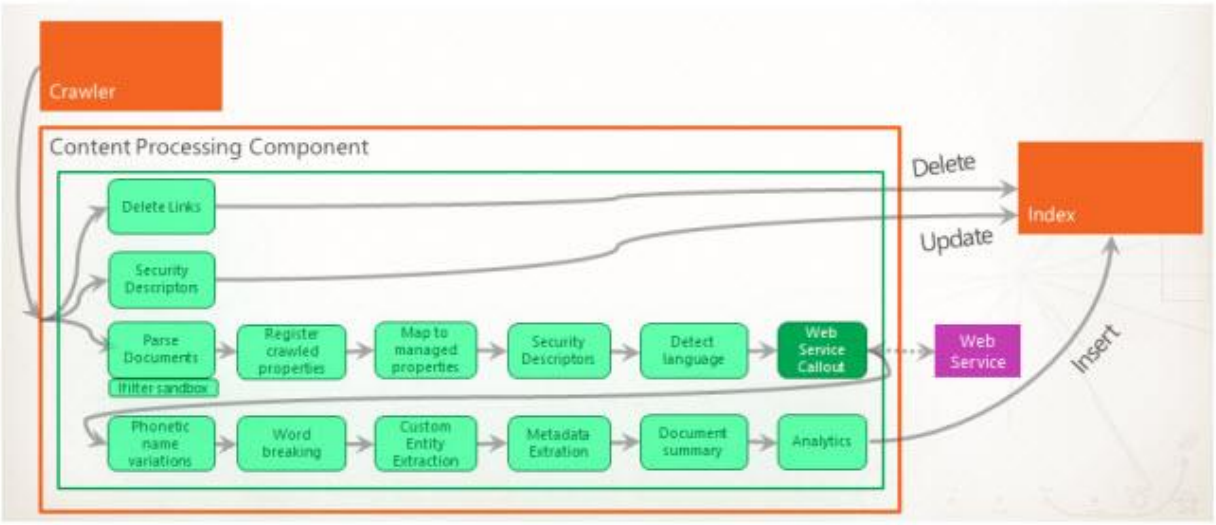

Figure 18.19: Flows and operators in CPC

Once content is processed by the CPC, the index component receives the processed items from the CPC and writes them to the search index. The index component also handles incoming queries, retrieves information from the search index, and sends back the result set to the query processing component.

Whether it is the CPC that fails to process the content rapidly or the index component that writes to the index slowly, what suffers is the end-user's experience with Sharepoint search! To ensure that Sharepoint delivers to users a fast and flawless searching experience, administrators should not only be able to detect slowdowns before they impact query processing, but also tell where the slowdown originated – is it with the CPC or the index component? The **Search Flows** test answers this question accurately! This test monitors the flows on CPC, keeps track of documents that are in queue waiting to be processed by the flows, and reports how quickly the CPC and the index component process the enqueued contents. While at it, the test points to potential bottlenecks in content processing and accurately isolates the source of the bottleneck – is it the CPC or the index component?

Purpose	Monitors the flows on CPC, keeps track of documents that are in queue waiting to be processed by the flows, and reports how quickly the CPC and the index component process the enqueued contents. While at it, the test points to potential bottlenecks in content processing and accurately isolates the source of the bottleneck – is it the CPC or the index component?		
Target of the test	A Sharepoint Server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the Sharepoint server monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Total inbound items: Indicates the total number of items placed on input queues.	Number	
	Items queued for processing: Indicates the number of items that are currently in queues in front of input operators that are ready for processing.	Number	A high value or a consistent increase in the value of this measure is indicative of bottlenecks in content processing.
	Active threads: Indicates the number of threads that are currently active.	Number	
	Input queue empty time: Indicates the total time spent by input operators waiting for items.	Millisecs	If this value is low (say, less than a thousand), it indicates that the input queues are rarely ever empty! You may then want to check the processor usage on the CPC component. If this is very high, it is a clear indication that the CPC is stressed and could be the key contributor to the slowdown in content processing. On the other hand, if the value of this measure is high (say, over a thousand) , it indicates that the input queues are empty for long time spells. This implies that the CPC is processing content quickly. In this case, check the disk I/O and latency on the index component. If these parameters are high, it implies that the index component is stressed and is unable to handle the load imposed by the CPC. You can thus conclude that the bottleneck lies with the index component.

	Input queue full time: Indicates the total time spent waiting for space to become available on input queues.	Millisecs	If this value is high (say, over a thousand), it indicates that the CPC is taking a long time to process the contents in the input queues and free up the queues! You may then want to check the processor usage on the CPC component. If this is very high, it is a clear indication that the CPC is stressed and could be the key contributor to the slowdown in content processing. On the other hand, if the value of this measure is low (say, less than a thousand), it indicates that the input queues are getting cleared very quickly. This implies that the CPC is processing content quickly. In this case, check the disk I/O and latency on the index component. If these parameters are high, it implies that the index component is stressed and is unable to handle the load imposed by the CPC. You can thus conclude that the bottleneck lies with the index component.

18.2.4 The Sharepoint Documents Layer

Using the tests mapped to this layer, you can closely monitor the growth in the number and size of document libraries, documents, and lists.

Figure 18.20: The tests mapped to the Sharepoint Documents Layer

18.2.4.1 Sharepoint Document Libraries and Lists Test

Document libraries are collections of files that you can share with team members on a Web based on Microsoft Windows SharePoint Services. For example, you can create a library of common documents for a project, and team members can use their Web browsers to find the files, read them, and make comments. Users with Microsoft Office 2003 can check out and edit the files as if they resided on a local or network drive.

A list in SharePoint is used to store data across columns in separate rows. You can think of a list as a table in a database that will have columns and rows. You can think of a list as a table in a database that will have columns and rows. You can also think of it as a spreadsheet with columns and rows. Items such as issues, software bugs, employee addresses, phone numbers, web site links or pretty much anything else can be stored.

To ensure that all the web applications deployed on the Sharepoint farm have adequate storage resources at their disposal, administrators must make sure that document libraries and lists used by the web applications do not grow

uncontrollably, both in number and in size. For this, administrators must keep a close watch on the growth of the document libraries and lists. This is where the **Sharepoint Document Libraries and Lists** test helps! This test reports the total number of document libraries and lists created on Sharepoint, tracks the rate at which these numbers are growing, and promptly alerts administrators to an abnormal increase in the number of document libraries and lists. In addition, the test also measures the size of document libraries from time to time, and intimates administrators if the size increases unexpectedly! The detailed diagnosis of this test also reports the top-10 document libraries and lists in terms of size, thus leading administrators to those libraries and lists that could be draining the storage resources of Sharepoint.

Purpose	Reports the total number of document libraries and lists created on Sharepoint, tracks the rate at which these numbers are growing, and promptly alerts administrators to an abnormal increase in the number of document libraries and lists. In addition, the test also measures the size of document libraries from time to time, and intimates administrators if the size increases unexpectedly! The detailed diagnosis of this test also reports the top-10 document libraries and lists in terms of size, thus leading administrators to those libraries and lists that could be draining the storage resources of Sharepoint.		
Target of the test	A Sharepoint Server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **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**. 5. **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.		
Outputs of the test	One set of results for the Sharepoint Server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Number of document libraries: Indicates the total number of document libraries on the Sharepoint server.	Number	A consistent increase in the value of this measure could indicate that new document libraries are regularly created on Sharepoint. You may want to check how much space these new libraries are consuming to understand the true impact of this addition on storage resources. You can use the detailed diagnosis of this measure to identify the top-10 document libraries in terms of size – i.e., space usage.
	Documents in document libraries: Indicates the total number of documents in all document libraries on Sharepoint.	Number	A consistent increase in the value of this measure could indicate the influx of new documents into existing document libraries or the creation of new libraries with a new set of documents. You may want to check how much space these new documents are consuming to understand the true impact of this addition on storage resources.
	Size of document libraries: Indicates the total size of all the document libraries on Sharepoint.	MB	A consistent increase in the value of this measure could be attributed to the addition of new document libraries, new documents, and large-sized documents.
	Average number of documents per document library: Indicates the average number of documents per library.	Number	
	Document library growth rate: Indicates the percentage growth in the number of document libraries handled by Sharepoint, since the last measurement period.	Percent	A consistent increase in the value of this measure could indicate that new document libraries are regularly created on Sharepoint. You may want to check how much space these new libraries are consuming to understand the true impact of this addition on storage resources.
	Lists count: Indicates the number of lists on Sharepoint.	Number	A consistent increase in the value of these measures could indicate that new lists are regularly created on Sharepoint. You may want to check how much space these new lists are consuming to understand the true impact of this addition on storage resources. You can use the detailed diagnosis of the *Lists* measure to identify the top-10 Sharepoint lists in terms of size – i.e., space usage.
	Lists growth rate: Indicates the percentage growth in the number if lists on Sharepoint, since the last measurement period.	Number	

	Attachments: Indicates the number of attachments on Sharepoint.	Number	

The detailed diagnosis of the *Number of document libraries* measure lists the top 10 libraries in Sharepoint with the maximum number of documents. Using this information, you can quickly identify that document library with the highest document count and also figure out the **PARENTWEBURL** of the web application with which the library is associated. If that web application grows abnormally in size or count of documents, this information will lead administrators to the exact document library that is responsible for it.

List of Top 10 Document Library

TIME	TITLE	DESCRIPTION	ITEMCOUNT	PARENTWEBURL
Jan 30, 2014 06:42:11				
	Documents	-	8	/
	Documents	-	4	/
	Documents	-	3	/sites/mysites
	Documents	-	2	/sites/testcomplete
	Documents	-	1	/sites/eginnovations
	Documents	-	1	/site
	Documents	-	0	/sites/quota
	Documents	-	0	/sites/new_site_pravat
	Documents	-	0	/sites/test

Figure 18.21: The detailed diagnosis of the Number of document libraries measure

The detailed diagnosis of the *Lists count* measure displays the top 10 lists in Sharepoint with the maximum number of items. Using this information, you can quickly identify that list with is most heavily populated and also figure out the **PARENTWEBURL** of the web application with which the list is associated. If that web application grows abnormally, this information will lead administrators to the exact list that may be responsible for it.

List of Top 10 Lists

TIME	TITLE	DESCRIPTION	ITEMCOUNT	PARENTWEBURL
Jan 30, 2014 06:42:11				
	Composed Looks	Use this list to store composed looks. These looks can be applied to this site by navigating to Site Settings and choosing Change the look.	18	/my

Figure 18.22: The detailed diagnosis of the Lists count measure

18.2.4.2 Sharepoint Documents Test

Documents are stored within a document library in Sharepoint. Documents add to the size of the sites, site collections, and web applications they are associated with. Significant and rapid spikes in the number and size of documents on the Sharepoint server can hence cause sites, site collections, and ultimately, web applications to grow in size exponentially; in the long run, this may result in a severe space crunch in the content database. This is why, administrators need to keep a close watch on the number of documents handled by the Sharepoint server and the space resources they use. To achieve this, administrators can use the **Sharepoint Documents** test! This test periodically monitors the number and size of documents in the Sharepoint server, reports abnormal document growth, and thus warns administrators of potential space contentions well before they actually occur!

Purpose	Periodically monitors the number and size of documents in the Sharepoint server, reports abnormal document growth, and thus warns administrators of potential space contentions well before they actually occur		
Target of the test	A Sharepoint Server 2010/2013		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **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**. 5. **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.		
Outputs of the test	One set of results for the Sharepoint Server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

	Number of documents in Sharepoint: Indicates the total number of documents in the Sharepoint server.	Number	A consistent increase in the value of this measure could indicate that new documents are created in Sharepoint at regular intervals. You may want to check how much space these new documents are consuming to understand the true impact of this addition on storage resources.
	Versions: Indicates the total number of document versions in Sharepoint.	Number	A consistent increase in the value of this measure could indicate that newer versions of one/more existing documents are now available in Sharepoint. This in turn implies that many outdated/obsolete documents may also exist in Sharepoint. In the event of rapid growth in document count, you may want to delete the stale versions of documents so as to control the growth and make space for newer documents.
	Size of all documents: Indicates the total size of all the documents n Sharepoint.	MB	A consistent increase in the value of this measure could be attributed to the addition of new documents and/or large-sized documents.
	Average size of a document: Indicates the average size of a document.	MB	With the help of the value of this measure, you can ascertain whether/not Sharepoint is the container for documents of large sizes.
	Documents growth rate: Indicates the percentage growth in the number of documents in Sharepoint, since the last measurement period.	Percent	A consistent increase in the value of this measure could indicate there is a consistent addition of new documents to Sharepoint. Compare the value of this measure with that of the *Versions* measure to understand whether the addition of newer 'versions' of existing documents is in any way contributing to the growth rate. If so, you may want to delete older versions of documents and unnecessary documents to curb the growth.
	Number of file formats stored: Indicates the total number of file formats stored in Sharepoint.	Number	Use the detailed diagnosis of this measure to know which file formats are stored in Sharepoint.

18.2.5 The Sharepoint Objects Layer

The tests mapped to this layer promptly capture the sporadic spikes or steady growth in the contents of the critical Sharepoint data containers such as content databases, sites and site collections, and web applications. Overgrown applications and objects responsible for the uncontrollable growth can thus be isolated.

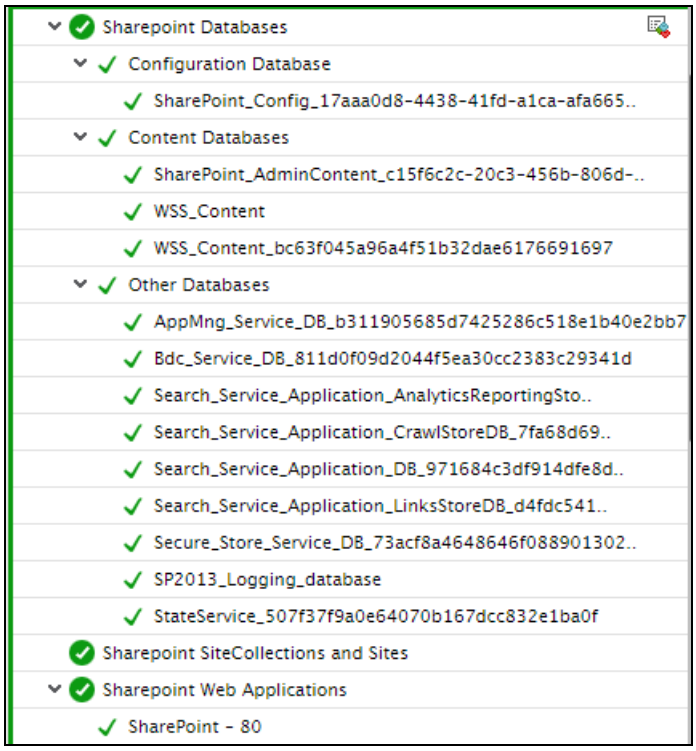

Figure 18.23: The tests mapped to the Sharepoint Objects layer

18.2.5.1 Sharepoint Databases Test

Different types of databases are typically installed for SharePoint. These database types are briefly discussed below:

- **Configuration Database:**

  The configuration database contains data about the following:

  - SharePoint databases
  - Internet Information Services (IIS) web sites
  - Web applications
  - Trusted solutions
  - Web Part packages
  - Site templates
  - Web applications

  The configuration database also contains specific data for SharePoint 2013 farm settings, such as default quota settings and blocked file types.

- **Content Database:**

  Content databases store all content for a site collection. This includes site documents or files in document libraries, list data, Web Part properties, audit logs, and sandboxed solutions, in addition to user names and rights.

  All of the files that are stored for a specific site collection are located in one content database on only one server. A content database can be associated with more than one site collection.

  Content databases also store user data for Power Pivot for SharePoint, if you installed it in your SharePoint Server 2013 environment.

- **Other Databases:**

  Service Application databases such as App Management database, Business Data Connectivity database, Search service application database, Secure store service database Usage and Health Data Collection database, and many more are typically grouped under **Other Databases**.

These databases can grow pretty quickly, and if this growth is not tracked and controlled, users may be left with no space for Sharepoint data. Sharepoint administrators should hence prudently and proactively plan their data storage needs, accordingly size the databases, and effectively manage the space available in the databases, so that manageability, performance, and reliability issues do not arise. This is where the **Sharepoint Databases** test helps!

Besides reporting the state of each database, this test also monitors the size, usage, and growth of every database, thus pointing administrators to those databases that are over-used or are exhibiting alarming growth patterns! In addition, the test provides hints for enhancing the overall performance of the content databases – will it help to cleanup the orphaned items? should the recycle bin storage space be reduced? should the content database host fewer site collections?

Purpose	Besides reporting the state of each database, this test also monitors the size, usage, and growth of every database, thus pointing administrators to those databases that are over-used or are exhibiting alarming growth patterns! In addition, the test provides hints for enhancing the overall performance of the content databases – will it help to cleanup the orphaned items? should the recycle bin storage space be reduced? should the content database host fewer site collections?
Target of the test	A Sharepoint Server 2010/2013
Agent deploying the test	An internal agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.
Outputs of the test	One set of results for each database of each type used by the Sharepoint Server being monitored First-level descriptor: Database type Second-level descriptor: Database name

<table>
<tr><th>Measurements made by the test</th><th>Measurement</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td>Is database in use?:
Indicates whether/not this database is in use.</td><td></td><td>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>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:
By default, the measure reports the Measure Values listed in the table above to indicate the usage state of the database. However, in the graph of this measure, the same will be represented using the numeric equivalents only.</td></tr>
<tr><td></td><td>Size:
Indicates the current size of this database.</td><td>GB</td><td>The size requirements typically vary with database type.

For instance, Configuration databases are unlikely to grow significantly, but Content databases are prone to rapid growth. Hence, Microsoft recommends that no content database be more than 200 GB in size.

Please refer to the following link: https://technet.microsoft.com/en-IN/library/cc678868.aspx, for detailed information on the size and scaling considerations of the different types of SharePoint databases.</td></tr>
</table>

	Disk space usage: Indicates the percentage disk space in the SQL server that is used by this database.	Percent	A high value for this measure is a cause for concern, as it indicates excessive disk space consumption by a database. Compare the value of this measure across databases of a type to identify that database which is eroding the disk space of the SQL server.
	Database growth rate: Indicates the percentage growth in the size of this database since the last measurement period.	Percent	A consistent rise in the value of this measure is a sign that the database is growing rapidly! Such rapid growth trends can be noticed more often in content databases. Since Microsoft recommends that no content database should be more than 200 GB in size, measures should be taken to control the growth of a content database. In this regard, you may want to consider the following measures: • Use an ootb Record Center as an archive for old content: The users must manually send each document to the RC using e.g. move and leave a link; note that only the latest major version with metadata is kept – all version history is lost. The information management policies supported by SharePoint for retention and disposition can be used to automate the cleanup. As the RC has its own content databases, the live collaboration databases will grow slower or even shrink as outdated information is moved to the archive. Keeping the live databases small ensures shorter recovery time; while the recovery time for the archived content can be considerable, but not business critical. Search must be configured appropriately to cover both live and archived content. • Use a third-party archiving solution for SharePoint. This has the same pros & cons as the previous option, but the functionality is probably better in relation to keeping version history and batch management of outdated content. Search must be configured appropriately to cover both live and archived content.

			• Use a third-party remote blob storage (RBS) solution for SharePoint so that documents are registered in the database, but not stored there. This gives smaller content databases, but more complicated backup and recovery as the content now resides both in databases and on disk. Provided that you don't lose both at the same time, the recovery time should be shorter. Search will work as before, as all content is still logically in the "database". • The databases size will shrink as data is actually deleted, and backup and recovery is more complicated as content is now both in the database and on disk. Search can be configured to also crawl and index the files on disk, but content ranking will suffer as the valuable metadata is lost. • Use powershell scripts or other code to implement the disposition of outdated content. The script can e.g. copy old documents to disk and delete old versions from the content database; the drawback being that all metadata will be lost and there is no link left in SharePoint. The databases size will shrink as data is actually deleted, and backup and recovery is more complicated as content is now both in the database and on disk. Search can be configured to also crawl and index the files on disk, but content ranking will suffer as the valuable metadata is lost.

	Total orphaned items: Indicates the number of orphaned sites in this content database.	Number	**This measure is reported only for 'Content Databases'.** An Orphaned Site is where SharePoint only has partial information and not a complete set of data for a given site collection in your Windows SharePoint Services or SharePoint Portal Server content databases or configuration databases. The site may in fact still be viewable via the browser, but you may notice that many things are broken. If the *Content database growth rate* measure is increasing consistently, you may want to check the variations in the value of this measure over the same time period to figure out whether/not the existence of too many orphan sites is contributing to the growth in the size of the content database. If so, you may want to cleanup the orphan sites to right-size your database and to ensure optimum performance.

	Site limit: Indicates the maximum number of site collections that this content database can host.	Number	**This measure is reported only for 'Content Databases'.** Microsoft strongly recommends limiting the number of site collections in a content database to 5,000. However, up to 10,000 site collections in a database are supported. Note that in a content database with up to 10,000 total site collections, a maximum of 2,500 of these can be non-Personal site collections. It is possible to support 10,000 Personal site collections if they are the only site collections within the content database. These limits relate to speed of upgrade. The larger the number of site collections in a database, the slower the upgrade with respect to both database upgrade and site collection upgrades. The limit on the number of site collections in a database is subordinate to the limit on the size of a content database that has more than one site collection. Therefore, as the number of site collections in a database increases, the average size of the site collections it contains must decrease. Exceeding the 5,000 site collection limit puts you at risk of longer downtimes during upgrades. If you plan to exceed 5,000 site collections, Microsoft recommends that you have a clear upgrade strategy to address outage length and operations impact, and obtain additional hardware to speed up the software updates and upgrades that affect databases.
	Configured site limit usage: Indicates the percentage of the configured site limit that is used by the content database.	Percent	**This measure is reported only for 'Content Databases'.** A value close to 100% indicates that the configured site limit is about to be reached. By comparing the value of this measure across content databases, you can easily identify the database that hosts too many site collections. You may then have to reassess the ability of that content database to handle additional site collections, and accordingly decide whether to reconfigure the site limit or reduce the number of site collections hosted by the database.

<table>
<tr>
<td></td>
<td>Needs upgrade?

Indicates whether/not this database needs to be upgraded.</td>
<td></td>
<td>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>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

This measure reports the Measure Values listed in the table above to indicate whether/not a database needs an upgrade. In the graph of this measure however, the same is represented using the numeric equivalents only.</td>
</tr>
<tr>
<td></td>
<td>Recycle bin storage space:

Indicates the space used by the items present in the second stage recycle bin of this database.</td>
<td>MB</td>
<td>Recycle Bins are used to help users protect and recover data. Microsoft SharePoint Server supports two stages of Recycle Bins: the first-stage Recycle Bin and second-stage Recycle Bin.

When a user deletes an item, the item is automatically sent to the first-stage Recycle Bin. By default, when an item is deleted from the first-stage Recycle Bin, the item is sent to the second-stage Recycle Bin.

A high value for this measure could indicate that a large amount of deleted data resides in the second stage recycle bin, unnecessarily consuming disk space and increasing the size of the database.</td>
</tr>
<tr>
<td></td>
<td>Recycle bin storage space growth rate:

Indicates the percentage growth in the space used in the second stage recycle bin of this database, since the last measurement period.</td>
<td>Percent</td>
<td>A consistent increase in the value of this measure indicates that deleted data is steadily accumulating in the recycle bin; this is a cause of concern, as data in the second stage recycle bin can add megabytes to the overall size of the database!

In case of content databases, every site collection has a second stage recycle bin and the size of this bin must not grow beyond 50 percent of the quota set for that site collection. You may want to reduce this percentage to ensure that the recycle bin does not grow too unwieldy and impact the size and performance of the content database.</td>
</tr>
</table>

18.2.5.2 SharePoint Farm Test

A SharePoint farm is a collection of SharePoint servers or SQL servers that work in concert to provide a set of basic SharePoint services that support a single site.

Since the primary purpose of any farm is to provide high availability to servers and services, administrators should be proactively alerted if that farm goes down. If not, end-users will be denied access to all servers and services riding on that farm for long periods of time!

Administrators will also require deep visibility into what servers and services make up the farm and what their current status is. Without this, unavailable servers/services can neither be identified, nor restored!

Moreover, until serious performance issues surface, administrators tend to remain clueless about which servers in a farm are updated/upgraded with critical patches/hot fixes and which are not. To avoid this, administrators should continuously track the upgrade status of the farm and of the servers in the farm.

The **SharePoint Farm** test addresses all these requirements! The test auto-discovers the SharePoint farm in which the monitored SharePoint server resides. Periodically, the test checks the status of this farm and alerts you if the farm goes offline. This way, you can initiate timely measures for restoring the farm to normalcy and in the process, ensure that users are able to access servers and services continuously. The test also reports the number, types, and names of servers in this farm and points you to the offline servers. Disabled service instances in the farm are also brought to light by this test. Additionally, the test also draws your attention to servers in a farm that need to be upgraded, so that you can quickly apply the required patches/hot fixes on those servers and maximize their performance.

Purpose	Auto-discovers the SharePoint farm in which the monitored SharePoint server resides. Periodically, the test checks the status of this farm and alerts you if the farm goes offline. This way, you can initiate timely measures for restoring the farm to normalcy and in the process, ensure that users are able to access servers and services continuously. The test also reports the number, types, and names of servers in this farm and points you to the offline servers. Disabled service instances in the farm are also brought to light by this test. Additionally, the test also draws your attention to servers in a farm that need to be upgraded, so that you can quickly apply the required patches/hot fixes on those servers and maximize their performance.		
Target of the test	A Sharepoint Server that is part of a Sharepoint farm		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **DOMAIN**, **DOMAIN USER**, **PASSWORD**, and **CONFIRM PASSWORD** – This test requires the credentials of a domain user to run and collect farm-related metrics from the target SharePoint server. Therefore, specify the domain to which that user belongs in the **DOMAIN** text box, and then, enter the credentials of the user in the **DOMAIN USER** and **PASSWORD** text boxes. To confirm the password, retype it in the **CONFIRM PASSWORD** text box. 5. **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.		
Outputs of the test	One set of results for the Sharepoint farm in which the monitored SharePoint server operates		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**

<table>
<tr>
<td></td>
<td>Status:
Indicates the current status of this farm.</td>
<td></td>
<td>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>Offline</td><td>1</td></tr><tr><td>Online</td><td>0</td></tr></table>
Note:
This measure reports the Measure Values listed in the table above to indicate the status of a farm. In the graph of this measure however, farm status is represented using the numeric equivalents only.</td>
</tr>
<tr>
<td></td>
<td>Needs upgrade?:
Indicates whether/not this farm needs an upgrade.</td>
<td></td>
<td>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>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:
This measure reports the Measure Values listed in the table above to indicate whether/not a farm needs an upgrade. In the graph of this measure however, the same is represented using the numeric equivalents only.
If this measure reports the value No, it could be because the patches were applied on the servers in the farm, but the SharePoint Products Configuration Wizard was not run after patch application on a few servers.</td>
</tr>
<tr>
<td></td>
<td>Total servers in farm:
Indicates the total number of servers in this farm.</td>
<td>Number</td>
<td>Use the detailed diagnosis of this measure to know which servers are in the farm, the status of each server, whether/not that server needs upgrade, and if so, whether/not it can be upgraded.</td>
</tr>
</table>

	Total service instances in farm: Indicates the total number of service instances in this farm.	Number	Within a farm, there are several services that run on one or more servers. These services provide basic functionality for SharePoint and regulate which services should run on which servers, in an effort to manage the impact on overall farm architecture and performance. Use the detailed diagnosis of this measure to know the services running in the farm and the servers they are running on. This way, when one or more of these servers go down, you will be able to identify the services that will be impacted.
	Servers online: Indicates the number of servers in this farm that are online currently.	Number	Use the detailed diagnosis of this measure to know which servers are online, whether/not that server needs upgrade, and if so, whether/not it can be upgraded.
	Servers offline: Indicates the number of servers in this farm that are offline currently.	Number	Ideally, the value of this measure should be 0. Use the detailed diagnosis of this measure to know which servers are offline, , whether/not that server needs upgrade, and if so, whether/not it can be upgraded.
	Servers that need upgrade: Indicates the number of servers in this farm that need to be upgraded.	Number	Use the detailed diagnosis of this measure to know which servers require an upgrade.
	Web front end servers: Indicates the number of web front end servers in this farm.	Number	Use the detailed diagnosis of this measure to know which are the web front end servers in the farm.
	Application servers: Indicates the number of application servers in this farm.	Number	Use the detailed diagnosis of this measure to know which are the application servers in the farm.
	Database servers: Indicates the number of application servers in this farm.	Number	Use the detailed diagnosis of this measure to know which database servers are in the farm.

	Online service instances: Indicates the number of service instances running in this farm that are currently online.	Number	Use the detailed diagnosis of this measure to know which services are online and which servers they are running on.
	Offline service instances: Indicates the number of service instances running in this farm that are currently offline.	Number	Ideally, the value of this measure should be 0. Use the detailed diagnosis of this measure to know which services are offline and which servers they are running on.
	Disabled service instances: Indicates the number of service instances running in this farm that are currently disabled.	Number	Use the detailed diagnosis of this measure to know which services are disabled and which servers they are running on.

Use the detailed diagnosis of the *Total servers in farm* measure to know which servers are in the farm, the status of each server, whether/not that server needs upgrade, and if so, whether/not it can be upgraded. Offline servers in the farm and the ones needing an upgrade can thus be identified.

Details of Servers in Farm

SERVER NAME	SERVER DISPLAY NAME	SERVER STATUS	SERVER ROLE	CAN UPGRADE?	NEEDS UPGRADE?
Nov 23, 2015 16:58:29					
MSPROJECT2K8R2	MSPROJECT2K8R2	Online	SingleServer	False	True

Figure 18.24: The detailed diagnosis of the Total servers in farm measure

Use the detailed diagnosis of the *Total service instances in farm* measure to know the services running in the farm and the servers they are running on. This way, when one or more of these servers go down, you will be able to identify the services that will be impacted.

Details of Service Instances in Farm

SERVER NAME	SERVICE INSTANCE NAME
Nov 23, 2015 16:58:29	
MSPROJECT2K8R2	Microsoft SharePoint Foundation Database
MSPROJECT2K8R2	Search Host Controller Service
MSPROJECT2K8R2	Information Management Policy Configuration Service
MSPROJECT2K8R2	SharePoint Server Search
MSPROJECT2K8R2	App Management Service
MSPROJECT2K8R2	Managed Metadata Web Service
MSPROJECT2K8R2	Access Services
MSPROJECT2K8R2	Microsoft SharePoint Foundation Usage
MSPROJECT2K8R2	Business Data Connectivity Service
MSPROJECT2K8R2	Search Administration Web Service

Figure 18.25: The detailed diagnosis of the Total service instances in farm measure

Use the detailed diagnosis of the *Servers online* measure to know which servers are online, whether/not that server needs upgrade, and if so, whether/not it can be upgraded.

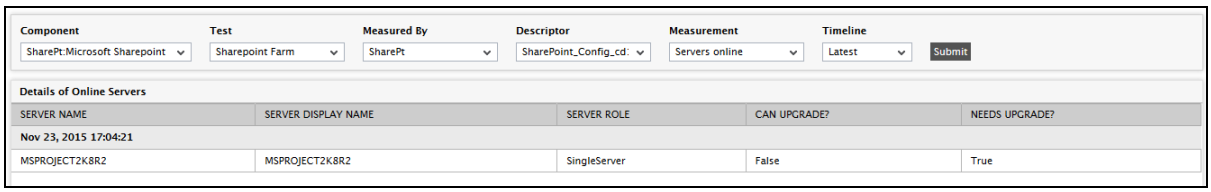

Figure 18.26: The detailed diagnosis of the Servers online measure

Use the detailed diagnosis of the *Servers that need upgrade* measure to know which servers require an upgrade.

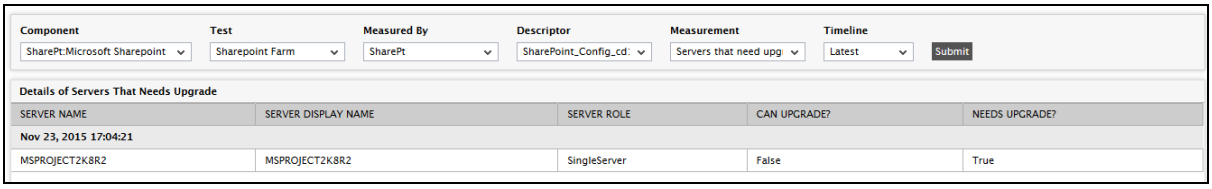

Figure 18.27: The detailed diagnosis of the Servers that need upgrade measure

Use the detailed diagnosis of the *Web front end servers* measure to know which are the web front end servers in the farm.

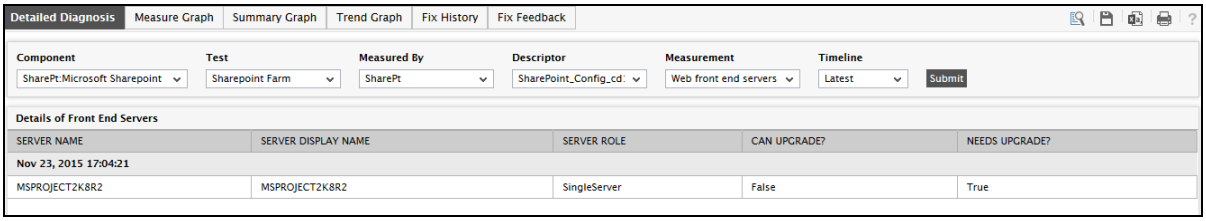

Figure 18.28: The detailed diagnosis of the Web Front end servers

To know which are the application servers in the farm, use the detailed diagnosis of the *Application servers* measure.

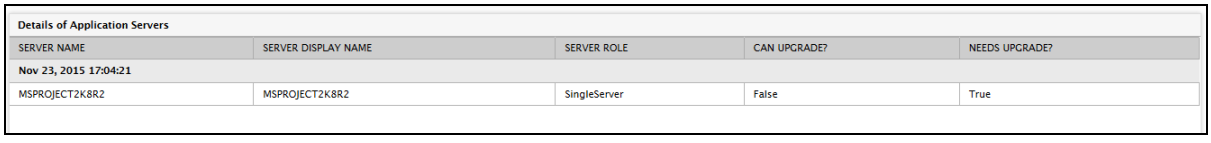

Figure 18.29: The detailed diagnosis of the Application servers measure

To identify the database servers in the farm, use the detailed diagnosis of the *Database servers* measure.

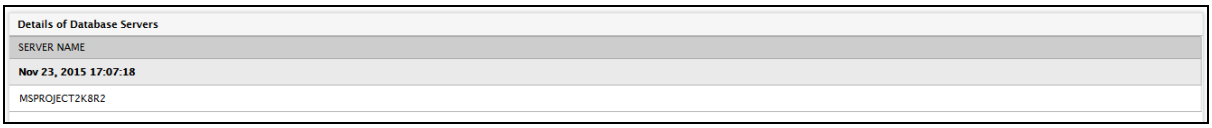

Details of Database Servers

SERVER NAME
Nov 23, 2015 17:07:18
MSPROJECT2K8R2

Figure 18.30: The detailed diagnosis of the Database servers measure

Use the detailed diagnosis of the *Online service instances* measure to know which services are online and which servers they are running on.

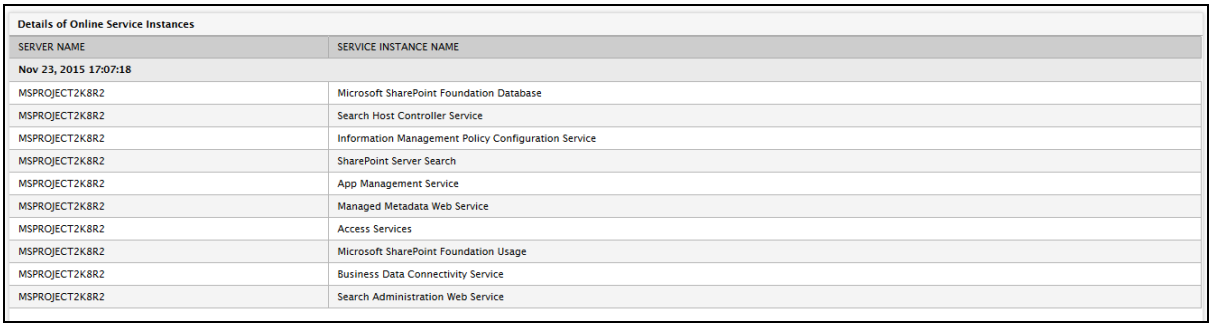

Details of Online Service Instances

SERVER NAME	SERVICE INSTANCE NAME
Nov 23, 2015 17:07:18	
MSPROJECT2K8R2	Microsoft SharePoint Foundation Database
MSPROJECT2K8R2	Search Host Controller Service
MSPROJECT2K8R2	Information Management Policy Configuration Service
MSPROJECT2K8R2	SharePoint Server Search
MSPROJECT2K8R2	App Management Service
MSPROJECT2K8R2	Managed Metadata Web Service
MSPROJECT2K8R2	Access Services
MSPROJECT2K8R2	Microsoft SharePoint Foundation Usage
MSPROJECT2K8R2	Business Data Connectivity Service
MSPROJECT2K8R2	Search Administration Web Service

Figure 18.31: The detailed diagnosis of the Online service instances measure

Use the detailed diagnosis of the *Disabled service instances* measure to know which services are disabled and which servers they are running on.

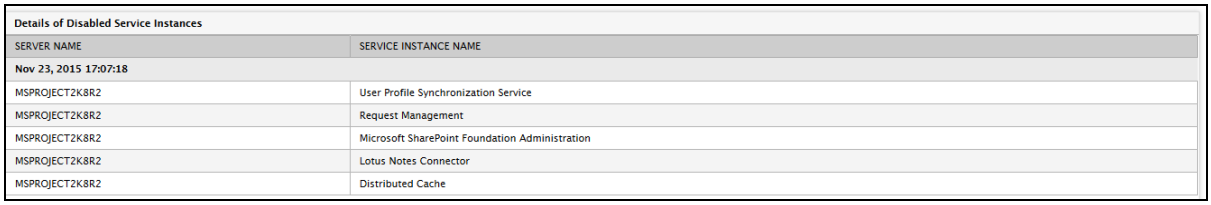

Details of Disabled Service Instances

SERVER NAME	SERVICE INSTANCE NAME
Nov 23, 2015 17:07:18	
MSPROJECT2K8R2	User Profile Synchronization Service
MSPROJECT2K8R2	Request Management
MSPROJECT2K8R2	Microsoft SharePoint Foundation Administration
MSPROJECT2K8R2	Lotus Notes Connector
MSPROJECT2K8R2	Distributed Cache

Figure 18.32: The detailed diagnosis of the Disabled service instances measure

18.2.5.3 Sharepoint Site Collections and Sites Test

A site collection is made up of one top-level site and all sites below it. As shown in the following figure, it is the top level of organization in a SharePoint 2013 web application.

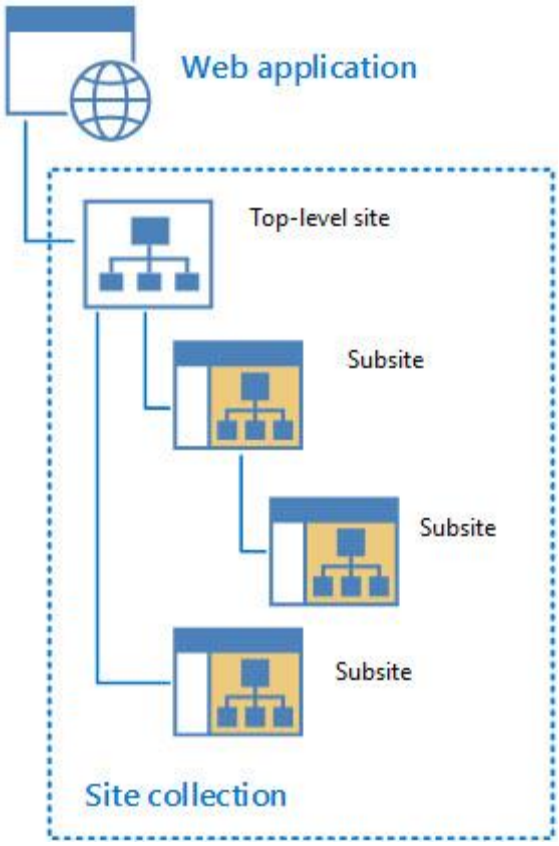

Figure 18.33: Site Collections and Sites

The number of site collections you can have in a single web application depends on the capacity of your server infrastructure.

From an architecture standpoint, all the content of a site collection must be stored in a single content database. You cannot have a site collection's content spread out across multiple content databases. Content databases scale with your infrastructure capacity so site collections can share a content database. A site collection can exist in only one content database, but one content database can host the content for multiple site collections. Similarly, any given SharePoint 2013 site can only exist in one site collection, but a site collection can host a multitude of sites. A site cannot exist outside of a site collection.

The number of site collections and sites sharing a single content database can impact the size of the database and its performance; administrators should therefore exercise restraint when associating sites and site collections with a content database. In addition, the amount of content that the sites and site collections store in their content database is also a key factor influencing the size of the content database. Variations to these two parameters – count and size - hence need to be closely monitored, so that administrators can proactively detect abnormal growth in the size of the content databases, isolate the site collections and sites that may be contributing to this, and take measures to fine-tune the site and site collection configurations to ensure peak performance of the content databases. The **Sharepoint Site Collections and Sites** test aids administrators in this endeavor!

This test captures the total number of site collections and sites on the Sharepoint server / farm and reports whether/not these numbers exceed the permissible limits. In addition, the test also tracks changes in the size of the site collections and sites over time, and promptly intimates administrators if the actual size is about to reach/exceed the size quota set for the site collection. In the process, the test points you to those site collections that are growing rapidly and the sites that may be contributing to their growth. If administrators initiate measures to curb the abnormal growth in the number or the size of the site collections and sites, they can once again take the help of this test to understand which sites and site collections are the least popular, so that such sites and site collections can be marked as probable targets for deletion or trimming.

Purpose	Captures the total number of site collections and sites on the Sharepoint server / farm and reports whether/not these numbers exceed the permissible limits. In addition, the test also tracks changes in the size of the site collections and sites over time, and promptly intimates administrators if the actual size is about to reach/exceed the size quota set for the site collection. In the process, the test points you to those site collections that are growing rapidly and the sites that may be contributing to their growth. If administrators initiate measures to curb the abnormal growth in the number or the size of the site collections and sites, they can once again take the help of this test to understand which sites and site collections are the least popular, so that such sites and site collections can be marked as probable targets for deletion or trimming
Target of the test	A Sharepoint Server 2010/2013
Agent deploying the test	An internal agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the HOST.
4. LEAST ACTIVE SITE COLLECTION DAYS – If a site collection is not modified for a duration exceeding the value (in days) specified here, then this test will count that site collection as a Least active site collection.
5. LEAST ACTIVE SITE DAYS - If a site is not modified for a duration exceeding the value (in days) specified here, then this test will count that site as a Least active site.
6. 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.
7. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for the Sharepoint Server being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

	Site collections: Indicates the number of site collections in the Sharepoint environment.	Number	The maximum recommended number of site collections per farm is: **Personal Sites -** 500,000, **Other site templates -** 250,000. The Sites can all reside on one web application, or can be distributed across multiple web applications. Note that this limit is affected by other factors that might reduce the effective number of site collections that can be supported by a given content database. Care must be exercised to avoid exceeding supported limits when a container object, such as a content database, contains a large number of other objects. For example, if a farm contains a smaller total number of content databases, each of which contains a large number of site collections, farm performance might be adversely affected long before the supported limit for the number of site collections is reached.
	Total size of site collections: Indicates the total size of all site collections in the Sharepoint environment.	MB	A site collection can be as large as the content database size limit for the applicable usage scenario. For more information about the different content database size limits for specific usage scenarios, see the Content database limits discussed in the **Interpretation** column of the *Cotent database size* measure of the **Sharepoint Content Database** test. In general, Microsoft strongly recommends limiting the size of site collections to 100 GB for the following reasons: • Certain site collection actions, such as site collection backup/restore, cause large SQL Server operations which can affect performance or fail if other site collections are active in the same database. • SharePoint site collection backup and restore is only supported for a maximum site collection size of 100 GB. For larger site collections, the complete content database must be backed up. If multiple site collections larger than 100 GB are contained in a single content database, backup and restore operations can take a long time and are at risk of failure.

	Site collections exceeding quota limit: Indicates the number of site collections that are of a size that is greater than the configured quota template.	Number	A Quota Template allows Sharepoint administrators to specify the maximum amount of content that can be stored within a Site Collection. This way, administrators can exercise greater control on the amount of content that a site collection can store in the content database, which in turn, makes for better performance and a high quality user experience with Sharepoint. A non-zero value for this measure is indicative of the fact that one/more site collections are consuming more storage resources than they should. The detailed diagnosis of this measure will lead you to those errant site collections, so that you can figure out which sites on those collections are violating the set storage thresholds.
	Least active site collections: Indicates the number of site collections that are not frequently used.	Number	This measure reports the count of those sites that were not modified for a duration greater than the value of the **LEAST ACTIVE SITE COLLECTION DAYS** parameter. You can use the detailed diagnosis of this measure to know which site collections are seldom used. If the value of the **Site collections** measure appears to be rapidly approaching the maximum recommended site collection limit, then the detailed metrics will help you identify those site collections that are rarely used and are hence candidates for removal.
	Most active site collections: Indicates the number of site collections that were modified even yesterday.	Number	Use the detailed diagnosis of this measure to identify those site collections that are very actively used.
	Users in site collections: Indicates the number of users in site collections.	Number	Besides storage, quota templates can also restrict the number of users who can be added to the Active Directory directory service from a single site collection. When the maximum number of users for a site collection has been reached, no additional user accounts can be added unless one or more user accounts are deleted from the site collection. It is hence good practice to keep an eye on the changes to this measure, so as to proactively detect a potential user quota violation.

	Number of sites: Indicates the total number of sites in site collections.	Number	Microsoft recommends the creation of a maximum of 250,000 sites and subsites per site collection. You can create a very large total number of web sites by nesting subsites. For example, in a shallow hierarchy with 100 sites, each with 1,000 subsites, you would have a total of 100,000 web sites. Compare the value of this measure across site collections to know which collection consists of the maximum number of sites.
	Total size of sites: Indicates the total size of the sites in site collections.	MB	Typically, the value of this measure will be the same as that of the *Total size of site collections* measure. A site collection can be as large as the content database size limit for the applicable usage scenario. For more information about the different content database size limits for specific usage scenarios, see the Content database limits discussed in the **Interpretation** column of the *Cotent database size* measure of the **Sharepoint Content Database** test. In general, we strongly recommend limiting the size of site collections to 100 GB for the following reasons: • Certain site collection actions, such as site collection backup/restore, cause large SQL Server operations which can affect performance or fail if other site collections are active in the same database. • SharePoint site collection backup and restore is only supported for a maximum site collection size of 100 GB. For larger site collections, the complete content database must be backed up. If multiple site collections larger than 100 GB are contained in a single content database, backup and restore operations can take a long time and are at risk of failure.

	Most active sites: Indicates the number of sites that were accessed even yesterday.	Number	Use the detailed diagnosis of this measure to identify those site collections that are very actively used.
	Least active sites: Indicates the number of sites that are not used frequently.	Number	This measure reports the count of those sites that were not modified for a duration greater than the value of the **LEAST ACTIVE SITE DAYS** parameter. You can use the detailed diagnosis of this measure to know sites are seldom used. If the value of the *Number of sites* measure appears to be rapidly approaching the maximum recommended site limit, then the detailed metrics will help you identify those sites that are rarely used and are hence candidates for removal.

The detailed diagnosis of the *Least active site collections* measure reveals the top 10 site collections that were used the least. In times of rapid web application growth, this list will indicate those site collections that can be removed to curb the growth.

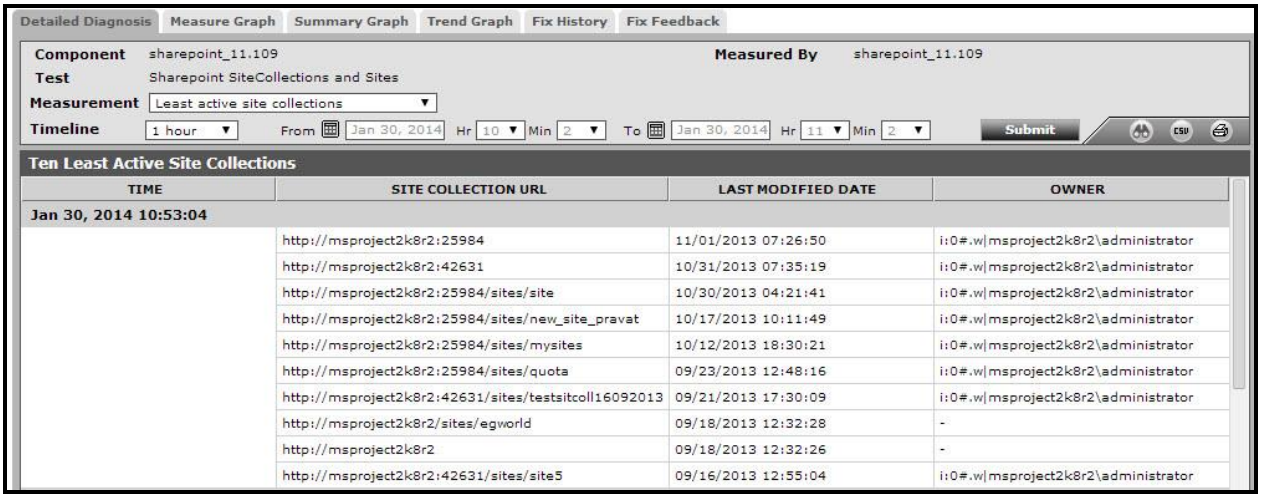

Figure 18.34: The detailed diagnosis of the Least active site collections measure

The detailed diagnosis of the *Least active sites* measure reveals the top 10 sites that were used the least. In times of rapid growth in the size of a site collection, this list will indicate those sites that can be removed to curb the growth.

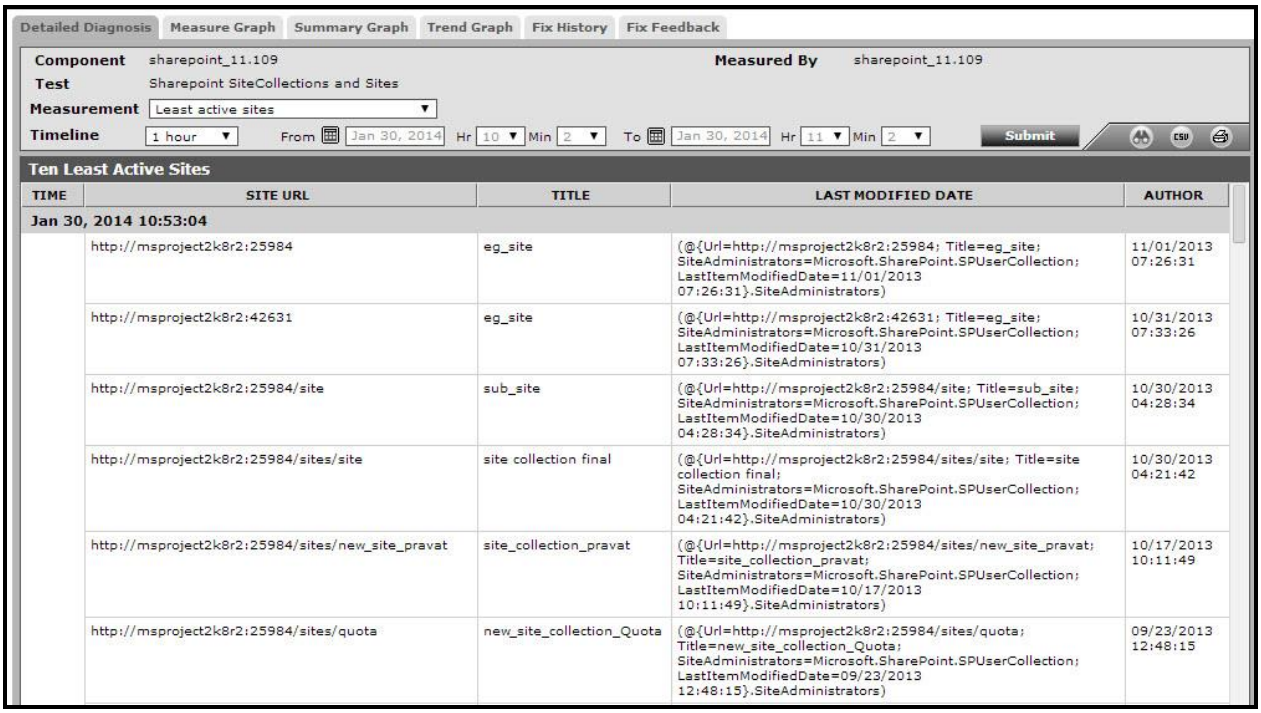

Figure 18.35: The detailed diagnosis of the Least active sites measure

18.2.5.4 Sharepoint Web Applications Test

Web Applications (WAs) are top-level containers for content in a SharePoint farm, and are typically the interface through which a user interacts with SharePoint - site collections, lists, and libraries come underneath the web application. A web application is associated with a set of access mappings or URLs which are defined in the SharePoint central management console, then automatically replicated into the IIS configuration of every server configured in the farm. WAs are typically independent of each other, have their own application pools, and can be restarted independently in Internet Information Services. Web Applications provide the ability to isolate content, processes, features and users. For example, you can separate the content anonymous users can see vs. what authenticated users can see by hosting the same content in different web apps.

A web application can grow in size over time! If this growth is not kept under control, then you may end up with a situation where a few web applications are hogging the storage resources provided by the Sharepoint environment, leaving the other web applications with limited to no resources! To avoid this, administrators need to be able to quickly isolate the web applications that are growing rapidly, understand their composition, and isolate the reasons for the abnormal growth. The **Sharepoint Web Applications** test helps administrators with this! For each web application deployed on a Sharepoint server, this test monitors the current size of that web application and captures a consistent increase in the size of the same, thus pointing you to those web applications that are growing in size at a steady pace and the content databases they are using. In addition, the test also leads you to the probable reasons for the abnormal size of the web application – is it because the web application is handling documents of huge sizes? or is it because the web application is storing too many versions of a document, which is in fact adding to its size? Or is it owing to the numerous sites, site collections, and document libraries that are being hosted by that web application?

Purpose	For each web application deployed on a Sharepoint server, this test monitors the current size of that web application and captures a consistent increase in the size of the same, thus pointing you to those web applications that are growing in size at a steady pace and the content databases they are using

<table>
<tr><td>Target of the test</td><td colspan="3">A Sharepoint Server 2010/2013</td></tr>
<tr><td>Agent deploying the test</td><td colspan="3">An internal agent</td></tr>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the HOST.
4. 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.
5. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each web application on the Sharepoint Server being monitored</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Size of this web application:
Indicates the current size of this web application.</td><td>GB</td><td></td></tr>
</table>

	Web application growth rate: Indicates the percentage growth in the size of this web application since the last measurement period.	Percent	Compare the value of this measure across web applications to know which web application has grown the maximum since the previous measurement period. By closely tracking the variations in this measure for that web application over time, you can determine whether/not the web application is growing rapidly in size! If so, it is a cause for concern, as it indicates that that web application has the potential of consuming all available storage resources! In such a situation, you may want to reset the size limit for the site collections that are within the web application, so as to curb its growth. A site collection can be as large as the content database size limit for the applicable usage scenario. For more information about the different content database size limits for specific usage scenarios, see the Content database limits discussed in the **Interpretation** column of the *Cotent database size* measure of the **Sharepoint Content Database** test. In general, Microsoft strongly recommends limiting the size of site collections to 100 GB for the following reasons: • Certain site collection actions, such as site collection backup/restore, cause large SQL Server operations which can affect performance or fail if other site collections are active in the same database. • SharePoint site collection backup and restore is only supported for a maximum site collection size of 100 GB. For larger site collections, the complete content database must be backed up. If multiple site collections larger than 100 GB are contained in a single content database, backup and restore operations can take a long time and are at risk of failure.

	Users in this web application: Indicates the number of users in this web application.	Number	Compare the value of this measure across web applications to identify that application which has the maximum number of users.
	Content databases used by this web application: Indicates the number of content databases that were used by this web application.	Number	
	Site collections part of this web application: Indicates the number of site collections in this web application.	Number	The maximum recommended number of site collections per farm is: **Personal Sites -** 500,000, **Other site templates -** 250,000. The Sites can all reside on one web application, or can be distributed across multiple web applications. Compare the value of this measure across web applications to know which application consists of the maximum number of site collections. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of site collections has anything to do with the increase in web application size.
	Sites part of this web application: Indicates the total number of sites in the site collections that are part of this web application.	Number	Microsoft recommends the creation of a maximum of 250,000 sites and subsites per site collection. You can create a very large total number of web sites by nesting subsites. For example, in a shallow hierarchy with 100 sites, each with 1,000 subsites, you would have a total of 100,000 web sites. Compare the value of this measure across web applications to know which application consists of the maximum number of sites. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of sites has anything to do with the increase in web application size.

	Number of document libraries: Indicates the number of document libraries in this web application.	Number	Document libraries are collections of files that you can share with team members on a Web based on Microsoft Windows SharePoint Services. By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of document libraries. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of document libraries has anything to do with the increase in web application size.
	Lists in this web application: Indicates the number of lists in this web application.	Number	A list in SharePoint is used to store data across columns in separate rows. By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of Sharepoint lists. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of lists has in any way impacted the web application size.
	Attachments: Indicates the number of attachments in this web application.	Number	By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of attachments. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of attachments has in any way impacted the web application size.
	Documents in this web application: Indicates the total number of documents in this web application.	Number	By comparing the value of this measure across web applications, you can figure out which web application has the maximum number of documents. In the event of a sudden increase in the size of a web application, you can check how the value of this measure has grown over the same period to figure out whether/not the addition of documents has in any way impacted the web application size.

	Size of documents: Indicates the total size of all documents that are available in this web application.	GB	Compare the value of this measure across web applications to identify that application with the maximum document size. This can be attributed to the existence of one/more large-sized documents or many moderately sized documents in the web application. If that web application appears to be growing in size rapidly, you may want to keep an eye on this measure to figure out if it is owing to the increase in document size.
	Document versions: Indicates the number of document versions in this web application.	Number	Typically, Sharepoint can support a maximum of 40,000 major versions and 511 minor versions of documents. If this limit is exceeded basic file operations—such as file open or save, delete, and viewing the version history— may not succeed.
	Average number of documents per document library: Indicates the average number of documents per library in this web application.	Number	

18.2.5.5 SharePoint Web Parts Test

By using web parts, you can modify the content, appearance, and behavior of pages of a SharePoint site by using a browser. Web parts are server-side controls that run inside a web part page: they're the building blocks of pages that appear on a SharePoint site.

For problem detection and troubleshooting purposes, administrators should know which web parts operate within a web application, which ones are open presently, and which ones are closed. The **SharePoint Web Parts** test provides this insight to the administrators. For each web application on SharePoint, this test reports the count of web parts in that web application, and the number of open and closed web parts in that web application. Detailed diagnosis of this test also reveals the names of the open and closed web parts.

Purpose	For each web application on SharePoint, this test reports the count of web parts in that web application, and the number of open and closed web parts in that web application. Detailed diagnosis of this test also reveals the names of the open and closed web parts.
Target of the test	A Sharepoint Server 2010/2013
Agent deploying the test	An internal/remote agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the HOST.
4. DOMAIN, DOMAIN USER, PASSWORD, and CONFIRM PASSWORD – This test requires the credentials of a farm administrator to run and collect metrics from the target SharePoint server. Therefore, specify the domain to which that farm administrator belongs in the DOMAIN text box, and then, enter the credentials of the farm administrator in the DOMAIN USER and PASSWORD text boxes. To confirm the password, retype it in the CONFIRM PASSWORD text box.
5. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each web application in the SharePoint server</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Total web parts:
Indicates the number of web parts in this web application.</td><td>Number</td><td></td></tr>
<tr><td></td><td>Open web parts:
Indicates the number of open web parts in this web application.</td><td>Number</td><td>Use the detailed diagnosis of this measure to determine which web parts are open in the target web application.</td></tr>
<tr><td></td><td>Closed web parts:
Indicates the number of closed web parts in this web application.</td><td>Number</td><td>Use the detailed diagnosis of this measure to determine which web parts are closed in the target web application.</td></tr>
</table>

Use the detailed diagnosis of the *Open web parts* measure to determine which web parts are open in the target web application.

Details of Open Webparts

SITE URL	DISPLAY TITLE	TYPE
Dec 23, 2015 03:07:57		
http://2k8r2sp2k13:18180/GettingStarted.aspx	Get started with your site	Microsoft.SharePoint.WebPartPages.GettingStartedWebPart
http://2k8r2sp2k13:18180/SitePages/Home.aspx	Site Assets	Microsoft.SharePoint.WebPartPages.XsltListViewWebPart
http://2k8r2sp2k13:18180/SitePages/Home.aspx	Documents	Microsoft.SharePoint.WebPartPages.XsltListViewWebPart

Figure 18.36: The detailed diagnosis of the Open web parts measure

Use the detailed diagnosis of the *Closed web parts* measure to determine which web parts are closed in the target web application.

Details of Closed Webparts

SITE URL	DISPLAY TITLE	TYPE
Dec 23, 2015 03:07:57		
http://2k8r2sp2k13:18180/SitePages/Home.aspx	Site Pages	Microsoft.SharePoint.WebPartPages.XsltListViewWebPart
http://2k8r2sp2k13:18180/SitePages/Home.aspx	Documents	Microsoft.SharePoint.WebPartPages.XsltListViewWebPart

Figure 18.37: The detailed diagnosis of the Closed web parts measure

18.2.6 The Sharepoint Usage Analytics Layer

The tests mapped to this layer report a wide variety of usage analytics that measure the experience of users of web sites, web applications, browsers, distributed cache, and web parts on SharePoint, and reports abnormalities.

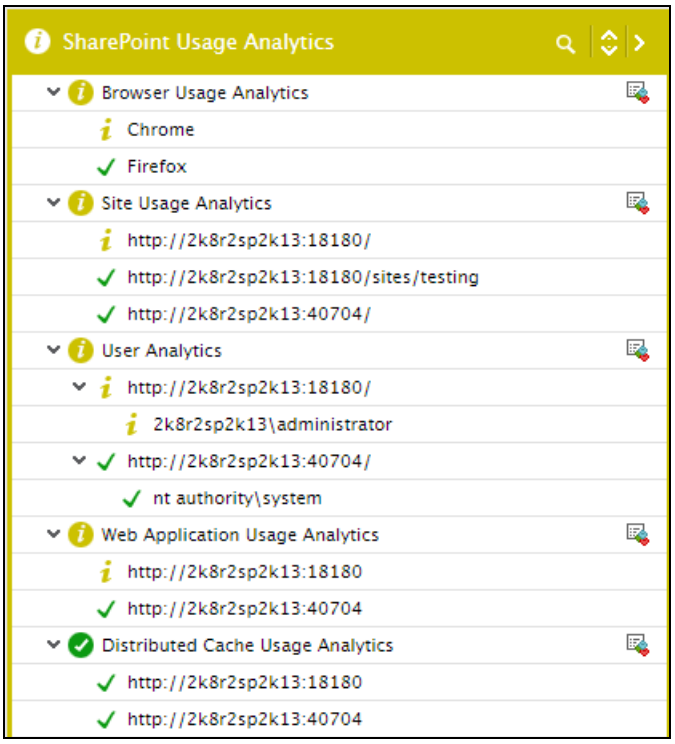

Figure 18.38: The tests mapped to the Sharepoint Usage Analytics layer

18.2.6.1 Site Usage Analytics Test

Enterprises typically use SharePoint to create web sites and web applications. The success of the SharePoint platform therefore hinges on how happy users are when interacting with the web sites that it helped create. If the number of visitors to a web site keeps increasing, it is indicative of an increase in the web site's popularity, which directly translates into 'many happy users'! Likewise, if users to a web site constantly complain of slowness when browsing that web site, it indicates that user experience with the web site is unsatisfactory – meaning, 'many unhappy users'. This in turn can hit user productivity badly, escalate troubleshooting time and costs of the enterprise, and adversely impact its revenues and reputation! To improve user experience with SharePoint sites and to build user confidence in the SharePoint platform, administrators should be able to quickly identify slow web sites and precisely pinpoint the reason for the slowness.

This is where the **Site Usage Analytics** test helps! This test queries the SharePoint **usage database** at configured intervals and collects metrics on web site usage that is stored therein – this includes the web sites accessed, count of hits to each web site, users who browsed every site, the browsers that were used for web site access, web pages requested, the time taken for the requested pages to load, where page views spent time and how much, error responses returned, resources consumed, and many more. For each web site configured for monitoring, the test then reports the average time taken by that site to load pages. In the process, the test points administrators to slow web sites and also leads them to the probable source of the slownes – is it owing to a latent web front end? is it because of slow service calls? Or is it due to inefficient queries to the backend database?

Sometimes, poor user experience can be attributed to HTTP errors. This is why, this test instantly alerts administrators to HTTP error responses, thus ensuring their timely intervention and rapid resolution of the error conditions.

This way, the **Site Usage Analytics** test enables administrators to detect web site slowness well before users notice, helps them promptly and accurately diagnose the source of the poor user experience with a web site, and thus

ensures that they initiate measures to enhance user experience and pre-empt the damage that may be caused to revenue and reputation.

Note that this test will run only if a SharePoint Usage and Health Service application is created and is configured to collect usage and health data. To know how to create and configure this application, follow the steps detailed in Section 18.2.6.1.1.

Purpose	Enables administrators to detect web site slowness well before users notice, promptly and accurately diagnose the source of the poor user experience with a web site, and initiate measures to enhance user experience and pre-empt the damage that may be caused to revenue and reputation
Target of the test	A Sharepoint Server 2010/2013
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **SQL PORT NUMBER** – Specify the port number of the SQL server hosting the usage database. 5. **INSTANCE** – If the SQL server hosting the usage database is instance-based, then provide the instance name here. If not, then set this to *none*. 6. **SSL** – If the SQL server hosting the usage database is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**. 7. **ISNTLMV2** - In some Windows networks, **NTLM** (**NT LAN Manager**) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the **ISNTLMV2** flag is set to **No**, indicating that NTLMv2 is not enabled by default on the SQL server hosting the usage database. Set this flag to **Yes** if NTLMv2 is enabled on that SQL server. 8. **DATABASE SERVER NAME** – Specify the name of Microsoft SQL server that hosts the usage database to be accessed by this test. 9. **DATABASE NAME** – Specify the name of the usage database that this test should access. 10. **DATABASE USER NAME, DATABASE PASSWORD, CONFIRM PASSWORD** - Specify the credentials of a user who has read-only access to the usage database configured, in the **DATABASE USER NAME** and **DATABASE PASSWORD** text boxes. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box. 11. **SITE** – Configure a comma-separated list of web site URLs that you want this test to monitor. For eg., *http://www.msproject28rk2:11982,http://www.mydocs.com* 12. **SLOW TRANSACTION CUTOFF (MS)** - This test reports the count of slow page views and also pinpoints the pages that are slow. To determine whether/not a page is slow, this test uses the **SLOW TRANSACTION CUTOFF** parameter. By default, this parameter is set to *4000 millisecs* (i.e., 4 seconds). This means that, if a page takes more than 4 seconds to load, this test will consider that page as a slow page by default. You can increase or decrease this slow transaction cutoff according to what is 'slow' and what is 'normal' in your environment.  Note The default value of this parameter is the same as the default *Maximum threshold* setting of the *Avg page load time* measure – i.e., both are set to *4000 millisecs* by default. While the former helps eG to distinguish between slow and healthy page views for the purpose of providing detailed diagnosis, the latter tells eG when to generate an alarm on *Avg page load time*. For best results, it is recommended that both these settings are configured with the same value at all times. Therefore, if you change the value of one of these configurations, then make sure you update the value of the other as well. For instance, if the **SLOW TRASACTION CUTOFF** is changed to *6000 millisecs*, change the *Maximum Threshold* of the *Avg page load time* measure to *6000* millsecs as well.

	13. **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.		
Outputs of the test	One set of results for each **SITE** configured for monitoring		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Unique users: Indicates the number of unique users of this web site.	Number	The detailed diagnosis of this measure reveals the names of the unique users and the number of requests from each user to the web site being monitored. From this, you can identify those users who are actively using the web site.
	Unique visitors: Indicates the number of unique visitors to this web site.	Number	SharePoint authenticated users and anonymous users (using IP address) are counted as visitors. Compare the value of this measure across sites to identify the most popular SharePoint site. You can use the detailed diagnosis of this measure to know who are the unique visitors to the web site and the number of requests from each visitor to the web site. This way, you can identify that visitor who visits the web site most frequently.
	Unique destinations: Indicates the number of unique destinations of this site.	Number	To know the most popular destination URLs of this site, use the detailed diagnosis of this measure. Here, you will find the top-10 destinations in terms of the number of hits.

	Unique browsers: Indicates the number of unique browsers used for accessing this site.	Number	To know which browsers are commonly used to access this web site, use the detailed diagnosis of this measure. Here, the unique browsers will be listed and the number of hits to the web site from each browser will be displayed alongside, so that you can instantly identify that browser that has been widely used to access the web site.
	Unique referrers: Indicates the number of unique URLs external to this site (parent site is treated as external as well), from where the users navigated to this site.	Number	To know which referrer URL was responsible for the maximum hits to this web site, use the detailed diagnosis of this measure. The top-10 unique referrer URLs in terms of the number of hits they generated will be displayed as part of the detailed diagnostics.

	Apdex score: Indicates the apdex score of this site.	Number	**Apdex (Application Performance Index)** is an open standard developed by an alliance of companies. It defines a standard method for reporting and comparing the performance of software applications in computing. Its purpose is to convert measurements into insights about user satisfaction, by specifying a uniform way to analyze and report on the degree to which measured performance meets user expectations. The Apdex method converts many measurements into one number on a uniform scale of 0-to-1 (0 = no users satisfied, 1 = all users satisfied). The resulting Apdex score is a numerical measure of user satisfaction with the performance of enterprise applications. This metric can be used to report on any source of end-user performance measurements for which a performance objective has been defined. The Apdex formula is: $Apdex_t = (Satisfied\ Count + Tolerating\ Count / 2) / Total\ Samples$ This is nothing but the number of satisfied samples plus half of the tolerating samples plus none of the frustrated samples, divided by all the samples. A score of 1.0 means all responses were satisfactory. A score of 0.0 means none of the responses were satisfactory. Tolerating responses half satisfy a user. For example, if all responses are tolerating, then the Apdex score would be 0.50. Ideally therefore, the value of this measure should be *1.0*. A value less than *1.0* indicates that the user experience with the web site has been less than satisfactory.

	Satisfied page views: Indicates the number of times pages in this web site were viewed without any slowness.	Number	A page view is considered to be slow when the average time taken to load that page exceeds the **SLOW TRANSACTION CUTOFF** configured for this test. If this **SLOW TRANSACTION CUTOFF** is not exceeded, then the page view is deemed to be 'satisfactory'. Ideally, the value of this measure should be high. If the value of this measure is much lesser than the value of the *Tolerating page views* and the *Frustrated page views*, it is a clear indicator that the experience of the users of this web site is below-par. In such a case, use the detailed diagnosis of the *Tolerating page views* and *Frustrated page views* measures to know which pages are slow.
	Tolerating page views: Indicates the number of tolerating page views to this web site.	Number	If the *Average page load time* of a page exceeds the **SLOW TRANSACTION CUTTOFF** configuration of this test, but is less than 4 times the **SLOW TRANSACTION CUTOFF** (i.e., < 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Tolerating page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the overall user experience from this browser is less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.
	Frustrated page views: Indicates the number of frustrated page views to this web site.	Number	If the *Average page load time* of a page is over 4 times the **SLOW TRANSACTION CUTTOFF** configuration of this test (i.e., > 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Frustrated page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the experience of users using this browser has been less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.

	Average page load time: Indicates the average time taken by the pages in this web site to load completely.	Secs	This is the average interval between the time that a user initiates a request and the completion of the page load of the response in the user's browser. If the value of this measure is consistently high for a web site, there is reason to worry. This is because, it implies that the web site is slow in responding to requests. If this condition is allowed to persist, it can adversely impact user experience with the web site. You may want to check the *Apdex score* in such circumstances to determine whether/not user experience has already been affected. Regardless, you should investigate the anomaly and quickly determine where the bottleneck lies – is it with the web front-end? is it owing to slow service calls? Or is it because of inefficient queries to the backend? - so that the problem can be fixed before users even notice any slowness! For that, you may want to compare the values of the *Web front-end processing time*, *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test. If the *Average front end time* is the highest, it indicates that the problem is with the web site/web application front end – this can be attributed to a slowdown in page rendering or in DOM building. If the *Average server connection time* is the highest, it denotes that the network is the problem source. This in turn can be caused by TCP connection latencies and delays in domain look up. On the other hand, if the *Average response available time* measure registers the highest value, it indicates that the problem lies with the web site/web application backend – i.e., the web/web application server that is hosting the web site/web application being monitored.

	Web front-end processing time: Indicates the average time in milliseconds it took for the web front end server to process the requests to this web site.	Msecs	If the *Avg page load time* of a web site is abnormally high, then you can compare the value of this measure with that of the *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average service calls duration: Indicates the time taken by this web site to generate service calls.	Msecs	If the *Avg page load time* of a web site is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average IIS latency: Indicates the average time requests to this web site took in the frontend web server after the requests were received by the frontend web server but before this web site began processing the requests.	MSecs	If the *Avg page load time* of a web site is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average CPU duration*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average CPU duration: Indicates the average time for which requests to this web site used the CPU.	Msecs	If the *Avg page load time* of a web site is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	SQL logical reads: Indicates the total number of 8 kilobyte blocks that this web site read from storage on the back-end database server.	Number	

	Average CPU megacycles: Indicates the average number of CPU mega cycles spent processing the requests to this web site in the client application on the front end web server.	Number	
	Total queries: Indicates the total number of database queries generated for this site.	Number	
	Average query duration: Indicates the average time taken for all backend database queries generated for this site.	Msecs	If the *Avg page load time* of a web site is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration*, *Average IIS latency*, and *Average CPU duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average data consumed: Indicates the average bytes of data downloaded by requests to this web site.	KB	
	GET requests: Indicates the number of GET requests to this web site.	Number	
	POST requests: Indicates the number of POST requests to this web site.	Number	
	OPTION requests: Indicates the number of OPTION request to this web site.	Number	
	300 responses: Indicates the number of responses to requests to this web site with a status code in the 300-399 range	Number	300 responses could indicate page caching on the client browsers. Alternatively 300 responses could also indicate redirection of requests. A sudden change in this value could indicate a problem condition.

	400 errors: Indicates the number responses to requests to this web site that had a status code in the range 400-499.	Number	A high value indicates a number of missing/error pages. Use the detailed diagnosis of this measure to know when each of the 400 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.
	500 errors: Indicates the number of responses to the requests to this web site that had a status code in the range 500-599.	Number	Since responses with a status code of 500-600 indicate server side processing errors, a high value reflects an error condition. Use the detailed diagnosis of this measure to know when each of the 500 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.

The detailed diagnosis of the *Unique users* measure reveals the names of the unique users and the number of requests from each user to the web site being monitored. From this, you can identify those users who are actively using the web site.

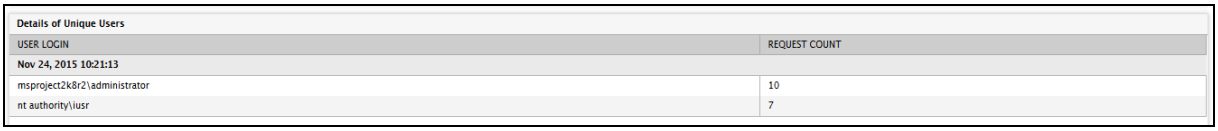

Details of Unique Users

USER LOGIN	REQUEST COUNT
Nov 24, 2015 10:21:13	
msproject2k8r2\administrator	10
nt authority\iusr	7

Figure 18.39: The detailed diagnosis of the Unique users measure

You can use the detailed diagnosis of the *Unique visitors* measure to know who are the unique visitors to the web site and the number of requests from each visitor to the web site. This way, you can identify that visitor who visits the web site most frequently.

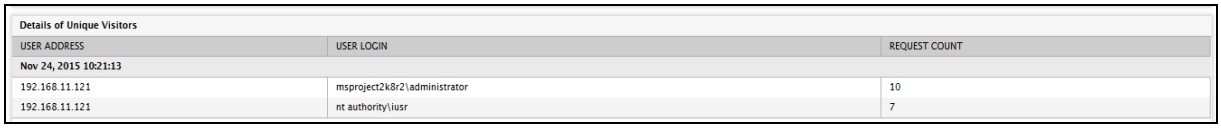

Details of Unique Visitors

USER ADDRESS	USER LOGIN	REQUEST COUNT
Nov 24, 2015 10:21:13		
192.168.11.121	msproject2k8r2\administrator	10
192.168.11.121	nt authority\iusr	7

Figure 18.40: The detailed diagnosis of the Unique visitors measure

To know the most popular destination URLs of this site, use the detailed diagnosis of the *Unique destinations* measure. Here, you will find the top-10 destinations in terms of the number of hits.

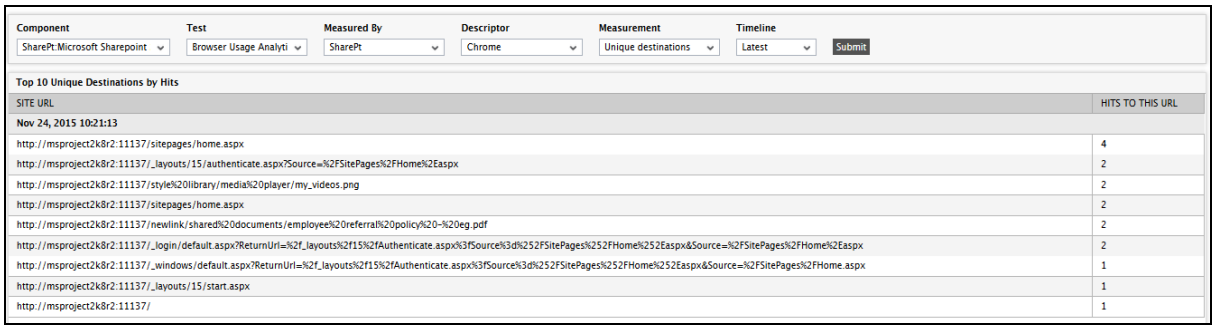

Component: SharePt:Microsoft Sharepoint | Test: Browser Usage Analyti | Measured By: SharePt | Descriptor: Chrome | Measurement: Unique destinations | Timeline: Latest | Submit

Top 10 Unique Destinations by Hits

SITE URL	HITS TO THIS URL
Nov 24, 2015 10:21:13	
http://msproject2k8r2:11137/sitepages/home.aspx	4
http://msproject2k8r2:11137/_layouts/15/authenticate.aspx?Source=%2FSitePages%2FHome%2Easpx	2
http://msproject2k8r2:11137/style%20library/media%20player/my_videos.png	2
http://msproject2k8r2:11137/sitepages/home.aspx	2
http://msproject2k8r2:11137/newlink/shared%20documents/employee%20referral%20policy%20-%20eg.pdf	2
http://msproject2k8r2:11137/_login/default.aspx?ReturnUrl=%2f_layouts%2f15%2fAuthenticate.aspx%3fSource%3d%252FSitePages%252FHome%252Easpx&Source=%2FSitePages%2FHome%2Easpx	2
http://msproject2k8r2:11137/_windows/default.aspx?ReturnUrl=%2f_layouts%2f15%2fAuthenticate.aspx%3fSource%3d%252FSitePages%252FHome%252Easpx&Source=%2FSitePages%2FHome.aspx	1
http://msproject2k8r2:11137/_layouts/15/start.aspx	1
http://msproject2k8r2:11137/	1

Figure 18.41: The detailed diagnosis of the Unique destinations measure

To know which referrer URL was responsible for the maximum hits to this web site, use the detailed diagnosis of the *Unique referrers* measure. The top-10 unique referrer URLs in terms of the number of hits they generated will be displayed as part of the detailed diagnostics.

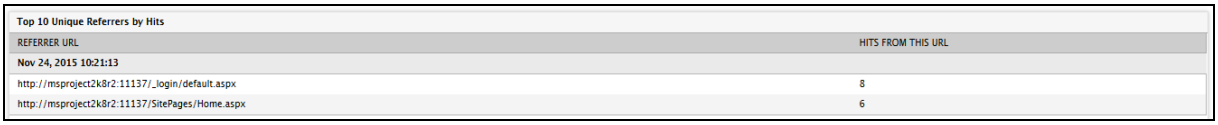

Top 10 Unique Referrers by Hits

REFERRER URL	HITS FROM THIS URL
Nov 24, 2015 10:21:13	
http://msproject2k8r2:11137/_login/default.aspx	8
http://msproject2k8r2:11137/SitePages/Home.aspx	6

Figure 18.42: The detailed diagnosis of the Unique referrers measure

If the *Tolerating page views* measure reports a non-zero value, then use the detailed diagnosis of this measure to view the top-10 pages in terms of page load time. From the detailed metrics, you can rapidly identify the URL of the page that took the longest to load, the load time of that page, when the slowness occurred, and which user's access was impacted by the slowness. Additionally, usage analytics such as the count of requests to the slow page, the count of queries run by the page, the amount of data consumed, and the status of the HTTP access to the page are also revealed as part of the detailed diagnosis.

Top 10 Tolerating Page Views by Load Time

LOG TIME	USER ADDRESS	USER LOGIN	SITE URL	REFERRER URL	PAGE LOAD TIME(SEC)	MACHINE NAME	BROWSER	REQUEST COUNT	QUERY COUNT	REQUEST TYPE	BYTES CONSUMED	HTTP STATUS
Nov 24, 2015 10:21:13												
2015-11-23 07:24:35.37	192.168.11.121	msproject2k8r2\administrator	http://msproject2k8r2:11137/newlink /shared%20documents /employee%20referral%20policy%20-%20eg.pdf	http://msproject2k8r2:11137 /SitePages/Home.aspx	13	MSPROJECT2K8R2	Chrome	0	2	GET	0	200

Figure 18.43: The detailed diagnosis of the Tolerating page views measure

If the *Frustrated page views* measure reports a non-zero value, then use the detailed diagnosis of this measure to view the top-10 pages in terms of page load time. From the detailed metrics, you can rapidly identify the URL of the page that took the longest to load, the load time of that page, when the slowness occurred, and which user's access was impacted by the slowness. Additionally, usage analytics such as the count of requests to the slow page, the count of queries run by the page, the amount of data consumed, and the status of the HTTP access to the page are also revealed as part of the detailed diagnosis.

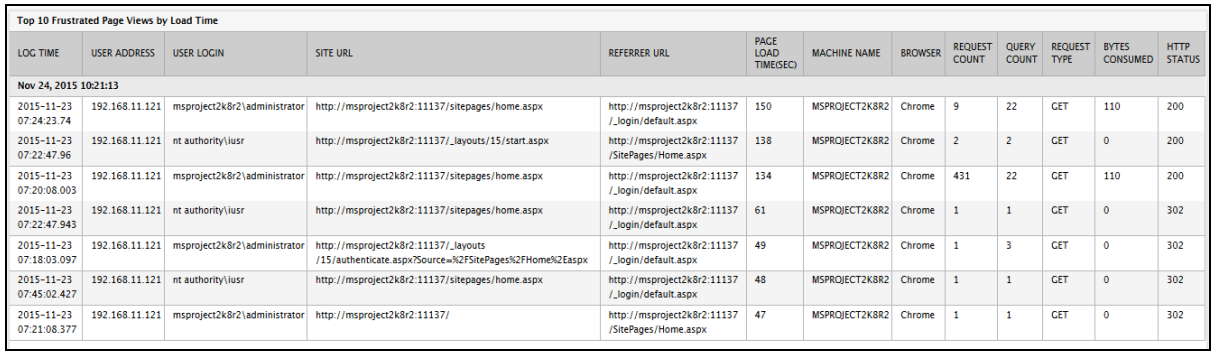

Top 10 Frustrated Page Views by Load Time

LOG TIME	USER ADDRESS	USER LOGIN	SITE URL	REFERRER URL	PAGE LOAD TIME(SEC)	MACHINE NAME	BROWSER	REQUEST COUNT	QUERY COUNT	REQUEST TYPE	BYTES CONSUMED	HTTP STATUS
Nov 24, 2015 10:21:13												
2015-11-23 07:24:23.74	192.168.11.121	msproject2k8r2\administrator	http://msproject2k8r2:11137/sitepages/home.aspx	http://msproject2k8r2:11137/_login/default.aspx	150	MSPROJECT2K8R2	Chrome	9	22	GET	110	200
2015-11-23 07:22:47.96	192.168.11.121	nt authority\iusr	http://msproject2k8r2:11137/_layouts/15/start.aspx	http://msproject2k8r2:11137/SitePages/Home.aspx	138	MSPROJECT2K8R2	Chrome	2	2	GET	0	200
2015-11-23 07:20:08.003	192.168.11.121	msproject2k8r2\administrator	http://msproject2k8r2:11137/sitepages/home.aspx	http://msproject2k8r2:11137/_login/default.aspx	134	MSPROJECT2K8R2	Chrome	431	22	GET	110	200
2015-11-23 07:22:47.943	192.168.11.121	nt authority\iusr	http://msproject2k8r2:11137/sitepages/home.aspx	http://msproject2k8r2:11137/_login/default.aspx	61	MSPROJECT2K8R2	Chrome	1	1	GET	0	302
2015-11-23 07:18:03.097	192.168.11.121	msproject2k8r2\administrator	http://msproject2k8r2:11137/_layouts/15/authenticate.aspx?Source=%2FSitePages%2FHome%2Easpx	http://msproject2k8r2:11137/_login/default.aspx	49	MSPROJECT2K8R2	Chrome	1	3	GET	0	302
2015-11-23 07:45:02.427	192.168.11.121	nt authority\iusr	http://msproject2k8r2:11137/sitepages/home.aspx	http://msproject2k8r2:11137/_login/default.aspx	48	MSPROJECT2K8R2	Chrome	1	1	GET	0	302
2015-11-23 07:21:08.377	192.168.11.121	msproject2k8r2\administrator	http://msproject2k8r2:11137/	http://msproject2k8r2:11137/SitePages/Home.aspx	47	MSPROJECT2K8R2	Chrome	1	1	GET	0	302

Figure 18.44: The detailed diagnosis of the Frustrated page views measure

Use the detailed diagnosis of the *400 errors* and *500 errors* measures to know when each of the 400 or 500 errors (as the case may be) occurred, which user experienced the error, when, using what browser, from which machine. This information will greatly aid troubleshooting.

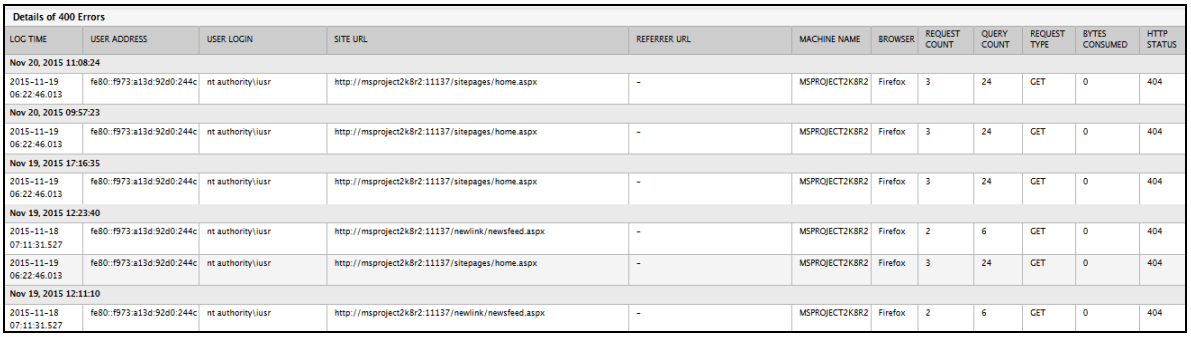

Details of 400 Errors

LOG TIME	USER ADDRESS	USER LOGIN	SITE URL	REFERRER URL	MACHINE NAME	BROWSER	REQUEST COUNT	QUERY COUNT	REQUEST TYPE	BYTES CONSUMED	HTTP STATUS
Nov 20, 2015 11:08:24											
2015-11-19 06:22:46.013	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/sitepages/home.aspx	-	MSPROJECT2K8R2	Firefox	3	24	GET	0	404
Nov 20, 2015 09:57:23											
2015-11-19 06:22:46.013	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/sitepages/home.aspx	-	MSPROJECT2K8R2	Firefox	3	24	GET	0	404
Nov 19, 2015 17:16:35											
2015-11-19 06:22:46.013	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/sitepages/home.aspx	-	MSPROJECT2K8R2	Firefox	3	24	GET	0	404
Nov 19, 2015 12:23:40											
2015-11-18 07:11:31.527	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/newlink/newsfeed.aspx	-	MSPROJECT2K8R2	Firefox	2	6	GET	0	404
2015-11-19 06:22:46.013	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/sitepages/home.aspx	-	MSPROJECT2K8R2	Firefox	3	24	GET	0	404
Nov 19, 2015 12:11:10											
2015-11-18 07:11:31.527	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/newlink/newsfeed.aspx	-	MSPROJECT2K8R2	Firefox	2	6	GET	0	404

Figure 18.45: The detailed diagnosis of the 400 errors measure

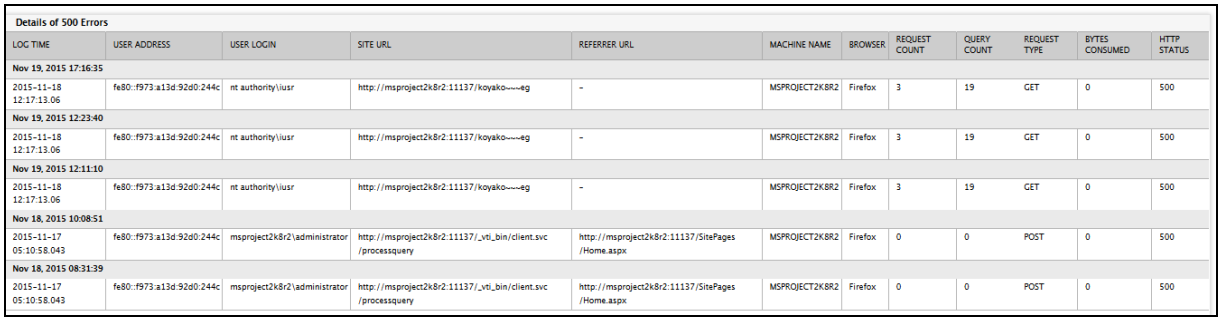

Details of 500 Errors

LOG TIME	USER ADDRESS	USER LOGIN	SITE URL	REFERRER URL	MACHINE NAME	BROWSER	REQUEST COUNT	QUERY COUNT	REQUEST TYPE	BYTES CONSUMED	HTTP STATUS
Nov 19, 2015 17:16:35											
2015-11-18 12:17:13.06	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/koyako~~~eg	-	MSPROJECT2K8R2	Firefox	3	19	GET	0	500
Nov 19, 2015 12:23:40											
2015-11-18 12:17:13.06	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/koyako~~~eg	-	MSPROJECT2K8R2	Firefox	3	19	GET	0	500
Nov 19, 2015 12:11:10											
2015-11-18 12:17:13.06	fe80::f973:a13d:92d0:244c	nt authority\iusr	http://msproject2k8r2:11137/koyako~~~eg	-	MSPROJECT2K8R2	Firefox	3	19	GET	0	500
Nov 18, 2015 10:08:51											
2015-11-17 05:10:58.043	fe80::f973:a13d:92d0:244c	msproject2k8r2\administrator	http://msproject2k8r2:11137/_vti_bin/client.svc/processquery	http://msproject2k8r2:11137/SitePages/Home.aspx	MSPROJECT2K8R2	Firefox	0	0	POST	0	500
Nov 18, 2015 08:31:39											
2015-11-17 05:10:58.043	fe80::f973:a13d:92d0:244c	msproject2k8r2\administrator	http://msproject2k8r2:11137/_vti_bin/client.svc/processquery	http://msproject2k8r2:11137/SitePages/Home.aspx	MSPROJECT2K8R2	Firefox	0	0	POST	0	500

Figure 18.46: The detailed diagnosis of the 500 errors measure

18.2.6.1.1 Configuring the eG Agent to Collect Usage Analytics

SharePoint Usage and Health Service application is a feature to analyze usage of SharePoint environment or troubleshooting SharePoint Issues. This application can be configured to collect two types of data: **Usage data** and **Health data**. The following tests use the **Usage data** collected by the application to report metrics:

- Site Usage Analytics test
- Web Application Usage Analytics test
- Browser Usage Analytics test
- User Analytics test
- Distributed Cache Usage Analytics test

Usage data is about usage on SharePoint Farm, like page requests, feature use, search query latency, etc. This data is similar to IIS log, however unlike IIS logs this has additional SharePoint specific data collected like Application ID, Site ID, Web ID, Correlation ID etc. Usage data is initially stored in Usage Log file (.USAGE) on SharePoint Server under logging directory which is later processed by Microsoft SharePoint Foundation Usage Data Import Timer job into **Usage Database**. Each of the tests above query the **Usage Database** at configured intervals to collect the metrics they require.

For these tests to run, the following pre-requisites should be fulfilled:

- A **SharePoint Usage and Health Service Application** should be created on the target SharePoint server;
- Usage and Health data collection should be enabled for this application
- The eG agent on the SharePoint server should be allowed to query the Usage database.

Each of these steps are detailed below.

Creating a SharePoint Usage and Health Service Application

For this, do the following:

1. Login to the target SharePoint server as a user with **Farm administrator** privileges.
2. Open the SharePoint management shell.
3. Then, run the following commands one after another:

```
Add-PSSnapin Microsoft.SharePoint.PowerShell
New-SPUsageApplication -Name "<Name_of_application>"
```

For example, your command can be:

```
Add-PSSnapin Microsoft.SharePoint.PowerShell
New-SPUsageApplication -Name "SharePointUsageApp"
```

4. If the command executes successfully, then your output will reveal the name of the application you created, the application type, and the application ID.

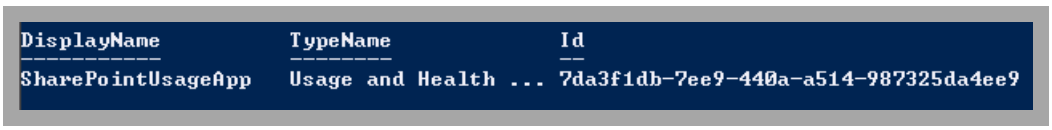

Figure 18.47: Output of the command issued for creating a SharePoint Usage and Health application

5. Next, open the SharePoint management console and follow the node sequence, *Central Administration -> Application Management -> Manage service applications*, on the console. Figure 18.48 will then appear. Click the **Manage service applications** option under **Service Applications** in Figure 18.48.

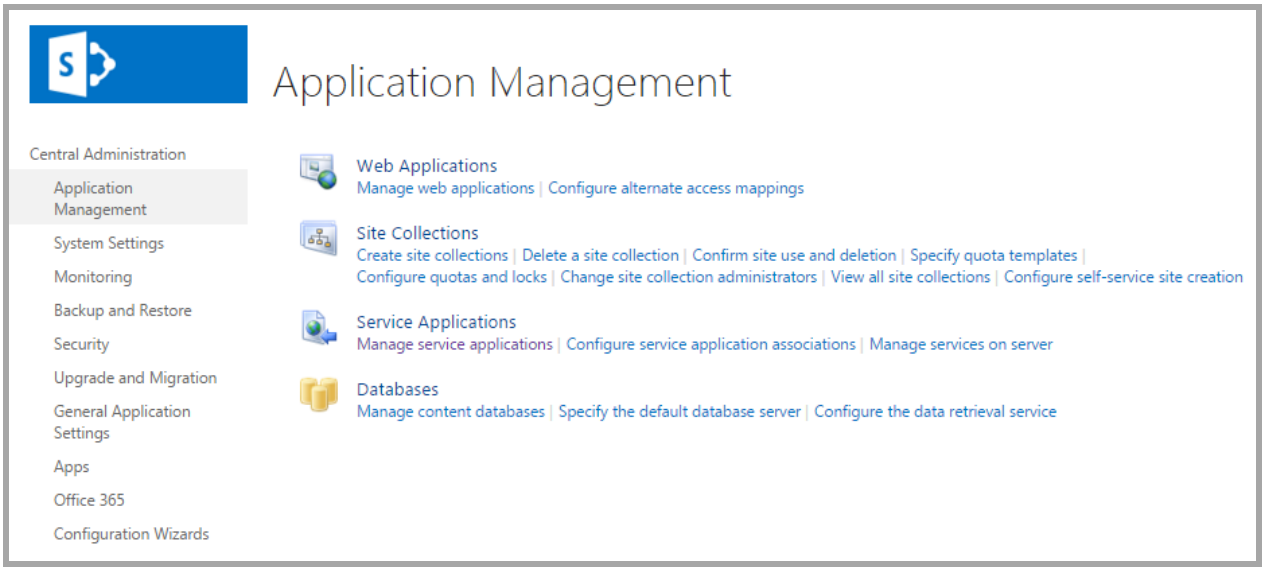

Figure 18.48: Selecting the Manage service applications option

6. When Figure 18.49 appears, look for an entry for the new Usage and Health application you created at step 3 above. If its available therein, it is a clear indicator that the new application has been created successfully.

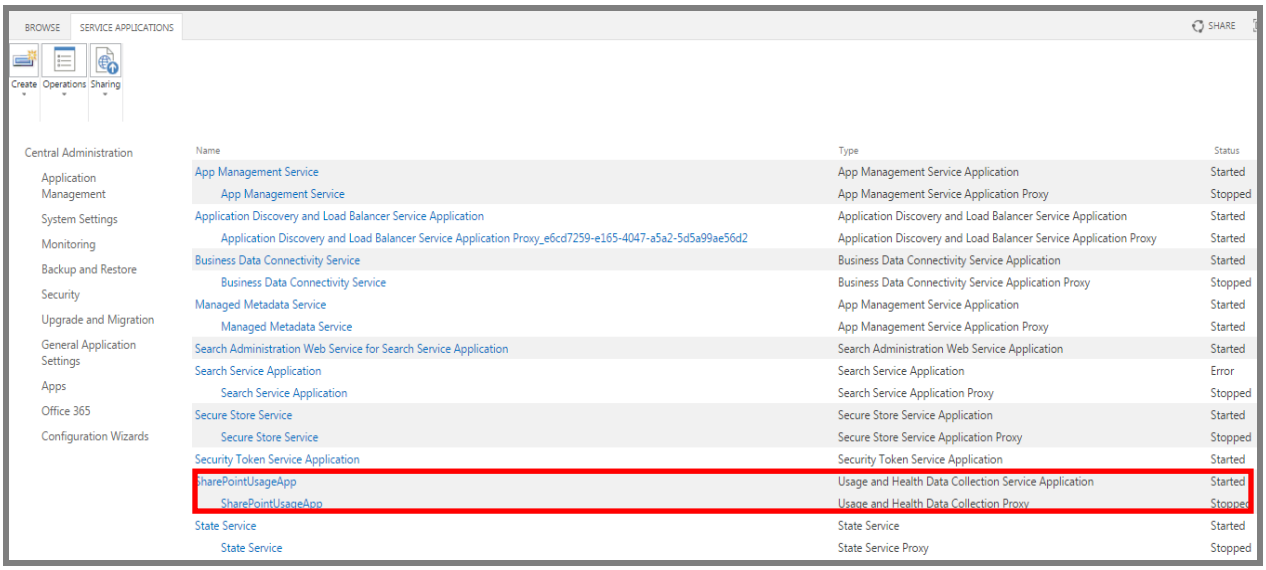

Figure 18.49: Looking for an entry for the new Usage and Health application you created

Enabling Usage and Health Data Collection

To achieve this, follow the steps below:

1. In the SharePoint management console, select the **Monitoring** node under **Central Administration.** Then, click on the **Configure usage and health data collection** option under **Reporting** (see Figure 18.50).

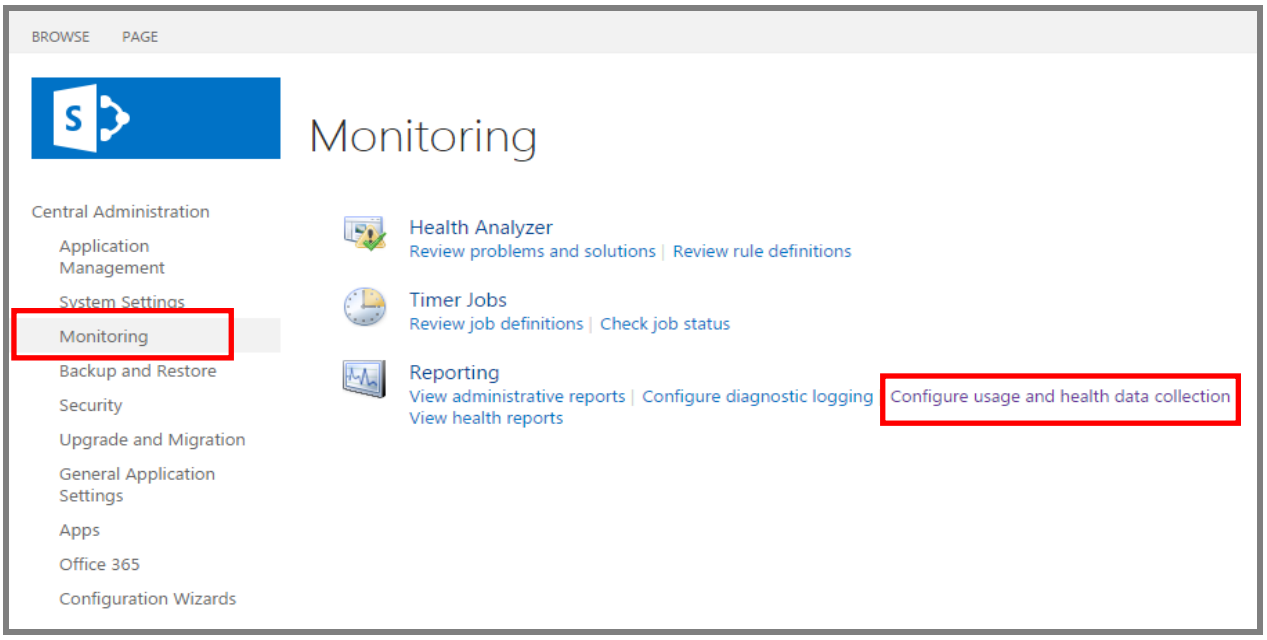

Figure 18.50: Selecting the Configure usage and health data collection option

2. When Figure 18.51 appears, select the **Enable usage data collection** checkbox therein.

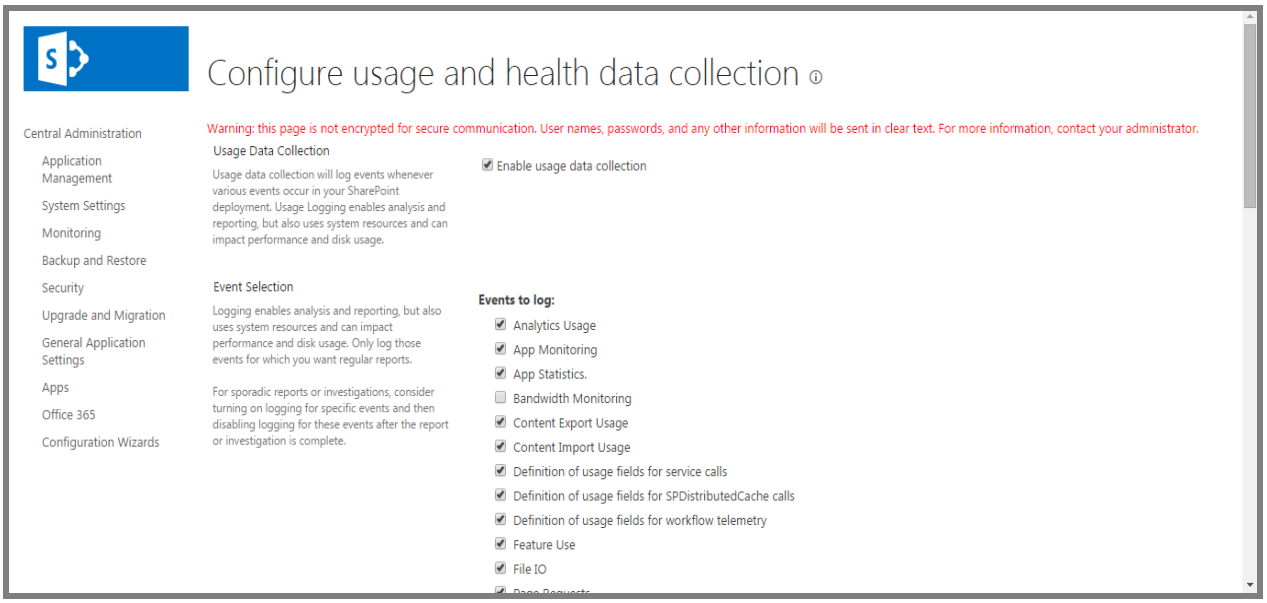

Figure 18.51: Enabling usage data collection

3. Under the **Events to log** section, make sure that the default selections (indicated by Figure 18.52 and Figure 18.53) are not disturbed.

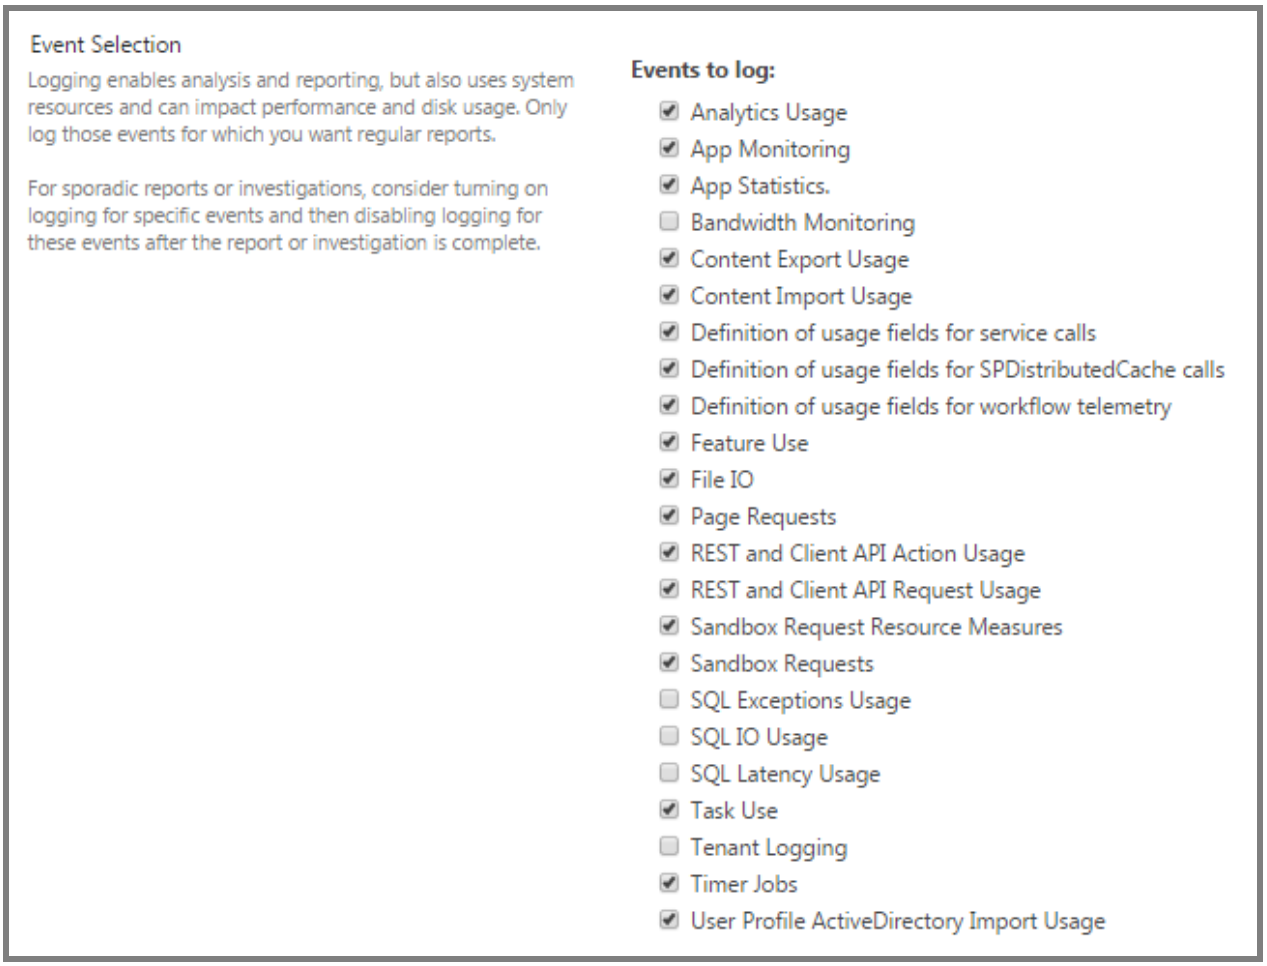

Figure 18.52: Retaining the default events to be logged

4. Scroll down Figure 18.13 and then select the **Enable health data collection** check box that becomes visible.

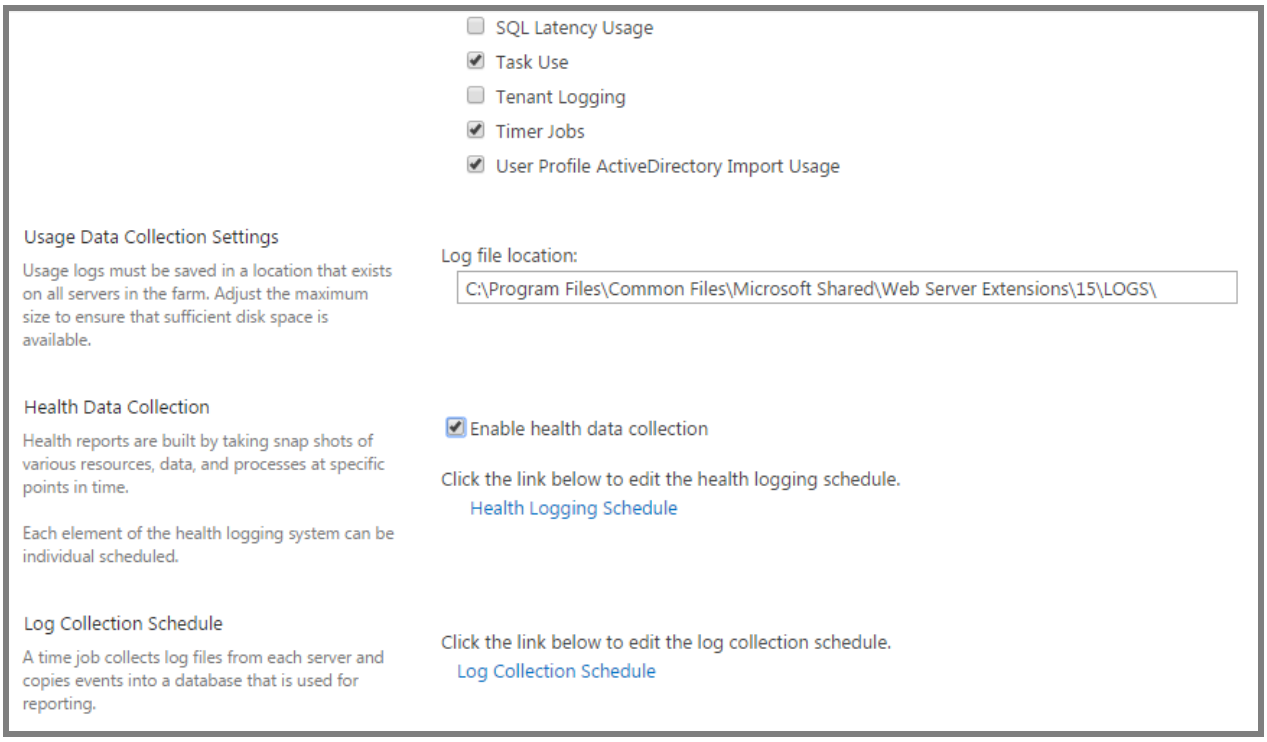

Figure 18.53: Enabling health data collection

5. Scroll further down Figure 18.53 until the **Database Server** and the **Database Name** fields visible (see Figure 18.54). Copy the values of these fields to a text editor. Make sure that the **DATABASE SERVER** and **DATABASE NAME** parameters of the analytics tests are configured with the copied values only.

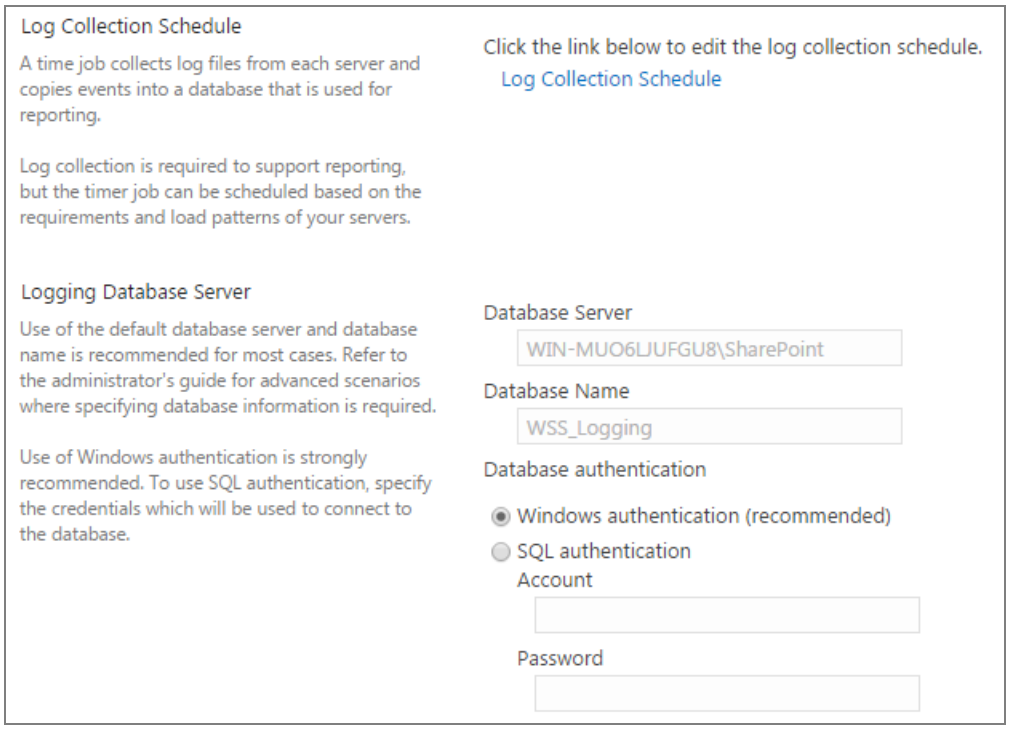

Figure 18.54: The name of the SQL server hosting the usage database and the name of the usage database

18.2.6.2 Web Application Usage Analytics Test

Enterprises typically use SharePoint to create web sites and web applications. The success of the SharePoint platform therefore hinges on how happy users are when interacting with the web applications that it helped create. If users of a web application constantly complain of slowness when browsing that web application, it indicates that user experience with the web application is sub-par. This in turn can hit user productivity badly, escalate troubleshooting time and costs of the enterprise, and adversely impact its revenues and reputation! To improve user experience with web applications and to build user confidence in the SharePoint platform, administrators should be able to quickly identify slow web applications and precisely pinpoint the reason for the slowness.

This is where the **Web Application Usage Analytics** test helps! This test queries the SharePoint usage database at configured intervals and collects metrics on web application usage that is stored therein – this includes the web applications accessed, count of users of each web application, the browsers that were used for web application access, web pages requested, the time taken for the requested pages to load, where page views spent time and how much, error responses returned, resources consumed, and many more. For each web application configured for monitoring, the test then reports the average time taken by that application to load pages. In the process, the test points administrators to slow web applications and also leads them to the probable source of the slowness – is it owing to a latent web front end? is it because of slow service calls? Or is it due to inefficient queries to the backend database?

Sometimes, poor user experience can be attributed to HTTP errors. This is why, this test instantly alerts administrators to HTTP error responses, thus ensuring their timely intervention and rapid resolution of the error conditions.

This way, the **Web Application Usage Analytics** test enables administrators to detect web application slowness well before users notice, helps them promptly and accurately diagnose the source of the poor user experience with a web

application, and thus ensures that they initiate measures to enhance user experience and pre-empt the damage that may be caused to revenue and reputation.

Note that this test will run only if a SharePoint Usage and Health Service application is created and is configured to collect usage and health data. To know how to create and configure this application, follow the steps detailed in Section 18.2.6.1.1.

Purpose	Enables administrators to detect web application slowness well before users notice, helps them promptly and accurately diagnose the source of the poor user experience with a web application, and thus ensures that they initiate measures to enhance user experience and pre-empt the damage that may be caused to revenue and reputation
Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **SQL PORT NUMBER** – Specify the port number of the Microsoft SQL server that is hosting the usage database. 5. **INSTANCE** – If the SQL server hosting the usage database is instance-based, then provide the instance name here. If not, then set this to *none*. 6. **SSL** – If the SQL server hosting the usage database is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**. 7. **ISNTLMV2** - In some Windows networks, **NTLM** (**NT LAN Manager**) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the **ISNTLMV2** flag is set to **No**, indicating that NTLMv2 is not enabled by default on the SQL server hosting the usage database. Set this flag to **Yes** if NTLMv2 is enabled on that SQL server. 8. **DATABASE SERVER NAME** – Specify the name of Microsoft SQL server that hosts the usage database to be accessed by this test. 9. **DATABASE NAME** – Specify the name of the usage database that this test should access. 10. **DATABASE USER NAME, DATABASE PASSWORD, CONFIRM PASSWORD** - Specify the credentials of a user who has read-only access to the configured usage, in the **DATABASE USER NAME** and **DATABASE PASSWORD** text boxes. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box. 11. **SLOW TRANSACTION CUTOFF (MS)** - This test reports the count of slow page views and also pinpoints the pages that are slow. To determine whether/not a page is slow, this test uses the **SLOW TRANSACTION CUTOFF** parameter. By default, this parameter is set to *4000 millisecs* (i.e., 4 seconds). This means that, if a page takes more than 4 seconds to load, this test will consider that page as a slow page by default. You can increase or decrease this slow transaction cutoff according to what is 'slow' and what is 'normal' in your environment.  Note The default value of this parameter is the same as the default *Maximum threshold* setting of the *Avg page load time* measure – i.e., both are set to *4000 millisecs* by default. While the former helps eG to distinguish between slow and healthy page views for the purpose of providing detailed diagnosis, the latter tells eG when to generate an alarm on *Avg page load time*. For best results, it is recommended that both these settings are configured with the same value at all times. Therefore, if you change the value of one of these configurations, then make sure you update the value of the other as well. For instance, if the **SLOW TRASACTION CUTOFF** is changed to *6000 millisecs*, change the *Maximum Threshold* of the *Avg page load time* measure to *6000* millsecs as well.

<table>
<tr><td></td><td colspan="3">12. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each web application on SharePoint</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><td></td><td>Unique users:
Indicates the number of unique users of this web application.</td><td>Number</td><td>The detailed diagnosis of this measure reveals the names of the unique users and the number of requests from each user to the web application being monitored. From this, you can identify those users who are actively using the web application.</td></tr>
<tr><td></td><td>Unique visitors:
Indicates the number of unique visitors to this web application.</td><td>Number</td><td>SharePoint authenticated users and anonymous users (using IP address) are counted as visitors.

Compare the value of this measure across web applications to identify the most popular one.

You can use the detailed diagnosis of this measure to know who are the unique visitors to the web application and the number of requests from each visitor to the web application This way, you can identify that visitor who visits the web application most frequently.</td></tr>
<tr><td></td><td>Unique destinations:
Indicates the number of unique destinations of this web application.</td><td>Number</td><td>To know the most popular destination URLs of this web application, use the detailed diagnosis of this measure. Here, you will find the top-10 destinations in terms of the number of hits.</td></tr>
</table>

	Unique browsers: Indicates the number of unique browsers used for accessing this web application.	Number	To know which browsers are commonly used to access this web application, use the detailed diagnosis of this measure. Here, the unique browsers will be listed and the number of hits to the web application from each browser will be displayed alongside, so that you can instantly identify that browser that has been widely used to access the web application.
	Unique referrers: Indicates the number of unique URLs external to this web application (parent web application is treated as external as well), from where the users navigated to this web application.	Number	To know which referrer URL was responsible for the maximum hits to this web application, use the detailed diagnosis of this measure. The top-10 unique referrer URLs in terms of the number of hits they generated will be displayed as part of the detailed diagnostics.

	Apdex score: Indicates the apdex score of this web application.	Number	**Apdex (Application Performance Index)** is an open standard developed by an alliance of companies. It defines a standard method for reporting and comparing the performance of software applications in computing. Its purpose is to convert measurements into insights about user satisfaction, by specifying a uniform way to analyze and report on the degree to which measured performance meets user expectations. The Apdex method converts many measurements into one number on a uniform scale of 0-to-1 (0 = no users satisfied, 1 = all users satisfied). The resulting Apdex score is a numerical measure of user satisfaction with the performance of enterprise applications. This metric can be used to report on any source of end-user performance measurements for which a performance objective has been defined. The Apdex formula is: $Apdex_t = (Satisfied\ Count + Tolerating\ Count / 2) / Total\ Samples$ This is nothing but the number of satisfied samples plus half of the tolerating samples plus none of the frustrated samples, divided by all the samples. A score of 1.0 means all responses were satisfactory. A score of 0.0 means none of the responses were satisfactory. Tolerating responses half satisfy a user. For example, if all responses are tolerating, then the Apdex score would be 0.50. Ideally therefore, the value of this measure should be *1.0*. A value less than *1.0* indicates that the user experience with the web application has been less than satisfactory.

	Satisfied page views: Indicates the number of times pages in this web application were viewed without any slowness.	Number	A page view is considered to be slow when the average time taken to load that page exceeds the **SLOW TRANSACTION CUTOFF** configured for this test. If this **SLOW TRANSACTION CUTOFF** is not exceeded, then the page view is deemed to be 'satisfactory'. Ideally, the value of this measure should be high. If the value of this measure is much lesser than the value of the *Tolerating page views* and the *Frustrated page views*, it is a clear indicator that the experience of the users of this web application is below-par. In such a case, use the detailed diagnosis of the *Tolerating page views* and *Frustrated page views* measures to know which pages are slow.
	Tolerating page views: Indicates the number of tolerating page views to this web application.	Number	If the *Average page load time* of a page exceeds the **SLOW TRANSACTION CUTTOFF** configuration of this test, but is less than 4 times the **SLOW TRANSACTION CUTOFF** (i.e., < 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Tolerating page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the overall user experience from this browser is less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.
	Frustrated page views: Indicates the number of frustrated page views to this web application.	Number	If the *Average page load time* of a page is over 4 times the **SLOW TRANSACTION CUTTOFF** configuration of this test (i.e., > 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Frustrated page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the experience of users using this browser has been less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.

	Average page load time: Indicates the average time taken by the pages in this web application to load completely.	Secs	This is the average interval between the time that a user initiates a request and the completion of the page load of the response in the user's browser. If the value of this measure is consistently high for a web application, there is reason to worry. This is because, it implies that the web application is slow in responding to requests. If this condition is allowed to persist, it can adversely impact user experience with the web application. You may want to check the *Apdex score* in such circumstances to determine whether/not user experience has already been affected. Regardless, you should investigate the anomaly and quickly determine where the bottleneck lies – is it with the web front-end? is it owing to slow service calls? Or is it because of inefficient queries to the backend? - so that the problem can be fixed before users even notice any slowness! For that, you may want to compare the values of the *Web front-end processing time*, *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test. If the *Average front end time* is the highest, it indicates that the problem is with the web site/web application front end – this can be attributed to a slowdown in page rendering or in DOM building. If the *Average server connection time* is the highest, it denotes that the network is the problem source. This in turn can be caused by TCP connection latencies and delays in domain look up. On the other hand, if the *Average response available time* measure registers the highest value, it indicates that the problem lies with the web site/web application backend – i.e., the web/web application server that is hosting the web site/web application being monitored.

	Web front-end processing time: Indicates the average time in milliseconds it took for the web front end server to process the requests to this web application.	Msecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average service calls duration: Indicates the time taken by this web application to generate service calls.	Msecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average IIS latency: Indicates the average time requests to this web application took in the frontend web server after the requests were received by the frontend web server but before this web application began processing the requests.	MSecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average CPU duration*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average CPU duration: Indicates the average time for which requests to this web application used the CPU.	Msecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	SQL logical reads: Indicates the total number of 8 kilobyte blocks that this web application read from storage on the back-end database server.	Number	

	Average CPU megacycles: Indicates the average number of CPU mega cycles spent processing the requests to this web application in the client application on the front end web server.	Number	
	Total queries: Indicates the total number of database queries generated for this web application.	Number	
	Average query duration: Indicates the average time taken for all backend database queries generated for this web application.	Msecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average IIS latency*, and *Average CPU duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average data consumed: Indicates the average bytes of data downloaded by requests to this web application.	KB	
	GET requests: Indicates the number of GET requests to this web application.	Number	
	POST requests: Indicates the number of POST requests to this web application.	Number	
	OPTION requests: Indicates the number of OPTION request to this web application.	Number	

	300 responses: Indicates the number of responses to requests to this web application with a status code in the 300-399 range	Number	300 responses could indicate page caching on the client browsers. Alternatively 300 responses could also indicate redirection of requests. A sudden change in this value could indicate a problem condition.
	400 errors: Indicates the number responses to requests to this web application that had a status code in the range 400-499.	Number	A high value indicates a number of missing/error pages. Use the detailed diagnosis of this measure to know when each of the 400 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.
	500 errors: Indicates the number of responses to the requests to this web application that had a status code in the range 500-599.	Number	Since responses with a status code of 500-600 indicate server side processing errors, a high value reflects an error condition. Use the detailed diagnosis of this measure to know when each of the 500 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.

18.2.6.3 Browser Usage Analytics Test

Different users use different browsers to access and browse the web sites and web applications created on SharePoint. Very often, user experience with a web site/application can vary with the browser being used! Using an obsolete or an unsupported browser can cause users to see errors or serious performance degradations when accessing web sites or mission-critical web application. This in turn can delay critical business operations, impair user productivity, and basically, be the reason for enterprises to incur huge penalities, mounting costs, and heavy losses! What administrators need to do therefore is to identify what browsers are being used by their users, see for themselves whether/not user experience changes with browser, and in the process, isolate those browsers that could be delivering a sub-par experience to their users.

This is where the **Browser Usage Analytics** test helps! This test queries the SharePoint usage database at configured intervals and collects metrics on browser usage that is stored therein. For each browser used, the test then reports the average time taken by that browser to load pages. In the process, the test points administrators to slow browsers and also leads them to the probable source of the slowness - is it owing to a latent web front-end? Is it because of slow service calls? Or is it due to inefficient queries to the backend database?

The test also captures HTTP errors that occurred when using each browser, thus enabling administrators to quickly detect browser-related issues and rapidly fix them before user experience is impacted.

This way, the **Browser Usage Analytics** test enables administrators to identify problematic browsers, helps them to try and enhance the experience of users using such browsers, or at least conclude which browsers are not ideal for usage with which web sites/web applications.

Note that this test will run only if a SharePoint Usage and Health Service application is created and is configured to collect usage and health data. To know how to create and configure this application, follow the steps detailed in Section 18.2.6.1.1.

Purpose	Enables administrators to identify problematic browsers, helps them to try and enhance the experience of users using such browsers, or at least conclude which browsers are not ideal for usage with which web sites/web applications
Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **SQL PORT NUMBER** – Specify the port number of the SQL server that hosts the usage database. 5. **INSTANCE** – If the SQL server that hosts the usage database is instance-based, then provide the instance name here. If not, then set this to *none*. 6. **SSL** – If the SQL server hosting the usage database is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**. 7. **ISNTLMV2** - In some Windows networks, **NTLM** (**NT LAN Manager**) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the **ISNTLMV2** flag is set to **No**, indicating that NTLMv2 is not enabled by default on the SQL server that hosts the usage database. Set this flag to **Yes** if NTLMv2 is enabled on that SQL server. 8. **DATABASE SERVER NAME** – Specify the name of Microsoft SQL server that hosts the usage database to be accessed by this test. 9. **DATABASE NAME** – Specify the name of the usage database that this test should access. 10. **DATABASE USER NAME, DATABASE PASSWORD, CONFIRM PASSWORD** - Specify the credentials of a user who has read-only access to the usage database configured, in the **DATABASE USER NAME** and **DATABASE PASSWORD** text boxes. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box. 11. **SLOW TRANSACTION CUTOFF (MS)** - This test reports the count of slow page views and also pinpoints the pages that are slow. To determine whether/not a page is slow, this test uses the **SLOW TRANSACTION CUTOFF** parameter. By default, this parameter is set to *4000 millisecs* (i.e., 4 seconds). This means that, if a page takes more than 4 seconds to load, this test will consider that page as a slow page by default. You can increase or decrease this slow transaction cutoff according to what is 'slow' and what is 'normal' in your environment.  Note The default value of this parameter is the same as the default *Maximum threshold* setting of the *Avg page load time* measure – i.e., both are set to *4000 millisecs* by default. While the former helps eG to distinguish between slow and healthy page views for the purpose of providing detailed diagnosis, the latter tells eG when to generate an alarm on *Avg page load time*. For best results, it is recommended that both these settings are configured with the same value at all times. Therefore, if you change the value of one of these configurations, then make sure you update the value of the other as well. For instance, if the **SLOW TRASACTION CUTOFF** is changed to *6000 millisecs*, change the *Maximum Threshold* of the *Avg page load time* measure to *6000* millsecs as well.

	12. **UNKNOWN BROWSERS** – By default, this flag is set to **No**. This means, by default, eG will monitor only those browsers that SharePoint can recognize. If users use browsers that SharePoint cannot recognize, then, usage analytics of such browsers will be grouped under the *Unknown* browser type in the **Usage database**. If you want to view metrics related to the *Unknown* browser type as well, then set this flag to **Yes**. In this case, *Unknown* will be displayed as an additional descriptor of this test. 13. **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.		
Outputs of the test	One set of results for each browser using which users are accessing SharePoint web applications		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Unique users: Indicates the number of unique users of this web application.	Number	Compare the value of this measure across browsers to identify the most popular one. The detailed diagnosis of this measure reveals the names of the unique users and the number of requests from each user to the browser. From this, you can identify those users who are actively using the browser.
	Unique visitors: Indicates the number of unique visitors using this browser.	Number	SharePoint authenticated users and anonymous users (using IP address) are counted as visitors. You can use the detailed diagnosis of this measure to know who are the unique visitors to the browser and the number of requests from each visitor to the browser. This way, you can identify that visitor who uses the browser most frequently.
	Unique destinations: Indicates the number of unique destinations of this browser.	Number	To know the most popular destination URLs, use the detailed diagnosis of this measure. Here, you will find the top-10 destinations in terms of the number of hits.

	Unique referrers: Indicates the number of unique URLs external to this browser (parent web application is treated as external as well), from where the users navigated to this browser.	Number	To know which referrer URL was responsible for the maximum hits, use the detailed diagnosis of this measure. The top-10 unique referrer URLs in terms of the number of hits they generated will be displayed as part of the detailed diagnostics.
	Apdex score: Indicates the apdex score of this browser.	Number	**Apdex (Application Performance Index)** is an open standard developed by an alliance of companies. It defines a standard method for reporting and comparing the performance of software applications in computing. Its purpose is to convert measurements into insights about user satisfaction, by specifying a uniform way to analyze and report on the degree to which measured performance meets user expectations. The Apdex method converts many measurements into one number on a uniform scale of 0-to-1 (0 = no users satisfied, 1 = all users satisfied). The resulting Apdex score is a numerical measure of user satisfaction with the performance of enterprise applications. This metric can be used to report on any source of end-user performance measurements for which a performance objective has been defined. The Apdex formula is: $Apdex_t = (Satisfied\ Count + Tolerating\ Count / 2) / Total\ Samples$ This is nothing but the number of satisfied samples plus half of the tolerating samples plus none of the frustrated samples, divided by all the samples. A score of 1.0 means all responses were satisfactory. A score of 0.0 means none of the responses were satisfactory. Tolerating responses half satisfy a user. For example, if all responses are tolerating, then the Apdex score would be 0.50. Ideally therefore, the value of this measure should be *1.0*. A value less than *1.0* indicates that the user experience with the browser has been less than satisfactory.

	Satisfied page views: Indicates the number of times pages were viewed in this browser without any slowness.	Number	A page view is considered to be slow when the average time taken to load that page exceeds the **SLOW TRANSACTION CUTOFF** configured for this test. If this **SLOW TRANSACTION CUTOFF** is not exceeded, then the page view is deemed to be 'satisfactory'. Ideally, the value of this measure should be high. If the value of this measure is much lesser than the value of the *Tolerating page views* and the *Frustrated page views*, it is a clear indicator that the experience of the users of this browser is below-par. In such a case, use the detailed diagnosis of the *Tolerating page views* and *Frustrated page views* measures to know which pages are slow.
	Tolerating page views: Indicates the number of tolerating page views in this browser.	Number	If the *Average page load time* of a page exceeds the **SLOW TRANSACTION CUTTOFF** configuration of this test, but is less than 4 times the **SLOW TRANSACTION CUTOFF** (i.e., < 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Tolerating page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the overall user experience from this browser is less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.
	Frustrated page views: Indicates the number of frustrated page views in this browser.	Number	If the *Average page load time* of a page is over 4 times the **SLOW TRANSACTION CUTTOFF** configuration of this test (i.e., > 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Frustrated page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the experience of users using this browser has been less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.

	Average page load time: Indicates the average time taken by the pages to load completely in this browser.	Secs	This is the average interval between the time that a user initiates a request and the completion of the page load of the response in the user's browser. If the value of this measure is consistently high for a browser, there is reason to worry. This is because, it implies that the browser is slow in responding to requests. If this condition is allowed to persist, it can adversely impact user experience. You may want to check the *Apdex score* in such circumstances to determine whether/not user experience has already been affected. Regardless, you should investigate the anomaly and quickly determine where the bottleneck lies – is it with the browser itself? is it with the web front-end? is it owing to slow service calls? Or is it because of inefficient queries to the backend? - so that the problem can be fixed before users even notice any slowness! For that, you may want to compare the values of the *Web front-end processing time*, *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test.
	Web front-end processing time: Indicates the average time in milliseconds it took for the web front end server to process the requests to this browser.	Msecs	If the *Avg page load time* of a browser is abnormally high, then you can compare the value of this measure with that of the *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average service calls duration: Indicates the time taken by this browser to generate service calls.	Msecs	If the *Avg page load time* of a browser is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?

	Average IIS latency: Indicates the average time requests to this browser took in the frontend web server after the requests were received by the frontend web server but before this browser began processing the requests.	MSecs	If the *Avg page load time* of a browser is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average CPU duration*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average CPU duration: Indicates the average time for which requests to this browser used the CPU.	Msecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	SQL logical reads: Indicates the total number of 8 kilobyte blocks that this browser read from storage on the back-end database server.	Number	
	Average CPU megacycles: Indicates the average number of CPU mega cycles spent processing the requests to this browser in the client application on the front end web server.	Number	
	Total queries: Indicates the total number of database queries generated by requests to this browser.	Number	
	Average query duration: Indicates the average time taken for all backend database queries generated by requests to this browser.	Msecs	If the *Avg page load time* of a browser is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average IIS latency*, and *Average CPU duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?

	Average data consumed: Indicates the average bytes of data downloaded by requests to this browser.	KB	
	GET requests: Indicates the number of GET requests to this web browser.	Number	
	POST requests: Indicates the number of POST requests to this web browser.	Number	
	OPTION requests: Indicates the number of OPTION requests to this browser.	Number	
	300 responses: Indicates the number of responses for requests to this browser that had a status code in the 300-399 range.	Number	300 responses could indicate page caching on the client browsers. Alternatively 300 responses could also indicate redirection of requests. A sudden change in this value could indicate a problem condition.
	400 errors: Indicates the number responses for requests to this browser that had a status code in the range 400-499.	Number	A high value indicates a number of missing/error pages. Use the detailed diagnosis of this measure to know when each of the 400 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.

	500 errors: Indicates the number of responses for requests to this browser that had a status code in the range 500-599.	Number	Since responses with a status code of 500-600 indicate server side processing errors, a high value reflects an error condition. Use the detailed diagnosis of this measure to know when each of the 500 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.

18.2.6.4 User Analytics Test

Enterprises typically use SharePoint to create web sites and web applications. The success of the SharePoint platform therefore hinges on the level of user satisfaction with the web sites and applications created on that platform. The key to ensuring high user satistaction lies in closely tracking user requests to the web sites/web applications on SharePoint, measuring the responsiveness of the web sites/web applications to the user requests, instantly detecting poor responsiveness, and accurately isolating which user's experience is being impacted by this slowness, well before that user notices! This can be achieved using the **User Analytics** test!

This test queries the SharePoint usage database at configured intervals and collects usage metrics that are stored therein – this includes the web sites/web applications accessed, count and names of users of each web site/web application, the browsers that were used for web site/web application access, web pages requested, the time taken for the requested pages to load, where page views spent time and how much, error responses returned, resources consumed, and many more. Using the query results, the test then auto-discovers the users accessing each of the web sites/web applications that are configured for monitoring. Then, for each such user, this test reports the average time taken by the corresponding site/web application to load pages. In the process, the test points administrators to slow web sites/web applications, reveals the exact user who has suffered the most owing to this slowness, and also leads them to the probable source of the slowness – is it owing to a latent web front end? is it because of slow service calls? Or is it due to inefficient queries to the backend database?

Sometimes, poor user experience can be attributed to HTTP errors. This is why, this test instantly alerts administrators to HTTP error responses, thus ensuring their timely intervention and rapid resolution of the error conditions.

This way, the **User Analytics** test enables administrators to proactively detect users who are experiencing or who will potentially experience performance issues with a web site/web application, helps them promptly and accurately diagnose the source of the poor user experience, and thus ensures that they initiate measures to enhance user experience and pre-empt the damage that may be caused to revenue and reputation.

Note that this test will run only if a SharePoint Usage and Health Service application is created and is configured to collect usage and health data. To know how to create and configure this application, follow the steps detailed in Section 18.2.6.1.1.

Purpose	Enables administrators to proactively detect users who are experiencing or who will potentially experience performance issues with a web site/web application, helps them promptly and accurately diagnose the source of the poor user experience, and thus ensures that they initiate measures to enhance user experience and pre-empt the damage that may be caused to revenue and reputation

Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **SQL PORT NUMBER** – Specify the port number of the SQL server that hosts the usage database. 5. **INSTANCE** – If the SQL server hosting the usage database is instance-based, then provide the instance name here. If not, then set this to *none*. 6. **SSL** – If the SQL server hosting the usage database is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**. 7. **ISNTLMV2** - In some Windows networks, **NTLM** (**NT LAN Manager**) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the **ISNTLMV2** flag is set to **No**, indicating that NTLMv2 is not enabled by default on the SQL server hosting the usage database. Set this flag to **Yes** if NTLMv2 is enabled on that SQL server. 8. **DATABASE SERVER NAME** – Specify the name of Microsoft SQL server that hosts the usage database to be accessed by this test. 9. **DATABASE NAME** – Specify the name of the usage database that this test should access. 10. **DATABASE USER NAME, DATABASE PASSWORD, CONFIRM PASSWORD** - Specify the credentials of a user who has read-only access to the usage database configured, in the **DATABASE USER NAME** and **DATABASE PASSWORD** text boxes. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box. 11. **SLOW TRANSACTION CUTOFF (MS)** - This test reports the count of slow page views and also pinpoints the pages that are slow. To determine whether/not a page is slow, this test uses the **SLOW TRANSACTION CUTOFF** parameter. By default, this parameter is set to *4000 millisecs* (i.e., 4 seconds). This means that, if a page takes more than 4 seconds to load, this test will consider that page as a slow page by default. You can increase or decrease this slow transaction cutoff according to what is 'slow' and what is 'normal' in your environment.  Note The default value of this parameter is the same as the default *Maximum threshold* setting of the *Avg page load time* measure – i.e., both are set to *4000 millisecs* by default. While the former helps eG to distinguish between slow and healthy page views for the purpose of providing detailed diagnosis, the latter tells eG when to generate an alarm on *Avg page load time*. For best results, it is recommended that both these settings are configured with the same value at all times. Therefore, if you change the value of one of these configurations, then make sure you update the value of the other as well. For instance, if the **SLOW TRASACTION CUTOFF** is changed to *6000 millisecs*, change the *Maximum Threshold* of the *Avg page load time* measure to *6000* millsecs as well.

	12. **SITE** – Configure a comma-separated list of web site URLs that you want this test to monitor. For eg., *http://www.msproject28rk2:11982,http://www.mydocs.com* 13. **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.		
Outputs of the test	One set of results for each user accessing every SharePoint **SITE** configured for monitoring First-level descriptor: Site URL Second-level descriptor: User name		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Unique sessions: Indicates the number of unique sessions for this user on this web site.	Number	Compare the value of this measure across users to identify the user who has the maximum number of open sessions on the site, and is hence, probably overloading the site. The detailed diagnosis of this measure reveals the unique client IP addresses from which the user launched his/her sessions and the number of requests received from each IP address.
	Unique destinations: Indicates the number of unique destinations for this site for this user.	Number	To know the most popular destination URLs for a user, use the detailed diagnosis of this measure. Here, you will find the top-10 destinations in terms of the number of hits.
	Unique referrers: Indicates the number of unique URLs external to this site (parent web application is treated as external as well), from where this user navigated to the browser.	Number	To know which referrer URL was responsible for the maximum hits, use the detailed diagnosis of this measure. The top-10 unique referrer URLs in terms of the number of hits they generated will be displayed as part of the detailed diagnostics.

	Apdex score: Indicates the apdex score of this user for this site.	Number	**Apdex (Application Performance Index)** is an open standard developed by an alliance of companies. It defines a standard method for reporting and comparing the performance of software applications in computing. Its purpose is to convert measurements into insights about user satisfaction, by specifying a uniform way to analyze and report on the degree to which measured performance meets user expectations. The Apdex method converts many measurements into one number on a uniform scale of 0-to-1 (0 = no users satisfied, 1 = all users satisfied). The resulting Apdex score is a numerical measure of user satisfaction with the performance of enterprise applications. This metric can be used to report on any source of end-user performance measurements for which a performance objective has been defined. The Apdex formula is: $Apdex_t = (Satisfied\ Count + Tolerating\ Count / 2) / Total\ Samples$ This is nothing but the number of satisfied samples plus half of the tolerating samples plus none of the frustrated samples, divided by all the samples. A score of 1.0 means all responses were satisfactory. A score of 0.0 means none of the responses were satisfactory. Tolerating responses half satisfy a user. For example, if all responses are tolerating, then the Apdex score would be 0.50. Ideally therefore, the value of this measure should be *1.0*. A value less than *1.0* indicates that this user's experience with the corresponding web site has been less than satisfactory.

	Satisfied page views: Indicates the number of times pages were viewed in this web site by this user without any slowness.	Number	A page view is considered to be slow when the average time taken to load that page exceeds the **SLOW TRANSACTION CUTOFF** configured for this test. If this **SLOW TRANSACTION CUTOFF** is not exceeded, then the page view is deemed to be 'satisfactory'. Ideally, the value of this measure should be high. If the value of this measure is much lesser than the value of the *Tolerating page views* and the *Frustrated page views*, it is a clear indicator that the experience of the user is below-par. In such a case, use the detailed diagnosis of the *Tolerating page views* and *Frustrated page views* measures to know which pages are slow.
	Tolerating page views: Indicates the number of tolerating page views for this user in this web site.	Number	If the *Average page load time* of a page exceeds the **SLOW TRANSACTION CUTTOFF** configuration of this test, but is less than 4 times the **SLOW TRANSACTION CUTOFF** (i.e., < 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Tolerating page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the user is less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.
	Frustrated page views: Indicates the number of frustrated page views for this user in this web site.	Number	If the *Average page load time* of a page is over 4 times the **SLOW TRANSACTION CUTTOFF** configuration of this test (i.e., > 4 * **SLOW TRANSACTION CUTOFF**), then such a page view is considered to be a **Frustrated page view**. Ideally, the value of this measure should be 0. A value higher than that of the *Satisfied page views* measure is a cause for concern, as it implies that the experience of the user has been less than satisfactory. To know which pages are contributing to this sub-par experience, use the detailed diagnosis of this measure.

	Average page load time: Indicates the average time taken by the pages in this site that are requested by this user to load completely.	Msecs	This is the average interval between the time that a user initiates a request and the completion of the page load of the response in the user's browser. If the value of this measure is consistently high for a user, it implies a degraded user experience. You may want to check the *Apdex score* in such circumstances to determine whether/not user experience has already been affected. Regardless, you should investigate the anomaly and quickly determine where the bottleneck lies – is it with the web front-end? is it owing to slow service calls? Or is it because of inefficient queries to the backend? - so that the problem can be fixed before users even notice any slowness! For that, you may want to compare the values of the *Web front-end processing time*, *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test.
	Web front-end processing time: Indicates the average time in milliseconds it took for the web front end server to process the requests of this user to this web site.	Msecs	If the *Avg page load time* of a user is abnormally high, then you can compare the value of this measure with that of the *Average service calls duration*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average service calls duration: Indicates the time taken by the requests of this user to this web site to generate service calls.	Msecs	If the *Avg page load time* of a user is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average CPU duration*, *Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow front-end web server? inefficient queries to the backend database? or slow service calls?

	Average IIS latency: Indicates the average time requests from this user took in the frontend web server after the requests were received by the frontend web server but before the browser began processing the requests.	MSecs	If the *Avg page load time* of a user is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average CPU duration*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow browser? A slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average CPU duration: Indicates the average time for which requests from this user to this site used the CPU.	Msecs	If the *Avg page load time* of a web application is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration, Average IIS latency*, and *Average query duration* measures of this test to know what exactly is delaying page loading – a slow browser? a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	SQL logical reads: Indicates the total number of 8 kilobyte blocks that this browser read from storage on the back-end database server.	Number	
	Average CPU megacycles: Indicates the average number of CPU mega cycles spent processing the requests to this browser in the client application on the front end web server.	Number	
	Total queries: Indicates the total number of database queries generated by requests to this browser.	Number	

	Average query duration: Indicates the average time taken for all backend database queries generated by requests from this user to this web site.	Msecs	If the *Avg page load time* of a browser is abnormally high, then you can compare the value of this measure with that of the *Web front-end processing time*, *Average service calls duration*, *Average IIS latency*, and *Average CPU duration* measures of this test to know what exactly is delaying page loading – a slow browser? a slow front-end web server? inefficient queries to the backend database? or slow service calls?
	Average data consumed: Indicates the average bytes of data downloaded by the requests of this user.	KB	
	GET requests: Indicates the number of GET requests from this user to this site.	Number	
	POST requests: Indicates the number of POST requests from this user to this site.	Number	
	OPTION requests: Indicates the number of OPTION requests from this user to this site.	Number	
	300 responses: Indicates the number of responses for requests from this user that had a status code in the 300-399 range.	Number	300 responses could indicate page caching on the client browsers. Alternatively 300 responses could also indicate redirection of requests. A sudden change in this value could indicate a problem condition.
	400 errors: Indicates the number responses for requests from this user that had a status code in the range 400-499.	Number	A high value indicates a number of missing/error pages. Use the detailed diagnosis of this measure to know when each of the 400 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.

	500 errors: Indicates the number of responses for requests from this user that had a status code in the range 500-599.	Number	Since responses with a status code of 500-600 indicate server side processing errors, a high value reflects an error condition. Use the detailed diagnosis of this measure to know when each of the 500 errors occurred, which user experienced the error, when using what browser, from which machine. This information will greatly aid troubleshooting.

18.2.6.5 Distributed Cache Usage Analytics Test

The Distributed Cache service, which is built on Windows Server AppFabric Cache, is set to run in a collocated mode on all SharePoint 2013 Servers by default. It's essential for maintaining the large amounts of information on your SharePoint Server, ensuring that the information is fresh and readily available for the end user.

Caching functionalities, provided by the Distributed Cache service, enable web applications deployed on SharePoint to quickly retrieve data without any dependency on databases stored in SQL Server, as everything is stored in memory.

Any SharePoint server in the farm running the Distributed Cache service is known as a cache host. Cache size is the memory allocated to the Distributed Cache service on the cache host.

At any given point in time, sufficient memory resources should be available to the Distributed cache service to ensure optimum cache usage and to assure SharePoint users of a satisfactory experience with their web applications. In the absence of adequate memory, cache lookups will be delayed or even missed, thus affecting overall SharePoint performance and adversely impacting the health of user interactions with the web applications.

It is hence imperative that administrators keep an eye on the usage of the cache service by each dependent web application, rapidly detect unexpected slowness in cache reads and writes, capture cache misses, and figure out if such anomalies are owing to the bad size of the Distributed cache service. This is what the **Distributed Cache Usage Analytics** test help administrators do!

This test queries the SharePoint usage database and retrieves metrics revealing how each web application uses the distributed cache, from it. The metrics so collected reveal the following:

- Is any web application reading from and/or writing to the cache slowly? If so, which host is slow?
- Is any web application overloading the cache with read/write requests?
- Which web application is experiencing many cache misses?
- Are there any cache failures? If so, which web application failed to read from or write to the cache?

This way, the test brings cache usage and sizing irregularities to light, pinpoints the exact web application that is being impacted by these abnormalities, and thus prompts administrators to right-size the cache to ensure peak application performance.

Note that this test will run only if a SharePoint Usage and Health Service application is created and is configured to collect usage and health data. To know how to create and configure this application, follow the steps detailed in Section 18.2.6.1.1.

Purpose	Brings cache usage and sizing irregularities to light, pinpoints the exact web application that is being impacted by these abnormalities, and thus prompts administrators to right-size the cache

	to ensure peak application performance
Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**. 4. **SQL PORT NUMBER** – Specify the port number of the SQL server hosting the usage database. 5. **INSTANCE** – If the SQL server hosting the usage database is instance-based, then provide the instance name here. If not, then set this to *none*. 6. **SSL** – If the SQL server hosting the usage database is SSL-enabled, then set this flag to **Yes**. If not, set it to **No**. 7. **ISNTLMV2** - In some Windows networks, **NTLM** (**NT LAN Manager**) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 ("NTLMv2") was concocted to address the security issues present in NTLM. By default, the **ISNTLMV2** flag is set to **No**, indicating that NTLMv2 is not enabled by default on the SQL server hosting the usage database. Set this flag to **Yes** if NTLMv2 is enabled on that SQL server. 8. **DATABASE SERVER NAME** – Specify the name of Microsoft SQL server that hosts the usage database to be accessed by this test. 9. **DATABASE NAME** – Specify the name of the usage database that this test should access. 10. **DATABASE USER NAME, DATABASE PASSWORD, CONFIRM PASSWORD** - Specify the credentials of a user who has read-only access to the usage database configured, in the **DATABASE USER NAME** and **DATABASE PASSWORD** text boxes. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box. 11. **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.

Outputs of the test	One set of results for each web application in the monitored SharePoint server		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Reads: Indicates the number of cache reads performed by this web application.	Number	
	Average read duration: Indicates the average time taken by this web application to read from the cache.	Msecs	A low value is desired for this measure. A consistent increase in the value of this measure could indicate a reading bottleneck. One of the reasons for reading delays is insufficient memory for the cache on the host. You may want to right-size the cache to make sure that your requests are serviced quickly and efficiently. Reading delays can also occur if the cache is overloaded with read requests or too much data is to be read. Whenever this measure registers an abnormally high value for a web application, look up the value reported by the *Reads* and *Average read size* measures for the same web application to determine whether/not the slowness can be attributed to the count and size of the reads.
	Average read size: Indicates how many kilobytes of data, on an average, are read from the cache by this web application.	KB	
	Writes: Indicates the number of writes to the cache by this web application.	Number	

	Average write duration: Indicates the average time taken by this web application to write to the cache.	Number	A low value is desired for this measure. A consistent increase in the value of this measure could indicate a writing bottleneck. One of the reasons for writing delays is insufficient memory for the cache. You may want to right-size the cache to make sure that your requests are serviced quickly and efficiently. Writing delays can also occur if the web application is overloading the cache with write requests or too much data is to be written. Whenever this measure registers an abnormally high value for a web application, look up the value reported by the *Writes* and *Average writes size* measure for the same application to determine whether/not the slowness can be attributed to the unusually high number and size of the writes.
	Average writes size: Indicates how many kilobytes of data, on an average, are written to the cache by this web application.	Number	
	Misses: Indicates the number of requests from this web application that were not serviced by the cache.	Number	Ideally, the value of this measure should be 0. If on the other hand, this measure value is close to the value of the *Objects requested* measure, it is a cause for serious concern, as it implies that almost all objects requested were not found in the cache. Under such circumstances, use the detailed diagnosis of this measure to know which web site addresses could not be found in the cache. One of the reasons for a high number of misses could be insufficient memory allocation to the cache service. In such a situation, you may want to increase the cache size by adding more memory.
	Hits: Indicates the number of requests from this web application that were successfully served by this cache.	Number	Ideally, the value of this measure should be the same as the value of the *Objects requested* measure. If not, check whether the cache has enough memory, and if required, add more memory to it.

	Failures: Indicates the number of cache failures experienced by this web application.	Number	Ideally, the value of this measure should be 0. If this measure reports a non-zero value, then use the detailed diagnosis of this measure to know which web site addresses were being looked up in the cache when the failures occurred.
	Objects requested: Indicates the number of objects requested by this web application.	Number	

Use the detailed diagnoss of the *Misses* measure to know which web site addresses could not be found in the cache.

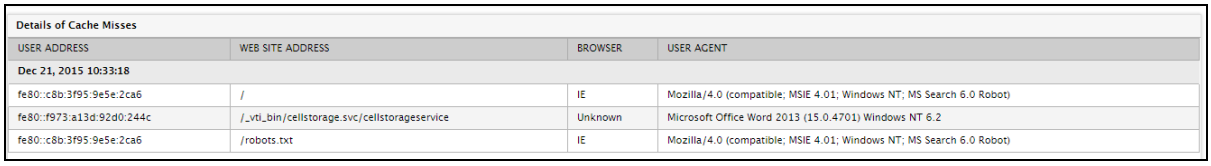

Details of Cache Misses

USER ADDRESS	WEB SITE ADDRESS	BROWSER	USER AGENT
Dec 21, 2015 10:33:18			
fe80::c8b:3f95:9e5e:2ca6	/	IE	Mozilla/4.0 (compatible; MSIE 4.01; Windows NT; MS Search 6.0 Robot)
fe80::f973:a13d:92d0:244c	/_vti_bin/cellstorage.svc/cellstorageservice	Unknown	Microsoft Office Word 2013 (15.0.4701) Windows NT 6.2
fe80::c8b:3f95:9e5e:2ca6	/robots.txt	IE	Mozilla/4.0 (compatible; MSIE 4.01; Windows NT; MS Search 6.0 Robot)

Figure 18.55: The detailed diagnosis of the Misses measure

Use the detailed diagnosis of the *Failures* measure to know which web site addresses were being looked up in the cache when the failures occurred.

Details of Cache Failures

USER ADDRESS	WEB SITE ADDRESS	BROWSER	USER AGENT
Dec 21, 2015 10:33:18			
fe80::f973:a13d:92d0:244c	/_vti_bin/sharedaccess.asmx	Unknown	Microsoft Office Word 2013 (15.0.4701) Windows NT 6.2

Figure 18.56: The detailed diagnosis of the Failures measure

18.2.6.6 Critical or Slow Web Parts Test

A web site/web application on SharePoint typically constitutes numerous web pages. The performance of each of these pages is a key determinant of user experience with the web site/web application as a whole! Each web page in turn is made up of multiple web parts, So, when a web page slows down, more often than not, one/more web parts that constitute that web page will be responsible for the slowness! This is why, when users complain that a web site/web application is slow, administrators should first check whether/not the web parts used to build that web site/web application are taking too long to load, and if so, why. This is where the **Critical or Slow Web Parts** test helps!

For each web site, this test reports the time taken by the web parts in that site to load. In the event of undue delay in web part loading, the test also points administrators to the probable cause of the delay – is it because of processing delays in the web part? Is it because the web parts took too long to generate service calls? Or is it owing to inefficient queries to the backend database? Detailed diagnosis of the test also leads you to the precise service calls and queries that could have contributed to the slowness.

For this test to run and report metrics, the following pre-requisites should be fulfilled:

- **A SharePoint Usage and Health Service application should be created and should be configured to collect usage and health data. To know how to create and configure this application, follow the steps detailed in Section 18.2.6.1.1.**
- **The SharePoint Developer Dashboard should be enabled. The steps for the same are detailed in Section 18.2.6.6.1 of this document.**

Purpose	For each web site, this test reports the time taken by the web parts in that site to load. In the event of undue delay in web part loading, the test also points administrators to the probable cause of the delay – is it because of processing delays in the web part? Is it because the web parts took too long to generate service calls? Or is it owing to inefficient queries to the backend database? Detailed diagnosis of the test also leads you to the precise service calls and queries that could have contributed to the slowness.
Target of the test	A Sharepoint Server
Agent deploying the test	An internal/remote agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">
1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the HOST.
4. SQL PORT NUMBER – Specify the port number of the SQL server hosting the usage database.
5. INSTANCE – If the SQL server hosting the usage database is instance-based, then provide the instance name here. If not, then set this to none.
6. SSL – If the SQL server hosting the usage database is SSL-enabled, then set this flag to Yes. If not, set it to No.
7. ISNTLMV2 - In some Windows networks, NTLM (NT LAN Manager) may be enabled. NTLM is a suite of Microsoft security protocols that provides authentication, integrity, and confidentiality to users. NTLM version 2 (“NTLMv2”) was concocted to address the security issues present in NTLM. By default, the ISNTLMV2 flag is set to No, indicating that NTLMv2 is not enabled by default on the SQL server hosting the usage database. Set this flag to Yes if NTLMv2 is enabled on that SQL server.
8. DATABASE SERVER NAME – Specify the name of Microsoft SQL server that hosts the usage database to be accessed by this test.
9. DATABASE NAME – Specify the name of the usage database that this test should access.
10. DATABASE USER NAME, DATABASE PASSWORD, CONFIRM PASSWORD - Specify the credentials of a user who has read-only access to the usage database configured, in the DATABASE USER NAME and DATABASE PASSWORD text boxes. Then, confirm the password by retyping it in the CONFIRM PASSWORD text box.
11. MAX ACCEPTABLE DURATION – By default, this parameter is set to 3 (seconds). This implies that this test, by default, will report metrics for only those web sites containing one/more web parts that process requests for a duration longer than the 3 seconds. You can increase or decrease the value of this parameter, depending upon what you think is ‘slow’ in your environment. This way, you can configure the test to focus on only those web sites that contain slow or critical web parts alone.
12. 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.
</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for each web site containing web parts that process requests for a duration longer than the configured MAX ACCEPTABLE DURATION</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

	Duration: Indicates the time taken by the web parts in this web site to load.	Secs	Compare the value of this measure across web sites to identify the slowest web site. Use the detailed diagnosis of this measure to know the details of requests that were processed slowly by the web parts in this web site.
	CPU duration: Indicates the time spent by the web parts in this web site processing requests.	Secs	If the value of the *Duration* measure is abnormally high, then compare the value of this measure with that of the *Total SQL duration* and *Total service call duration* measures to determine the accurate source of the slowness - – is it because of processing delays in the web parts? Is it because the web parts took too long to generate service calls? Or is it owing to inefficient queries run by the web parts in the backend database?
	SQL queries: Indicates the number of queries executed by the web parts in this web site.	Number	Use the detailed diagnosis of this measure to know which queries were run and the duration of each query. This way, long running queries can be identified.
	Total SQL duration: Indicates the total time taken by the web parts in this web site to run SQL queries.	Secs	If the value of the *Duration* measure is abnormally high, then compare the value of this measure with that of the *Total CPU duration* and *Total service call duration* measures to determine the accurate source of the slowness - – is it because of processing delays in the web parts? Is it because the web parts took too long to generate service calls? Or is it owing to inefficient queries run by the web parts? If queries are delaying web part operations, use the detailed diagnosis of the *SQL queries* measure to identify the long running / inefficient queries and then proceed to fine-tune them.
	Sharepoint requests: Indicates the number of requests to the web parts in this web site.	Number	The detailed diagnosis of this measure reveals the URLs of the SharePoint requests and the processing time of each request. Requests for which the web parts in a web site took too long to respond can be isolated in this manner.
	Asserts: Indicates the number of asserts performed by the web parts in this web site.	Number	

	Service calls: Indicates the number of service calls generated by the web parts in this web site.	Number	The detailed diagnosis of this measure reveals the exact service calls that were generated by the web parts and the duration of each service call. Slow service calls can thus be identified.
	Total service call duration: Indicates the total time taken by the web parts in this web site to generate service calls.	Secs	If the value of the *Duration* measure is abnormally high, then compare the value of this measure with that of the *Total CPU duration* and *Total SQL duration* measures to determine the accurate source of the slowness - – is it because of processing delays in the web parts? Is it because the web parts took too long to generate service calls? Or is it owing to inefficient queries run by the web parts? If service calls are delaying web part operations, use the detailed diagnosis of the *Service calls* measure to identify those service calls that were taking too long to generate and execute.

Use the detailed diagnosis of the *Duration* measure to know the details of requests that were processed slowly by the web parts in a web site. From these details, you can figure out the start time of each request, from which user the request was received, from which client the user connected to the web site, and on which server the web site was operating.

Details of Duration

CORRELATION ID	START TIME	CLIENT ADDRESS	SERVER NAME	USERNAME	MANAGED MEMORY	PAGE CHECKOUT LEVEL
Dec 23, 2015 11:12:32						
6CBF4D9D-24DD-0010-BD6A-60DE2E471742	Dec 22, 2015 13:03:48	fe80::f973:a13d:92d0:244c%11	2K8R2SP2K13	2k8r2sp2k13\administrator	N/A	Published
CCC04D9D-0492-0010-BD6A-6AF4FE5506AE	Dec 22, 2015 13:27:48	fe80::f973:a13d:92d0:244c%11	2K8R2SP2K13	2k8r2sp2k13\administrator	N/A	Published

Figure 18.57: The detailed diagnosis of the Duration measure

Use the detailed diagnosis of the *SQL queries* measure to know which queries were run and the duration of each query. This way, long running queries can be identified.

Details of SQL queries

CORRELATION ID	START TIME	END TIME	SCOPEID	NAME	NO OF CALLS	DURATION	READS	WRITES	SQL CPU (MS)	SQL DURATION (MS)
Dec 23, 2015 11:12:32										
6CBF4D9D-24DD-0010-BD6A-60DE2E471742	2015-12-22 13:04:19.6510904	Dec 22, 2015 13:04:20	2809344948174936	RawSqlText	1	1257.0436	63	2	140	1125
CCC04D9D-0492-0010-BD6A-6AF4FE5506AE	2015-12-22 13:28:30.6375499	Dec 22, 2015 13:28:32	1609232771186770	proc_GetContextObjectEventReceivers	1	1827.7537	6	0	47	351
6CBF4D9D-24DD-0010-BD6A-60DE2E471742	2015-12-22 13:04:13.1426493	Dec 22, 2015 13:04:13	2809344948174927	RawSqlText	1	60.2843	249	0	0	2
6CBF4D9D-24DD-0010-BD6A-60DE2E471742	2015-12-22 13:04:05.2776195	Dec 22, 2015 13:04:11	2809344948174886	proc_FetchDocForHttpGet	0	0	0	0	15	19
6CBF4D9D-24DD-0010-BD6A-60DE2E471742	2015-12-22 13:04:18.8149635	Dec 22, 2015 13:04:19	2809344948174933	proc_GetListWebParts	1	303.7063	57	0	0	0
CCC04D9D-0492-0010-BD6A-6AF4FE5506AE	2015-12-22 13:28:32.4984653	Dec 22, 2015 13:28:33	1609232771186771	proc_GetWebMetainfo	1	770.4265	32	0	0	3
6CBF4D9D-24DD-0010-BD6A-60DE2E471742	2015-12-22 13:04:13.7198720	Dec 22, 2015 13:04:17	2809344948174931	proc_GetListWebParts	1	3605.6354	51	0	63	2808
CCC04D9D-0492-0010-BD6A-6AF4FE5506AE	2015-12-22 13:29:08.0668470	Dec 22, 2015 13:29:09	1609232771186776	RawSqlText	1	1012.7146	20	0	0	1
CCC04D9D-0492-0010-BD6A-6AF4FE5506AE	2015-12-22 13:28:33.4327762	Dec 22, 2015 13:28:37	1609232771186773	RawSqlText	1	4373.9785	6	0	16	42

Figure 18.58: The detailed diagnosis of the SQL queries measure

The detailed diagnosis of the *SharePoint requests* measure reveals the URLs of the SharePoint requests and the processing time of each request. Requests for which the web parts in a web site took too long to respond can be isolated in this manner.

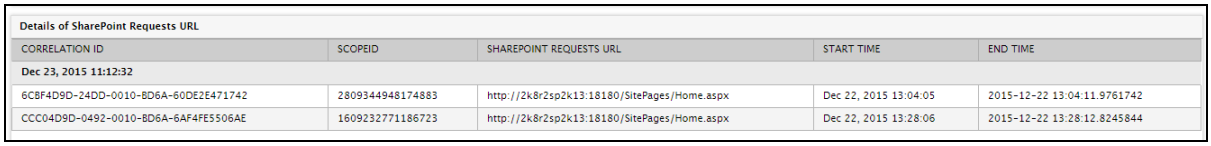

Details of SharePoint Requests URL

CORRELATION ID	SCOPEID	SHAREPOINT REQUESTS URL	START TIME	END TIME
Dec 23, 2015 11:12:32				
6C8F4D9D-24DD-0010-8D6A-60DE2E471742	2809344948174883	http://2k8r2sp2k13:18180/SitePages/Home.aspx	Dec 22, 2015 13:04:05	2015-12-22 13:04:11.9761742
CCC04D9D-0492-0010-8D6A-6AF4FE5506AE	1609232771186723	http://2k8r2sp2k13:18180/SitePages/Home.aspx	Dec 22, 2015 13:28:06	2015-12-22 13:28:12.8245844

Figure 18.59: The detailed diagnosis of the SharePoint requests measure

The detailed diagnosis of the *Service calls* measure reveals the exact service calls that were generated by the web parts and the duration of each service call. Slow service calls can thus be identified.

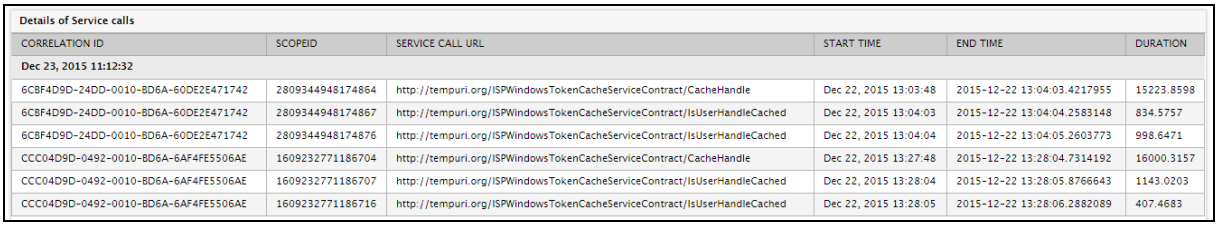

Details of Service calls

CORRELATION ID	SCOPEID	SERVICE CALL URL	START TIME	END TIME	DURATION
Dec 23, 2015 11:12:32					
6C8F4D9D-24DD-0010-8D6A-60DE2E471742	2809344948174864	http://tempuri.org/ISPWindowsTokenCacheServiceContract/CacheHandle	Dec 22, 2015 13:03:48	2015-12-22 13:04:03.4217955	15223.8598
6C8F4D9D-24DD-0010-8D6A-60DE2E471742	2809344948174867	http://tempuri.org/ISPWindowsTokenCacheServiceContract/IsUserHandleCached	Dec 22, 2015 13:04:03	2015-12-22 13:04:04.2583148	834.5757
6C8F4D9D-24DD-0010-8D6A-60DE2E471742	2809344948174876	http://tempuri.org/ISPWindowsTokenCacheServiceContract/IsUserHandleCached	Dec 22, 2015 13:04:04	2015-12-22 13:04:05.2603773	998.6471
CCC04D9D-0492-0010-8D6A-6AF4FE5506AE	1609232771186704	http://tempuri.org/ISPWindowsTokenCacheServiceContract/CacheHandle	Dec 22, 2015 13:27:48	2015-12-22 13:28:04.7314192	16000.3157
CCC04D9D-0492-0010-8D6A-6AF4FE5506AE	1609232771186707	http://tempuri.org/ISPWindowsTokenCacheServiceContract/IsUserHandleCached	Dec 22, 2015 13:28:04	2015-12-22 13:28:05.8766643	1143.0203
CCC04D9D-0492-0010-8D6A-6AF4FE5506AE	1609232771186716	http://tempuri.org/ISPWindowsTokenCacheServiceContract/IsUserHandleCached	Dec 22, 2015 13:28:05	2015-12-22 13:28:06.2882089	407.4683

Figure 18.60: The detailed diagnosis of the Service calls measure

18.2.6.6.1 Enabling the SharePoint Developer Dashboard

The **Developer Dashboard** is to help developers/ administrators with logging and debugging custom components that are added to SharePoint pages. It uses a simple process of elimination to isolate which component in a SharePoint page is the root cause for the slow performance. Developer Dashboard provides details on

- How long it took threads to process
- Details on the quantity and execution duration of SQL Server database calls and
- Details on Windows Communication Foundation (WCF) service calls
- Details on the URL
- Current user
- Load time etc.,

The Developer Dashboard can be turned on or enabled for an entire farm to get a snapshot of the activity and performance of a specific page request.

To enable this dashboard, run the following commands from the SharePoint management shell:

```
$content = ([Microsoft.SharePoint.Administration.SPWebService]::ContentService)
$appsetting =$content.DeveloperDashboardSettings
$appsetting.DisplayLevel =
[Microsoft.SharePoint.Administration.SPDeveloperDashboardLevel]::On
$appsetting.Update()
```

Chapter 19

Monitoring Microsoft Dynamics AX

Microsoft Dynamics® AX is an integrated, adaptable business management solution that streamlines financial, customer relationship, and supply chain processes. This ERP solution consolidates and standardizes processes, provides visibility across your organization, and simplifies compliance.

Since decision-makers rely on this solution for working efficiently and taking prompt and accurate decisions, slowdowns experienced by the solution and exceptions thrown by the AX portal can greatly impair the productivity and the decision-making ability of the users, and can ultimately affect revenues.

To avert this, the AX Application Object Server (AOS) and the AX portal need to be continuously monitored, and users promptly alerted to processing delays, overloads, and errors.

eG Enterprise provides a dedicated *Microsoft Dynamics AX* monitoring model that proactively detects and promptly alerts users to issues in the performance of the Dynamics AX solution.

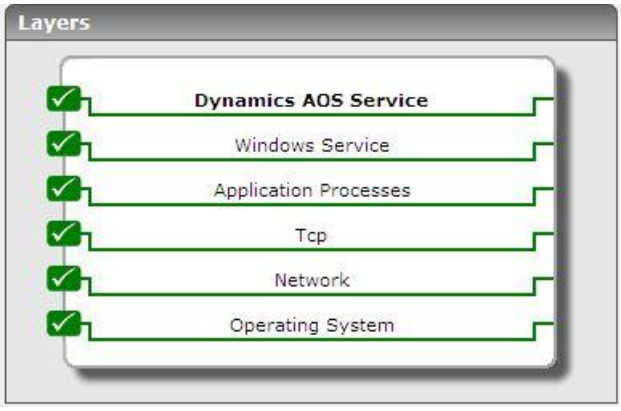

Figure 19.1: The layer model of the Microsoft Dynamics AX solution

Each layer in Figure 19.1 is mapped to a set of tests, which employ agent-based or agentless techniques to extract critical performance statistics from the AX solution. These metrics provide answers to the following key question:

- Is the AX server overloaded with requests?

- Is the server able to process the requests quickly?
- Has the AX Enterprise portal encountered any .NET business connector exceptions? If so, how many, and of what type?

Since the last 5 layers of Figure 19.1 have been discussed in-depth in the *Monitoring Unix and Windows Servers* document, this chapter will be discussing the top layer alone.

19.1 Dynamics AOS Service

The tests mapped to this layer monitors the load and the processing ability of the Application Object Server (AOS), and also captures exceptions (if any) that are encountered by the AX Enterprise Portal.

Figure 19.2: The tests mapped to the Dynamics AOS Service

19.1.1 AX Object Statistics Test

This test reports useful statistics with the help of which the session, request, and data load on the Application Object Server (AOS) can be ascertained.

Purpose	Reports useful statistics with the help of which the session, request, and data load on the Application Object Server (AOS) can be ascertained
Target of the test	A Microsoft Dynamics AX server
Agent deploying the test	An internal agent

Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Total sessions to AOS: Indicates the total number of active sessions on the server during the last measurement period.	Number	This is a good indicator of the session load on the server.
	Currently active sessions to AOS: Indicates the number of currently active server sessions.	Number	
	Client-to-server requests handled: Indicates the number of client-to-server requests during the last measurement period.	Number	This measure is a good indicator of the workload on the server.
	Client-to-server processing rate: The number of client-to-server requests processed per second	Reqs/Se	A low rate could indicate a processing bottleneck.
	Server-to-client requests processed: Indicates the number of server-to-client requests processed during the last measurement period.	Number	
	Data transmitted by server: Indicates the number of bytes sent by the server during the last measurement period.	Number	These measures are good indicators of the data load on the server.

	Data received by server: Indicates the number of bytes received by the server since the last measurement period.	Number	

19.1.2 AX Portal Statistics Test

This test reports critical statistics related to the .NET Business Connector sessions on the Microsoft Dynamics server.

Purpose	Reports critical statistics related to the .NET Business Connector sessions		
Target of the test	A Microsoft Dynamics AX server		
Agent deploying the test	An internal agent		
Configurable parameters for the test	1. **TEST PERIOD** - How often should the test be executed 2. **HOST** - The host for which the test is to be configured 3. **PORT –** Refers to the port used by the **HOST**.		
Outputs of the test	One set of results for the AX portal being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active .NET business connector sessions to portal: number of currently active .NET Business Connector sessions.	Number	This is a good indicator of the session load on the server.
	Web part execution and rendering time: Indicates the time in seconds taken to execute and render a Web Part.	Secs	A high value indicates that Web Part renditions takes too long.
	Fatal .NET business connector session exceptions: Indicates the Fatal .NET business connector session exceptions.	Number	For Enterprise Portal, this means that the page was not rendered. A Windows Sharepoint Services error page was displayed to the user.

	Non-fatal .NET business connector session exceptions: Indicates the number of nonfatal .NET Business Connector session exceptions.	Number	For Enterprise Portal, this means that the page was rendered, but some Web Parts on the page were not rendered.
	X++ .NET session exceptions: Indicates the number of X++ .NET session exceptions.	Number	
	.NET business connector sessions allocated: Indicates the total number of .NET Business Connector sessions allocated during the last measurement period.	Number	
	.NET business connector sessions disposed: Indicates the total number of .NET Business Connector sessions disposed during the last measurement period.	Number	
	.NET business connector session allocation rate: Indicates the NET business connector session allocation rate.	Number/Sec	

Chapter 20

Monitoring the Microsoft RDS License Server

A Microsoft RDS License server is a computer on which the TS Licensing role service is installed. A license server stores all RDS CALs (Microsoft RDS server Client Access Licenses) that have been installed for a group of Microsoft RDS servers and tracks the RDS CALs that have been issued. One license server can serve many Microsoft RDS servers simultaneously. As clients connect to a Microsoft RDS server, the Microsoft RDS server determines if the client needs a RDS CAL, requests a RDS CAL from a license server, and then delivers that RDS CAL to the client. In the absence of RDS CALs, users will neither be able to connect to the Microsoft RDS server, nor access any of the applications published on it. To avoid this, you will have to continuously track license usage by the Microsoft RDS clients, proactively detect a potential contention for licenses, and ensure that the Terminal License server has adequate number of licenses to support the current and future load of the Microsoft RDS server.

eG Enterprise provides a *Microsoft RDS License* monitoring model that periodically monitors the usage of the licenses stored on the Microsoft RDS License server and promptly alerts administrators if the license server is about to run out of RDS CALs.

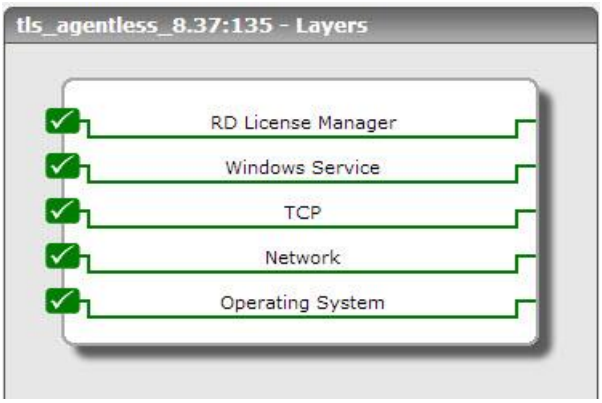

Figure 20.1: Layer model of the Microsoft RDS License server

Each layer of Figure 20.1 is mapped to a variety of tests that capture even the smallest of non-conformances that a Terminal license server experiences. Using the metrics reported, the following performance queries can be accurately answered:

- Is the Microsoft RDS License server available over the network? Is it responding quickly to requests?
- How many RDS CALs are managed by the Microsoft RDS License server per Microsoft RDS server?
- Are too many users connecting to any particular Microsoft RDS server causing excessive usage of TS CALs?
- Will any Microsoft RDS server require additional licenses to be installed? If so, which Microsoft RDS server is it?
- Is any license about to expire?
- Are there any inactive licenses? If so, which ones are they?

Since the four layers at the bottom of Figure 20.1 have already been dealt with in the *Monitoring Windows and Unix Servers* document, this chapter will discuss the first layer only.

20.1 RD License Manager Layer

This layer monitors license usage.

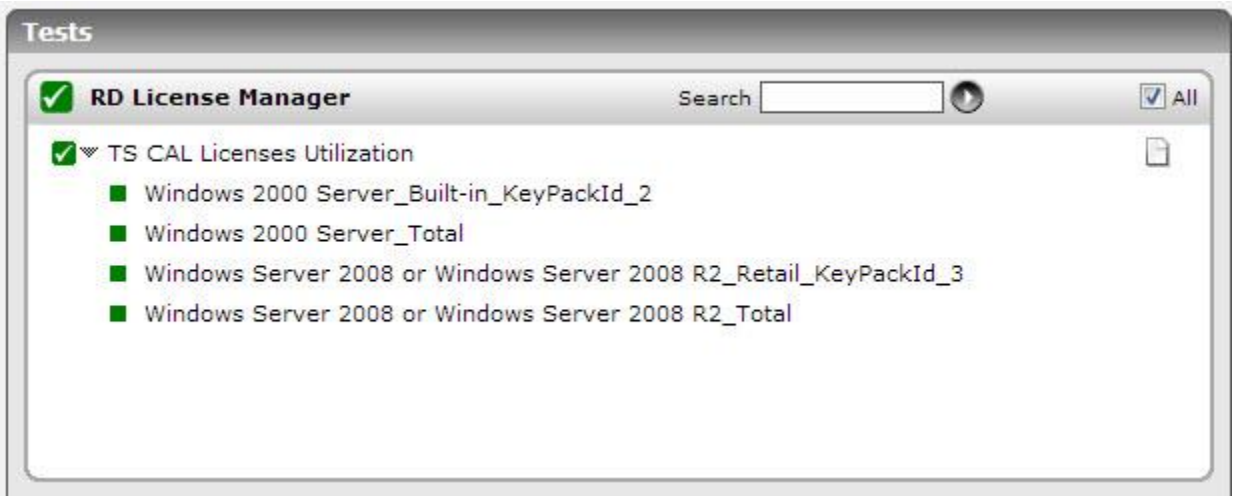

Figure 20.2: The tests mapped to the TS CAL Licenses Utilization test

20.1.1 TS CAL Licenses Utilization Test

Without an RDS CAL, a Microsoft RDS client cannot connect to a Microsoft RDS server and access the applications operating on that server. It is hence imperative that administrators periodically check whether/not the Microsoft RDS License server has enough RDS CALs to support the current and future user load of the Microsoft RDS server. To achieve this, administrators can use the **TS CAL Licenses Utilization** test. For every Microsoft RDS server that is managed by the license server, this test reports the number, type, and usage of RDS CALs installed on the Microsoft RDS License server under a particular *Key pack ID* and purchased under a specific *License program or Purchase method* (this can be, *Unknown*, *Retail, Built-in, Volume, Concurrent, Temporary, Open*). Optionally, you can also

group license usage by Microsoft RDS server alone (and not by key pack ID and purchase method). Using these statistics, you can rapidly detect probable license shortages and accurately point to the Microsoft RDS server that will potentially run out of licenses.

Purpose	For every Microsoft RDS server that is managed by the license server, this test reports the number, type, and usage of RDS CALs installed on the Microsoft RDS License server under a particular *Key pack ID* and purchased under a specific *License program or Purchase method* (this can be, *Unknown*, *Retail, Built-in, Volume, Concurrent, Temporary, Open*)
Target of the test	A Microsoft RDS License server
Agent deploying the test	An internal/remote agent

<table>
<tr><td>Configurable parameters for the test</td><td colspan="3">1. TEST PERIOD - How often should the test be executed
2. HOST - The host for which the test is to be configured
3. PORT – Refers to the port used by the HOST.
4. REPORT TOTAL - By default, this flag is set to Yes. This indicates that by default, the test reports license usage per <Microsoft_RDS_server>_<License_program>_<KeypackID> combination and also reports the total license usage across all key packs and license programs relevant to a particular Microsoft RDS server. This is why, by default, in addition to descriptors represented by a combination of <Microsoft_RDS_server>_<License_program>_<KeypackID>, a Total descriptor also appears for this test for every Microsoft RDS server. If you want the test to report metrics per <Microsoft_RDS_server>_<License_program>_<KeypackID> combination only, then set this flag to No.
5. REPORT ONLY TOTAL - If you want the test to report metrics for the Total descriptor (of every Microsoft RDS server) alone and not for each <Microsoft_RDS_server>_<License_program>_<KeypackID> combination, set this flag to Yes. In this case, the test will report metrics for the Licenses in use measure alone. By default, this flag is set to No.
6. IGNORE PER USER CALS - Microsoft RDS servers can operate in two licensing modes: Per Device (default factory setting) and Per User. A Per Device CAL gives each client computer or device the right to access a Microsoft RDS server. Using Per User licensing on the other hand, one user can access a Microsoft RDS server from an unlimited number of devices and only one CAL is needed instead of a CAL for each device.
If you want this test to ignore the CALs that have been installed in the 'Per User' mode when computing license usage, set this flag to Yes. By default, this flag is set to No, indicating that the test, by default, also considers the CALs installed in the per user mode when reporting license utilization.
7. IGNORE TEMPORARY LICENSES - By default, this flag is set to No. This implies that the test, by default, includes temporary licenses as well in the count of installed and used licenses. To make sure that the test disregards temporary licenses when computing license usage, set this flag to Yes.
8. 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.</td></tr>
<tr><td>Outputs of the test</td><td colspan="3">One set of results for every <Microsoft_RDS_server>_<License_program>_<KeypackID> combination</td></tr>
<tr><td>Measurements made by the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
</table>

	CAL Type: Indicates the CAL type of the licenses for this Microsoft RDS server, installed under this Key pack ID and license program.		CALs apply to either a "device" (as defined in the license agreement) or a "user". A business is free to choose either mode. In *Per-User* mode, a CAL is purchased to allow one user to connect to the server software. Any user can connect, but only one user may use a given CAL at any given time. Any number of CALs can be purchased to allow five, five hundred, or any number of users to simultaneously connect to the server. Any number of devices may connect to the server software, but only a set number of users can connect to it at once. Per-device mode operates in much the same way, but limits connections made by devices, rather than users. One CAL enables one device to connect to and use the server software, regardless of how many users are connecting. If the CALs apply to a *user*, then the value of this measure will be *Per User*. If the CALs apply to a *device*, then the value of this measure will be *Per Device*. In the case of invalid CALs, the value of this measure will be *Not Valid*. The numeric values that correspond to these measure values have been discussed in the table below: **Measure Value** / **Numeric Value** Per User / 1 Per Device / 0 or 3 Not Valid / 2 **Note:** By default, the measure reports the **Measure Values** listed in the table above to indicate the CAL type. However, in the graph of this measure, the same will be represented using the numeric equivalents only.

			This measure will not be available for the Total descriptor. You can use the detailed diagnosis of this measure to know the license program and expiration date of each license.
	Total license: Indicates the total number of licenses installed for this Microsoft RDS server under this key pack ID and license program.	Number	**This measure will not be available for the Total descriptor**.
	Available licenses: Indicates the number of licenses under this key pack ID and license program that are still to be used by this Microsoft RDS server.	Number	A high value is desired for this measure. **This measure will not be available for the Total descriptor**.
	Licenses in use: Indicates the number of licenses under this key pack ID and license program that are currently used by this Microsoft RDS server. For the **Total** descriptor, this measure reports the number of licenses currently used by this Microsoft RDS server across all relevant key pack IDs and license programs.	Number	A low value is desired for this measure. Compare the value of this measure for the **Total** descriptor across all Microsoft RDS servers to identify which Microsoft RDS server is over-utilizing the CALs. Using the detailed diagnosis of this measure, you can view the complete details of license usage. This includes the License ID of every license installed under a key pack, the license program under which each license was purchased, who it was issued to and when, the expiry date of license and its current status.
	License utilization: Indicates the percentage of licenses under this key pack ID and license program that are currently used by this Microsoft RDS server.	Percent	A value close to 100% indicates excessive CAL utilization. This in turn implies that too many users are connecting to the Microsoft RDS server. You may want to install additional licenses to ensure that subsequent users are able to connect to and work with the Microsoft RDS server. **This measure will not be available for the Total descriptor**.

The detailed diagnosis of the *CAL type* measure reveals the license program and expiration date of each license. If you have installed multiple licenses using a key pack ID, you can use the detailed diagnosis to know the purchase method and expiry date of every license under that key pack ID.

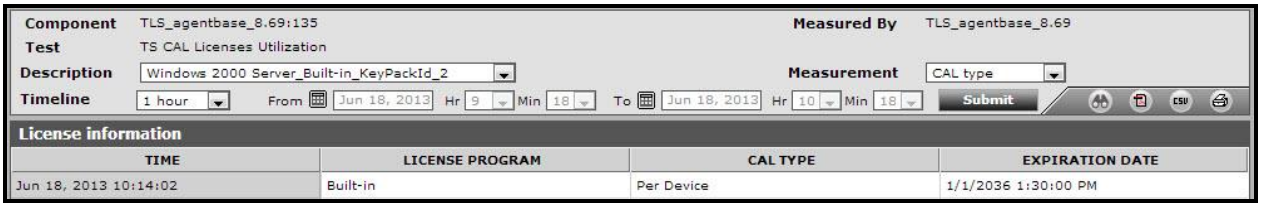

Figure 20.3: The detailed diagnosis of the CAL type measure

Using the detailed diagnosis of the *Licenses in use* measure, you can view the complete details of license usage. This includes the License ID of every license installed under a key pack, the license program under which each license was purchased, who it was issued to and when, the expiry date of license and its current status. If you notice abnormal license usage on a Terminal license server, you can use the detailed diagnosis to figure out which Microsoft RDS server was issued the maximum number of licenses. You can also identify licenses that are inactive currently, so that such licenses can be revoked and made available for the use of active connections to the Microsoft RDS server.

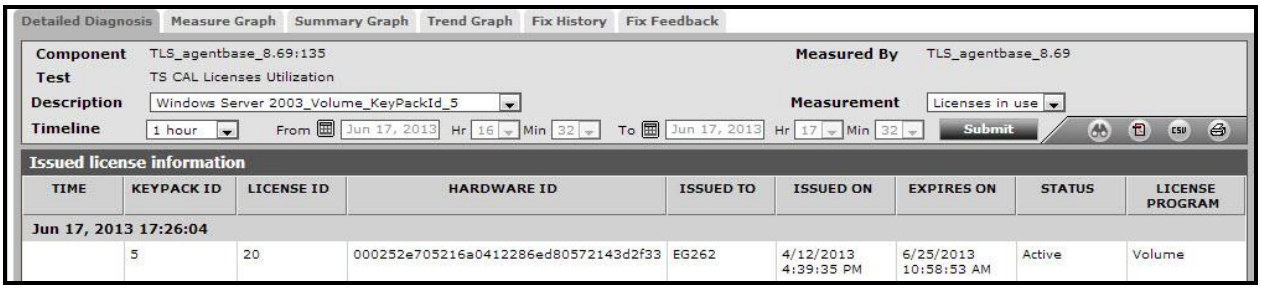

Figure 20.4: The detailed diagnosis of the Licenses in use measure

Chapter 21

Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to **Microsoft applications**. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com. We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.