Monitoring 2X ApplicationServer XG

eG Enterprise v6.1

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

x

Table of Figures

1

Introduction

2X server is an application that provides vendor independent desktops and applications. Accessible from anywhere, the 2X server allows you to publish full desktops and applications within a thin client or virtual environment, improving desktop manageability, security and performance.

The basic component of the 2X server is the LoadBalancer. This component is made up of two sub components; the Publishing Agent Service and the 2X Management Console. The Publishing Agent provides load balanced applications and desktop publishing services to clients. 2X Management Console is the GUI through which all configuration and management is performed.

The second component is the 2X Terminal Server Agent, which is typically installed on each Terminal server in the environment. The Terminal Server Agent's job is to collect resource availability information from the Terminal Services and then report that information to the LoadBalancer. That way, the ApplicationServer can always use the available Terminal Server resources efficiently.

The final component is the 2X Client Gateway. The 2X Client Gateway's job is to tunnel all of the traffic related to hosted applications over a single, secure port.

Clients connect to a 2x Client Gateway, which routes the request to the Publishing Agent. This Publishing Agent responds to the client request with the Terminal server that the client can use. The client then connects to the Terminal server either directly or via the Client Gateway to access the published applications or desktops (see Figure 1.1).

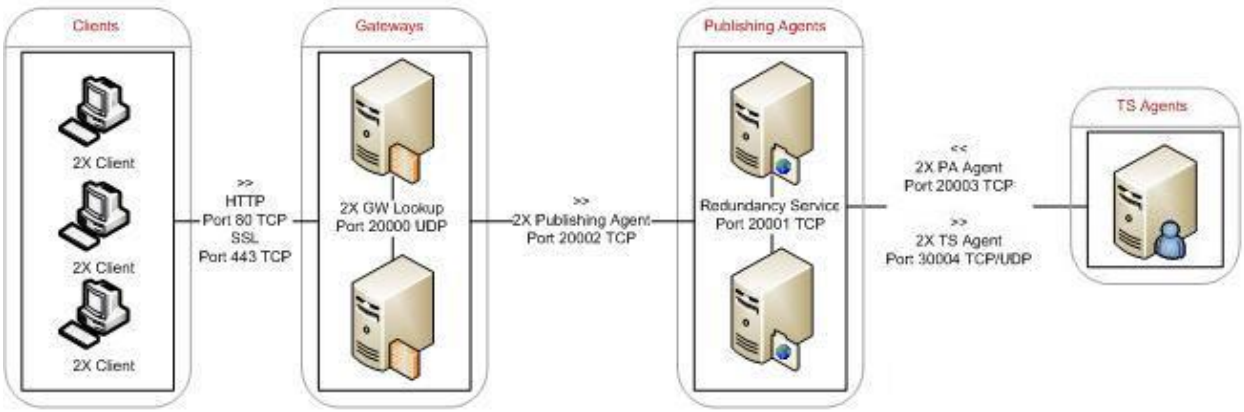

Figure 1.1: Components of the 2X server architeture

Quick and easy access to desktops, stable user sessions, and hassle-free use of published applications all make for a 'good user experience' with the 2X server. On the contrary, delayed access to desktops combined with application slow-downs can be the key spoilers of a user's experience with the 2X server. More frustrating to administrators is the fact that many times, the cause of such anomalies cannot be pinpointed, making recovery difficult. While standard monitoring solutions are available, most do not provide the critical in-depth analysis that administrators need, with many failing to provide comprehensive data into the network, storage, virtualization and application levels. Administrators thus far have been pressured to discover the reasons behind network performance failures, unable to pinpoint the problem to the network, profile server, Web access, virtualization platform or other components.

To meet the high standards of network administrators, eG 2X Monitoring provides total performance visibility for 2X installations of all types. As part of the eG Enterprise suite, eG 2X Monitoring is a comprehensive management solution for 2X farms, providing complete visibility and monitoring for all layers and tiers of the organization's 2X infrastructure, including the 2X Terminal server, the publishing agent, client gateway, network, storage and more.

Towards this end, eG Enterprise offers three specialized monitoring models - namely, the *2X Terminal Server* model, the *2X Publishing Agent* model, and the *2X Client Gateway* model - that focus on the overall performance and problems related to the 2X Terminal server, the publishing agent, and the client gateway components, respectively.

This document discusses all these models in detail.

2

Monitoring the 2X Terminal Server

eG Enterprise provides a 100%, web-based *2x Terminal Server* monitoring model that periodically checks the availability of the Terminal server, monitors user logins to the server and their authentication, discovers the published applications and their resource usage, and sends out proactive alerts to administrators if abnormalities are sensed in any of the monitored activities.

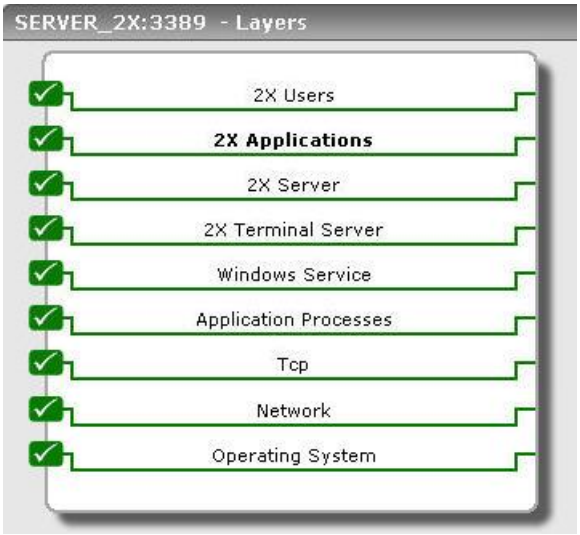

Figure 2.1: The layer model of the 2X Terminal server

Each layer of Figure 2.1 is mapped to a variety of tests that monitor various performance aspects of the 2X Terminal Server. Using the metrics reported by these tests, administrators can find quick and accurate answers to the following queries:

a. Is the Redirector component functioning properly? Are too many requests to the Redirector pending?

b. Are any user profiles on the Terminal server exceeding their prescribed quota?

c. Is login authentication taking too long?

d. Are any applications on the server consuming resources excessively? If so, which applications are these?

e. Is the session load too high on the server?

f. Are too many disconnected sessions running on the server?

g. Did too many sessions to the server suddenly log out?

h. Are clients able to connect to the server quickly?

i. Which users are currently connected to the server? Which one of these users is executing resource-intensive processes on the server?

j. Did any sessions to the server disconnect recently and reconnect soon after?

The sections that follow will discuss the top 4 layers of Figure 2.1, as the rest of the layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

2.1 The 2x Terminal Server Layer

Using the tests mapped to this layer, you can ascertain the following:

- Whether the Redirector component is functioning without a glitch
- Whether client logons are enabled/disabled

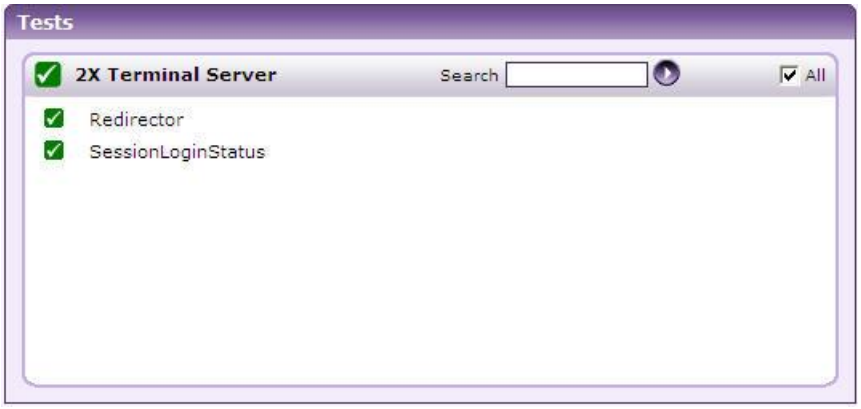

Figure 2.2: The tests mapped to the 2X Terminal Sever layer

2.1.1 Redirector Test

File serving very often is a much underestimated part of server-based computing environments. Improperly configured file serving components can wreak havoc on a server farm's performance.

File serving in server-based computing 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 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 server-based computing 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.

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	A 2X Terminal 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 Terminal server		
Outputs of the test	One set of results for the 2X Terminal 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.

2.1.2 Session Login Status Test

Administrators typically use the *Change logon* command line tool to enable / disable logons from client sessions to the server-based computing systems. 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 Terminal server.

Purpose	Periodically reports the status of logons from client sessions to the Terminal server		
Target of the test	A 2X Terminal 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** - 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 2X Terminal 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 all client logons are enabled. If the value of this measure is 0, it indicates that client logons are disabled.

2.2 The 2X Server Layer

Whenever users to the server-based computing environment complaint of delays in accessing their applications/desktops, you can use the tests associated to this layer to accurately isolate the reasons for such a slowdown - is it because of the large size of the user profile? or is it because the user credentials are taking too long to be authenticated?

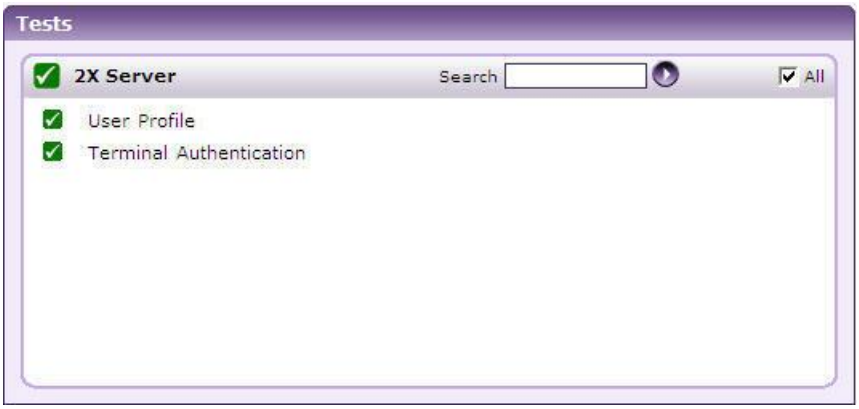

Figure 2.3: The tests mapped to the 2X Server layer

2.2.1 User Profile Test

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

Purpose	Monitors the size of the Terminal server user profiles and raises an alarm if the profile size exceeds the profile quota size
Target of the test	A 2X Terminal 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 Terminal 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. 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 enabled/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 2X Terminal 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>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.</td><td>Boolean</td><td>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.</td></tr>
<tr><td>Current profile size:
Indicates the current profile size.</td><td>MB</td><td></td></tr>
<tr><td></td><td>Number of files in user's profile:
Indicates the number of files available in the user profile.</td><td>Number</td><td></td></tr>
</table>

	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.

2.2.2 Terminal Connection Test

This test tracks various statistics pertaining to the 2X Terminal server connections to and from a host, from an external perspective.

Purpose	Tracks various statistics pertaining to the 2X Terminal server connections to and from a host, from an external perspective		
Target of the test	A 2X Terminal 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 connection to this port 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.

2.2.3 Terminal 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 local host and reports whether the login succeeded and how long it took		
Target of the test	A 2X Terminal 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 Terminal server 4. **USERNAME** - This test emulates a user logging into a Microsoft Windows domain or a local host. Therefore, specify the login name of the user here. 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. Multiple user names, passwords, and domains can be specified, separated by a comma.		
Outputs of the test	One set of results for every user account being checked		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Authentication status: Indicates whether the login was successful or not	Percent	A value of 100 % indicates that the login has succeeded. The value 0 is indicative of a failed login.
	Authentication time: Indicates the time it took to login	Secs	If this value is very high then it could be owing to a configuration issue (i.e. the domain might not be configured properly) or a slow-down/unavailability of the primary domain server.

2.2.4 ICA/RDP Listeners Test

The listener component runs on the XenApp/Terminal/2X server and is responsible for listening for and accepting new ICA/RDP client connections, thereby allowing users to establish new sessions on the XenApp/Terminal/2X 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 2X 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.

<table>
<tr><td>Purpose</td><td colspan="3">Tracks the status of the default listener ports and reports whether any of the ports is down</td></tr>
<tr><td>Target</td><td colspan="3">A 2X Terminal 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
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 rowspan="2">Measurements of the test</td><td>Measurement</td><td>Measurement Unit</td><td>Interpretation</td></tr>
<tr><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 The 2X Applications Layer

The applications deployed on the 2X Terminal Server and the resource usage of each application can be assessed using the tests mapped to this layer.

Figure 2.4: The tests mapped to the 2X Applications layer

2.3.1 2X Applications Test

This test reports statistics pertaining to the different applications deployed on the 2X Terminal Server and their usage by its clients.

Purpose	Reports statistics pertaining to the different applications deployed on the 2X Terminal Server and their usage by its clients
Target of the test	A 2X Terminal 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 Terminal server
4. APPS - Specify all here if you want to monitor all the applications that are executing on this server. To monitor only those applications that have been published on this server for use by clients, specify Published in this text box.
5. QUERY2XPATH - While monitoring 2X servers of v10.0 (and higher), this test will automatically discover the full path to the install directory of the server. This is why, this parameter is set to none by default. However, while monitoring older versions of the 2X server (i.e., versions prior to v10.0), you will have to explicitly specify the full path to the install directory of the 2X server in the query2xpath text box.
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 enabled/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 application</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>Processes running:
Number of instances of this application currently executing on the 2X Terminal Server.</td><td>Number</td><td>This value indicates if too many or too few instances corresponding to an application are executing on the host. The detailed diagnosis of this measure, if enabled, displays the complete list of processes executing, the users executing them, and their individual resource utilization.</td></tr>
<tr><td>Cpu usage:
Percentage of CPU used by this application.</td><td>Percent</td><td>A very high value could indicate that the specified application is consuming excessive CPU resources.</td></tr>
<tr><td></td><td>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.</td><td>Percent</td><td>A sudden increase in memory utilization for an application may be indicative of memory leaks in the application.</td></tr>
</table>

2.4 The 2X Users Layer

The tests mapped to this layer monitor the user activity on the server-based computing environment, and reveal the following:

k. Session overloads

l. Resource-intensive users

m. Sudden session disconnects

n. Unexpected session logouts

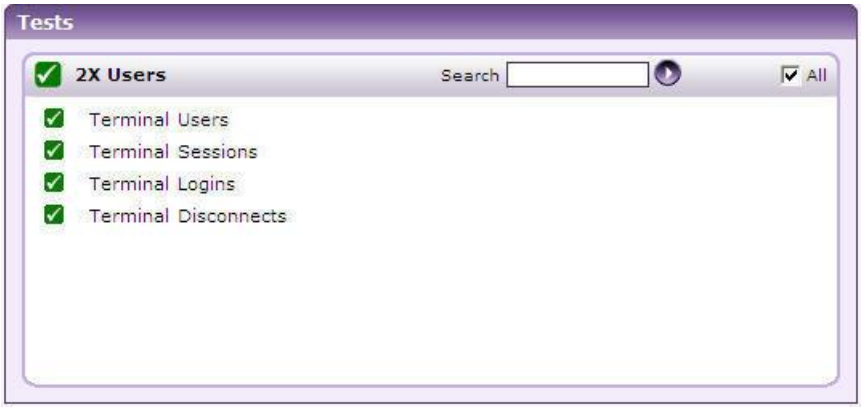

Figure 2.5: The tests mapped to the 2X Users layer

2.4.1 Terminal Sessions Test

This test reports performance statistics related to 2X Terminal server user sessions.

Purpose	Reports performance statistics related to 2X Terminal server user sessions
Target of the test	A 2X Terminal 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 Terminal 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 Terminal server's local time.
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="4">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, an administrator can obtain information that can help him/her plan the capacity of their 2X server farms. The detailed diagnosis capability, if enabled, lists the active and inactive sessions on the server.</td></tr>
<tr><td>Idle sessions:
Indicates the number of sessions that are initialized and are currently ready to accept connections.</td><td>Number</td><td>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 2X server sessions.</td></tr>
<tr><td>Connected sessions:
Indicates the current number of sessions that are connected, but no user has logged on to the server.</td><td>Number</td><td>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 2x server sessions.</td></tr>
</table>

	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 2X server 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 Terminal server cause excessive consumption of the server resources. To avoid this, a session limit is typically configured for disconnected sessions on the Terminal 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 2X server 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 2X server 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.
	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 Terminal 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.

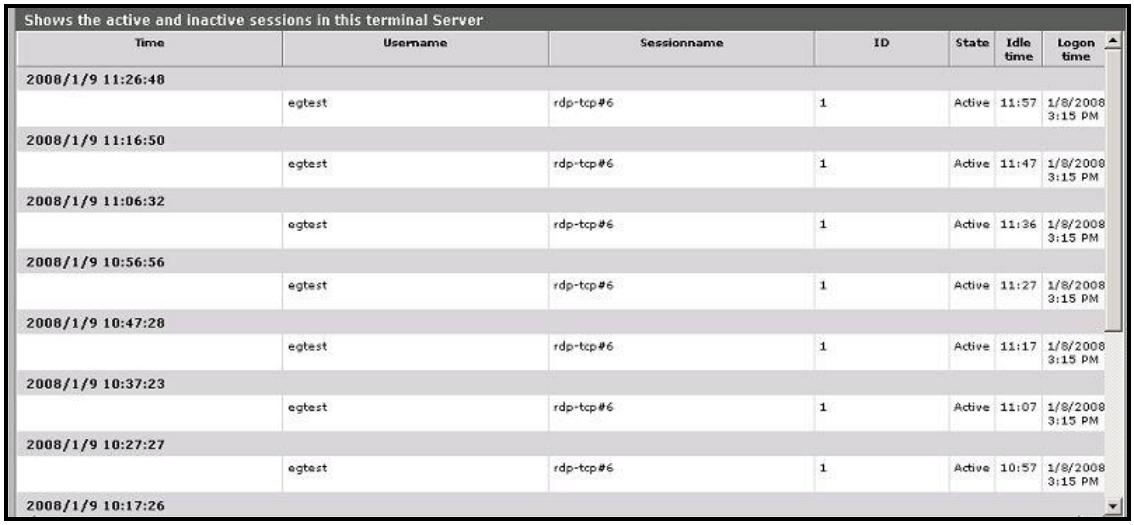

Shows the active and inactive sessions in this terminal Server

Time	Username	Sessionname	ID	State	Idle time	Logon time
2008/1/9 11:26:48						
	egtest	rdp-tcp#6	1	Active	11:57	1/8/2008 3:15 PM
2008/1/9 11:16:50						
	egtest	rdp-tcp#6	1	Active	11:47	1/8/2008 3:15 PM
2008/1/9 11:06:32						
	egtest	rdp-tcp#6	1	Active	11:36	1/8/2008 3:15 PM
2008/1/9 10:56:56						
	egtest	rdp-tcp#6	1	Active	11:27	1/8/2008 3:15 PM
2008/1/9 10:47:28						
	egtest	rdp-tcp#6	1	Active	11:17	1/8/2008 3:15 PM
2008/1/9 10:37:23						
	egtest	rdp-tcp#6	1	Active	11:07	1/8/2008 3:15 PM
2008/1/9 10:27:27						
	egtest	rdp-tcp#6	1	Active	10:57	1/8/2008 3:15 PM
2008/1/9 10:17:26						

Figure 2.6: The detailed diagnosis of the Active sessions measure

2.4.2 Terminal Logins Test

This test monitors the new logins to the 2X Terminal servers.

Purpose	Monitors the new logins to the 2X Terminal servers
Target of the test	A 2X Terminal 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 Terminal 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 Terminal server's local time.
5. 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.
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 enabled/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 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>New logins:
Indicates the number of new logins to the 2x Terminal servers in the last measurement period.</td><td>Number</td><td>A consistent zero value could indicate a connection issue.</td></tr>
<tr><td>Percent new logins:
Indicates the percentage of current sessions that logged in during the last measurement period.</td><td>Percent</td><td></td></tr>
<tr><td></td><td>Sessions logging out:
Indicates the number of sessions that logged out.</td><td>Number</td><td>If all the current sessions suddenly log out, it indicates a problem condition that requires investigation.</td></tr>
</table>

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

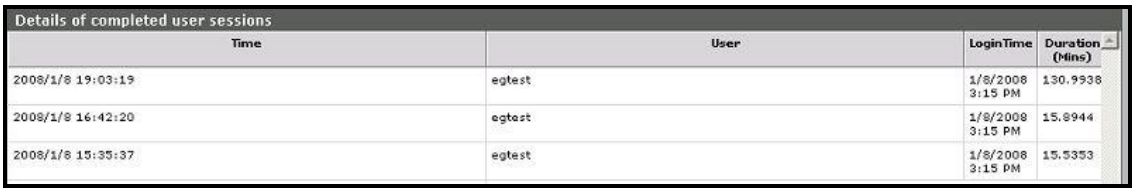

Details of completed user sessions

Time	User	LoginTime	Duration (Mins)
2008/1/8 19:03:19	egtest	1/8/2008 3:15 PM	130.9938
2008/1/8 16:42:20	egtest	1/8/2008 3:15 PM	15.8944
2008/1/8 15:35:37	egtest	1/8/2008 3:15 PM	15.5353

Figure 2.7: The detailed diagnosis of the Sessions logging out measure

2.4.3 Terminal Clients Test

This test measures the client connections to and from a 2X Terminal server.

Purpose	To monitor the client connections to and from a 2x Terminal server		
Target of the test	A 2X Terminal Server		
Agent deploying the test	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 Terminal server 4. **SERVERIP** - By default, the SERVERIP field will display the IP address of the Terminal 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.		
Outputs of the test	One set of results for every server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Current connections:** The number of TCP connections currently established by clients to the 2X Terminal server.	Number	This measure directly indicates the loading on the 2X Terminal server from clients. Typically one connection is established per active session to the server.
	New connections: The number of new TCP connections initiated by clients to the 2X Terminal server during the last measurement period.	Number	Tracking the new connections over time can provide an indication of when clients login to the 2X Terminal server. A spurt of connections and disconnections may be indicative of sporadic failures of the 2X Terminal server.
	Old connections removed: The number of TCP connections that were removed because the clients may have disconnected from the 2X Terminal server during the last measurement period	Number	A large number of sudden connection drops may be early warning indicators of problems with the 2X Terminal server.
	Avg connection duration: The average time from when a connection is established to when the corresponding connection is disconnected. The duration of a connection is measured from its start time to the current time. The accuracy of this measurement is limited by the frequency at which this test is run.	Secs	This value can provide an indicator of how long clients stay connected to a Terminal and a 2X server. This information together with the number of simultaneous clients can be useful for capacity planning in server-based computing environments.

2.4.4 Terminal Users Test

A 2X 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 this test assesses the traffic between the user terminal and the 2X Terminal 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.

Purpose	Tracks every user connection from the client to the 2X Terminal server, and monitors the resource utilization of every user on these servers
Target of the test	A 2X Terminal Server

<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 host for which the test is to be configured
3. PORT – Refers to the port used by the terminal 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 terminal server are to be monitored.
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 user logged into the 2X Terminal server</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>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 2X Terminal server.</td></tr>
<tr><td>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.</td><td>Percent</td><td>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.</td></tr>
</table>

	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 terminal 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	
	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	
	Output bandwidth: Indicates the average bandwidth used for server to client communications for all the sessions of a user	KB/Sec	
	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	
	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.	KB	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.

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 Terminal server, and reports the resource usage of each process (see Figure 2.8).

Lists the processes executed by a user on a Terminal server

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 2.8: The detailed diagnosis of the User sessions measure

2.4.5 Terminal Disconnects Test

A user session is terminated when a user logs off from the 2X/Terminal 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 2X/Terminal 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 Terminal server/2X server sessions		
Target of the test	A 2X Terminal 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 Terminal 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 Citrix server within the short period of time (in minutes) specified against reconnectperiod. 5. **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. 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 enabled/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 2X server being monitored		
Measurements made by the	**Measurement**	**Measurement Unit**	**Interpretation**

test	**Total disconnected sessions:** Indicates the total number of sessions that are in the disconnected state.	Number	
	New disconnects: Indicates the number of sessions that were disconnected in the last measurement period	Number	The detailed diagnosis of this measure, if enabled lists the users who have recently disconnected.
	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 2.9).

Details of disconnected sessions

Time	User	SessionID	ClientType
2008/1/8 19:00:30	egtest	1	rdpwd
2008/1/8 16:43:26	egtest	1	rdpwd
2008/1/8 15:30:56	egtest	1	rdpwd

Figure 2.9: 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 2.10).

Details of quick reconnected user sessions

Time	User	SessionID	ClientType	DisconnectTime	DisconnectDuration (mins)
2008/1/8 16:53:34					
	egtest	1	rdpwd	08/01/2008 16:43:26	10.13

Figure 2.10: The detailed diagnosis of the Quick reconnects measure

3

Monitoring the 2X Client Gateway

The 2X Client Gateway's job is to route requests it receives from thin clients to the 2X Publishing Agent. If HTTP/S connection to the client gateway is unavailable, then thin clients will not be able to access desktops/applications published on the 2X Terminal server.

eG Enterprise provides a specialized *2X Client Gateway* monitoring model, which periodically checks the availability of the HTTP/HTTPS connection to the 2X client gateway, and sends out prompt alerts to administrators when abnormalities are detected.

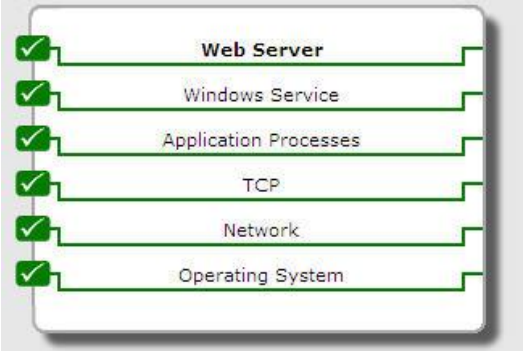

Figure 3.1: Layer model of the 2X Client Gateway

Each layer of Figure 3.1 is mapped to a series of tests that run frequent health checks on the 2X Client Gateway. Using these statistics, administrators can find accurate answers for the following performance queries:

- Can clients connect to the 2X Client Gateway via HTTP/HTTPS? If so, how quickly does the gateway respond to connection requests?
- Did too many errors occur while connecting to the gateway?
- Are all critical Windows services supporting the client gateway up and running?
- Are the critical 2X Client Gateway processes consuming resources excessively?

The section that follows will discuss the **Web Server** layer alone.

3.1 The Web Server Layer

The tests mapped to this layer reports the availability of the client gateway and the overall health of the IIS web server hosting the gateway.

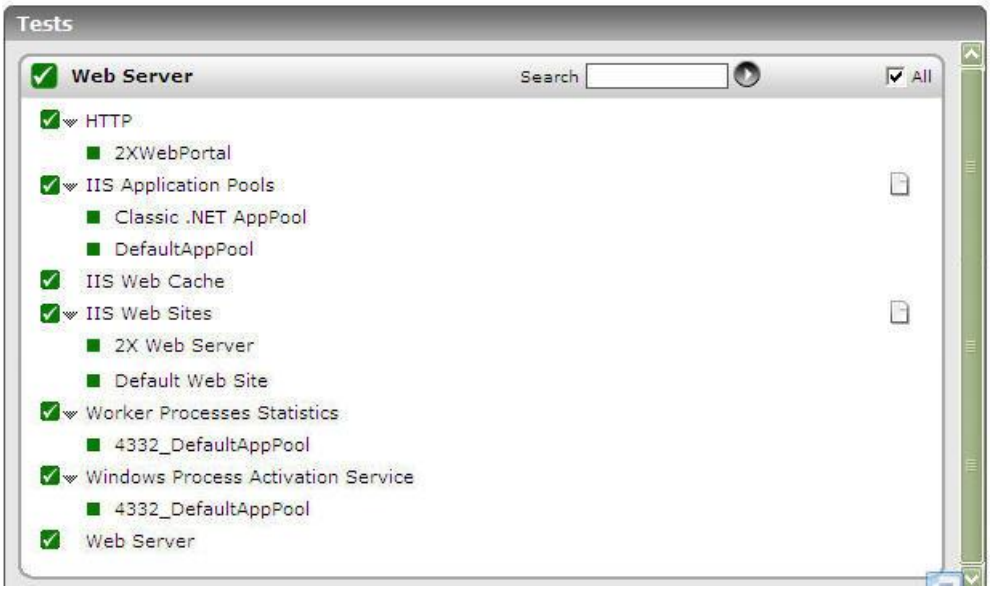

Figure 3.2: The tests mapped to the Web Server layer

Since the tests mapped to this layer have been dealt with in the *Monitoring Web Servers* doument, let us proceed to the next chapter.

4

Monitoring the 2X Publishing Agent

The Publishing Agent provides load balanced applications and desktop publishing services to clients. Whenever a client requests for a desktop/application, the 2X Client Gateway routes the request to the Publishing Agent, which responds to the client request with the Terminal server that the client can use. This means that the non-availability of or delays in the operations of the Publishing Agent can adversely impact the user experience with the 2X server.

The *2X Publishing Agent* monitoring model that eG Enterprise offers is quick to report intermittent/prolonged breaks in the availability of the 2X Client Gateway, and resource contentions that could be affecting a Terminal server's performance.

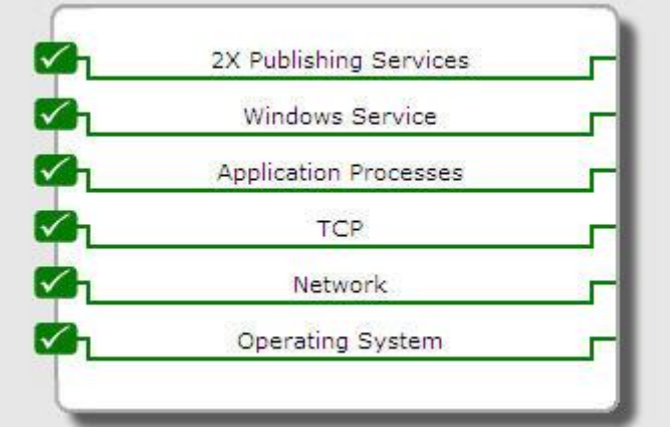

Figure 4.1: Layer model of the 2X Publishing Agent

Using the measures reported by the tests mapped to each layer of Figure 4.1 above, administrators can find quick and accurate answers to the following performance queries:

- Is any Terminal server consuming too much CPU and memory?
- Is any 2X Client Gateway unavailable currently?
- Is any 2X Client Gateway unresponsive to connection requests?
- Are all Windows services critical to the functioning of the Publishing Agent up and running?

- Is any application process on the 2X Publishing Agent consuming resources excessively?
- Is the Publishing Agent available over the network, or has any drop in network connectivity been noticed?

The sections that follow will only discuss the **2X Publishing Services** layer of Figure 4.1.

4.1 The 2X Publishing Services Layer

Using the tests mapped to this layer, resource-hungry Terminal servers and unavailable/unresponsive 2X Client Gateways can be isolated.

Figure 4.2: The tests mapped to the 2X Publishing Services layer

4.1.1 2X Terminal Server Information Test

This test measures the CPU and memory load on each of the Terminal servers. In addition, the test reports the version of the server and the maximum number of RDP and ICA connections tunnelled through the server.

Purpose	Measures the CPU and memory load on each of the Terminal servers. In addition, the test reports the version of the server and the maximum number of RDP and ICA connections tunnelled through the server
Target of the test	A 2X Publishing Agent
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** - Host name of the server for which the test is to be configured 3. **PORT** - Enter the port to which the specified host listens. The default port is 80. 4. **QUERY2XPATH** - While monitoring 2X servers of v10.0 (and higher), this test will automatically discover the full path to the install directory of the server. This is why, this parameter is set to *none* by default. However, while monitoring older versions of the 2X server (i.e., versions prior to v10.0), you will have to explicitly specify the full path to the install directory of the 2X server in the query2xpath text box.		
Outputs of the test	One set of results for each 2X Terminal server being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	CPU utilization: Indicates the overall percent of CPU utilized by the applications that are published in this server.	Percent	Ideally, this value must be low. A high value or a consistent increase in the value of this measure may indicate a performance bottleneck or a system slowdown.
	Memory utilization: Indicates the overall percent of memory utilized by the applications that are published in this server.	Percent	Ideally, this value must be low. A high value or a consistent increase in the value of this measure may indicate a performance bottleneck or a system slowdown.
	Version: Indicates the version of this server.	Number	
	Maximum RDP tunneled sessions: Indicates the maximum number of RDP sessions that are tunneled through the 2X Client Gateway of this server.	Number	
	Maximum ICA tunneled sessions: Indicates the maximum number of ICA sessions that are tunneled through the 2X Client Gateway of this server.	Number	

4.1.2 2X Gateway Status

This test periodically checks the availability of each 2X Gateway server configured in the environment, and in the process reveals which gateway is unavailable, and which gateway is not responding quickly to connection requests.

Purpose	Periodically checks the availability of each 2X Gateway server configured in the environment, and in the process reveals which gateway is unavailable, and which gateway is not responding quickly to connection requests		
Target of the test	A 2X Publishing Agent		
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** - Host name of the server for which the test is to be configured 3. **PORT** - Enter the port to which the specified host listens. The default port is 80. 4. **QUERY2XPATH** - While monitoring 2X servers of v10.0 (and higher), this test will automatically discover the full path to the install directory of the server. This is why, this parameter is set to *none* by default. However, while monitoring older versions of the 2X server (i.e., versions prior to v10.0), you will have to explicitly specify the full path to the install directory of the 2X server in the query2xpath text box.		
Outputs of the test	One set of results for every 2X Client Gateway configured in the environment		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Availability: Indicates whether this 2X Gateway is currently available or not	Percent	A value of *100* indicates that the 2X Gateway server is *Available* and the value of *0* indicates that the 2X Gateway server is *Not Available*.
	Response time: Indicates the time taken to connect to this 2X Gateway server.	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.

5

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 the **2X Terminal Server, the 2X Publishing Agent, and the 2X Client Gateway**. 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.