Monitoring Citrix NetScaler Insight

eG Enterprise v6.2

Table of contents

Table of Figures

1

Introduction

NetScaler Insight Center, a virtual appliance that runs on XenServer, VMWare ESX, or on Microsoft Hyper-V collects detailed information about web-application and virtual-desktop traffic, such as flow, user-session-level information, web page performance data, and database information flowing through the NetScaler ADCs, NetScaler Gateway appliances, or CloudBridge appliances at a site and provides actionable reports.

Figure 1.1 below depicts the architecture of the NetScaler Inisight Center.

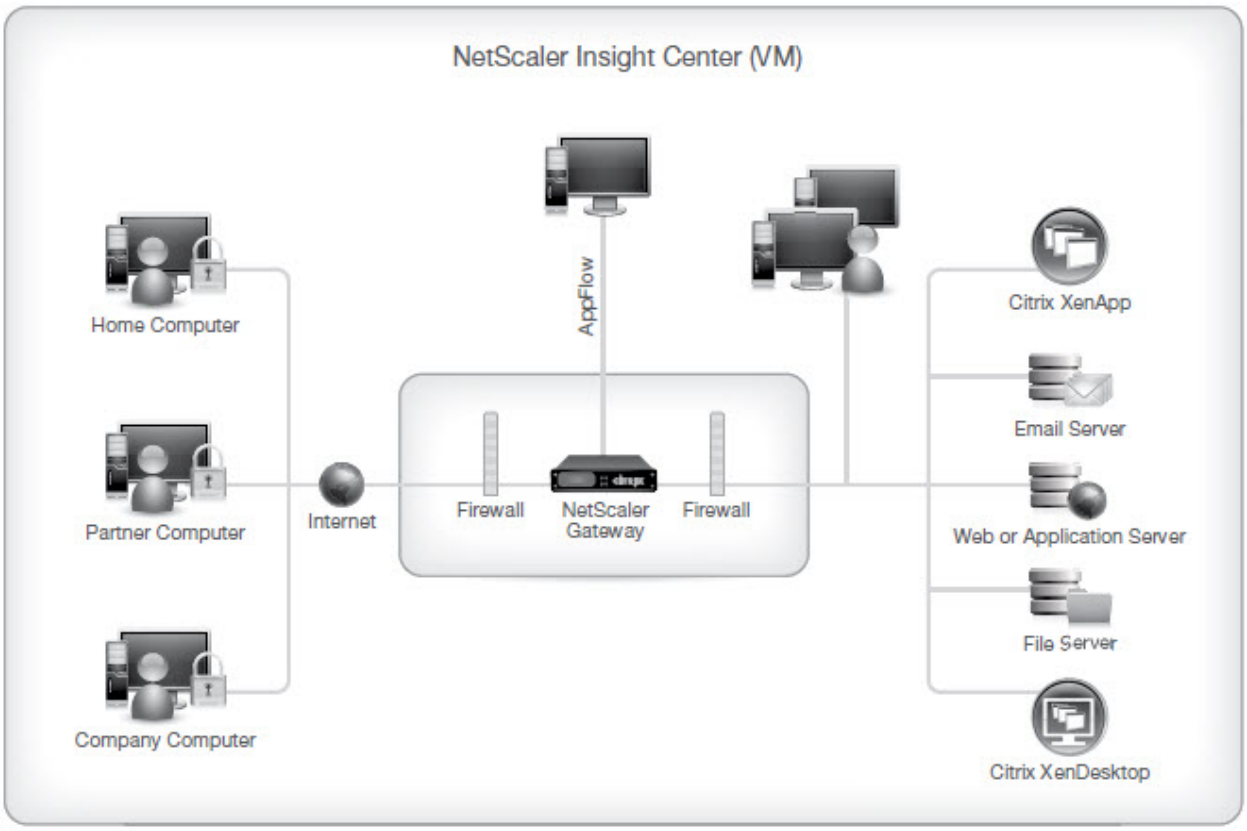

Figure 1.1: Architecture of the Citrix NetScaler Insight Center

NetScaler Insight Center has two main components:

- Web Insight that delivers data analytics for web traffic flowing through NetScaler ADCs.
- HDX Insight that delivers data analytics for XenApp and XenDesktop traffic flowing through NetScaler ADCs, NetScaler Gateway appliances, or CloudBridge appliances.

eG Enterprise integrates with the *Web Insight* and *HDX Insight* components of the NetScaler Insight Center, pulls the metrics captured by these components in real-time, and presents them in the eG monitoring console using *Citrix Web Insight* and *Citrix HDX Insight* monitoring models, respectively. Through this integration, eG

Enterprise leverages the monitoring capabilities of the NetScaler Insight Center, and provides administrators with valuable insights into the performance of applications and the user experience with applications/virtual desktops.

With that, eG can now integrate with both NetScaler Insight Center and XenDesktop Director, thus emerging as a 'one-stop-shop' solution that:

- Offers deep visibility into the availability, state, and health of desktops and applications in XenDesktop and XenApp (respectively) environments;
- Proactively alerts administrators to real/probable slowdowns in user accesses to desktops/applicaitons;
- Accurately diagnoses and pinpoints the root-cause of the slowdown, thus enabling speedy recovery from problems;
- Helps generate a wide variety of historical reports on performance, which facilitate an effective analysis of current resource capacity of servers/desktops, and precise prediction of future load and resource requirements;

Without eG, administrators may have to shuttle between multiple management tools and engage in manual correlation and analysis of metrics to isolate problem conditions and diagnose their root-cause.

This document deals with eG's *Citrix NetScaler HDX Insight* and *Citrix NetScaler Web Insight* monitoring capabilities.

2

Monitoring Citrix NetScaler HDX Insight

HDX Insight is part of the NetScaler Insight Center and provides end-to-end visibility into the Citrix ICA™ traffic that passes through the NetScaler™ or CloudBridge™ application networking fabric. In the process, HDX Insight delivers business intelligence and failure analysis capabilities for the network, virtual desktops, applications and application fabric.

eG Enterprise integrates with HDX Insight to emerge as a powerful monitoring solution that delivers exceptional user experience by closely tracking ICA sessions to virtual desktops/applications, rapidly isolating latent sessions and the clients/users who launched them, and accurately isolating the root-cause of the latency, much before users even notice any slowness!

Each layer of *eG's Citrix NetScaler HDX Insight* monitoring model has been mapped to tests that have been engineered to make **NITRO API** calls to HDX Insight to pull the metrics that Insight has captured. To run these tests, an eG agent needs to be deployed on a remote Windows/Unix host in the environment. This agent should also be configured with the credentials of a user with read-only permissions to HDX Insight.

Figure 2.1: The layer model of the Citrix HDX Insight component

Using the metrics so collected, administrators can find quick and precise answers for the following performance queries:

- Is any application user consuming too much bandwidth when communicating over ICA?
- Is any application user experiencing a time lag when accessing the application? If so, which user is affected and what could be delaying his/her access – the client side network? the server side network? the server hosting the applications? Or the NetScaler appliance?
- In which session was the application user's experience poor?
- Is any desktop user experiencing slowness during desktop accesses? If so, which user has been impacted and what could be impacting his/her experience with the virtual desktop? In which ICA session was this slowness observed?
- Is any application launching slowly?

- What is the current session load on the XenApp and the XenDesktop environment?
- Did any virtual desktop/application session log out suddenly?
- Is session load evenly distributed across NetScaler gateways or is any NetScaler gateway overloaded with sessions?
- Is any virtual ICA channel consuming more bandwidth than normal?
- Has NetScaler been configured with adequate SSL VPN licenses?

2.1 The HDX Licenses Layer

The test mapped to this layer tracks the license usage of the ICA sessions.

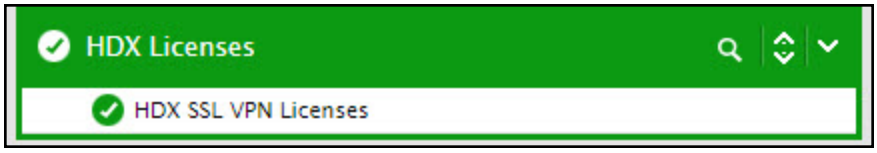

Figure 2.2: The tests mapped to the HDX Licenses layer

2.1.1 HDX SSL VPN Licenses Test

Using this test, administrators can understand how the SSL VPN licenses are being used, and accordingly plan license requirements for the future. Moreover, the test also promptly alerts administrators if the NetScaler appliance is about to run out of licenses.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the infrastructure monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes.**
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total licenses:	Indicates the total number of SSL VPN licenses configured.	Number	
Licenses in use:	Indicates the number of licenses currently utilized by the ICA sessions.	Number	If the number of Licenses in use is equal to or close to the Total licenses count, it is a cause for concern, as it indicates that the environment is running out of licenses.
License usage:	Indicates the percentage of licenses in use.	Percent	A value close to 100% is a cause for concern, as it indicates that the environment is running out of licenses.
Available licenses:	Indicates the number of licenses that are still be used.	Number	A high value is desired for this measure. A value close to 0 implies that not many licenses are available for the use of subsequent ICA sessions. You may want to buy more licenses to pre-empt this unpleasant outcome.

2.2 The HDX Gateways Layer

This layer is mapped to the **HDX Gateways** test, whicb reports load-balancing irregularities in the NetScaler appliances in an environment.

Note:

This layer will not appear if **Citrix NetScaler Insight v10.5** is monitored.

2.2.1 HDX Gateways Test

Multiple NetScaler appliances may be configured in an environment for the purpose of load-balancing. By tracking the ICA sessions and applications launched via every NetScaler appliance, administrators can effortlessly determine whether/not load is uniformly distributed across the appliances. This is what the **HDX Gateways** test helps administrators achieve. This test auto-discovers the NetScaler appliances configured in an environment and measures the session and application load handled by each appliance. In the process, the test indicates how effective the load-balancing algorithm is and whether/not it needs to be tweaked.

Note:

This test will not report metrics for a **Citrix NetScaler HDX Insight** component that is of v10.5 (or above).

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every NetScaler appliance configured in the environment

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes.**
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total session launches:	Indicates the total number of unique user sessions handled by this NetScaler appliance.	Number	Compare the value of this measure across NetScalers to know which NetScaler is overloaded with sessions. Use the detailed diagnosis of this measure to know the name and IP address of each NetScaler appliance configured in the environment and the count of unique sessions and applications that were launched through each.
Total applications :	Indicates the total number of unique appliacates launched via this NetScaler appliance.	Number	Compare the value of this measure across NetScalers to know which NetScaler was used to launch most of the applications.

2.3 The HDX Infrastructure Layer

Using the tests mapped to this layer, administrators can keep tabs on user accesses to virtual desktops and measure the experience of each user with the desktop.

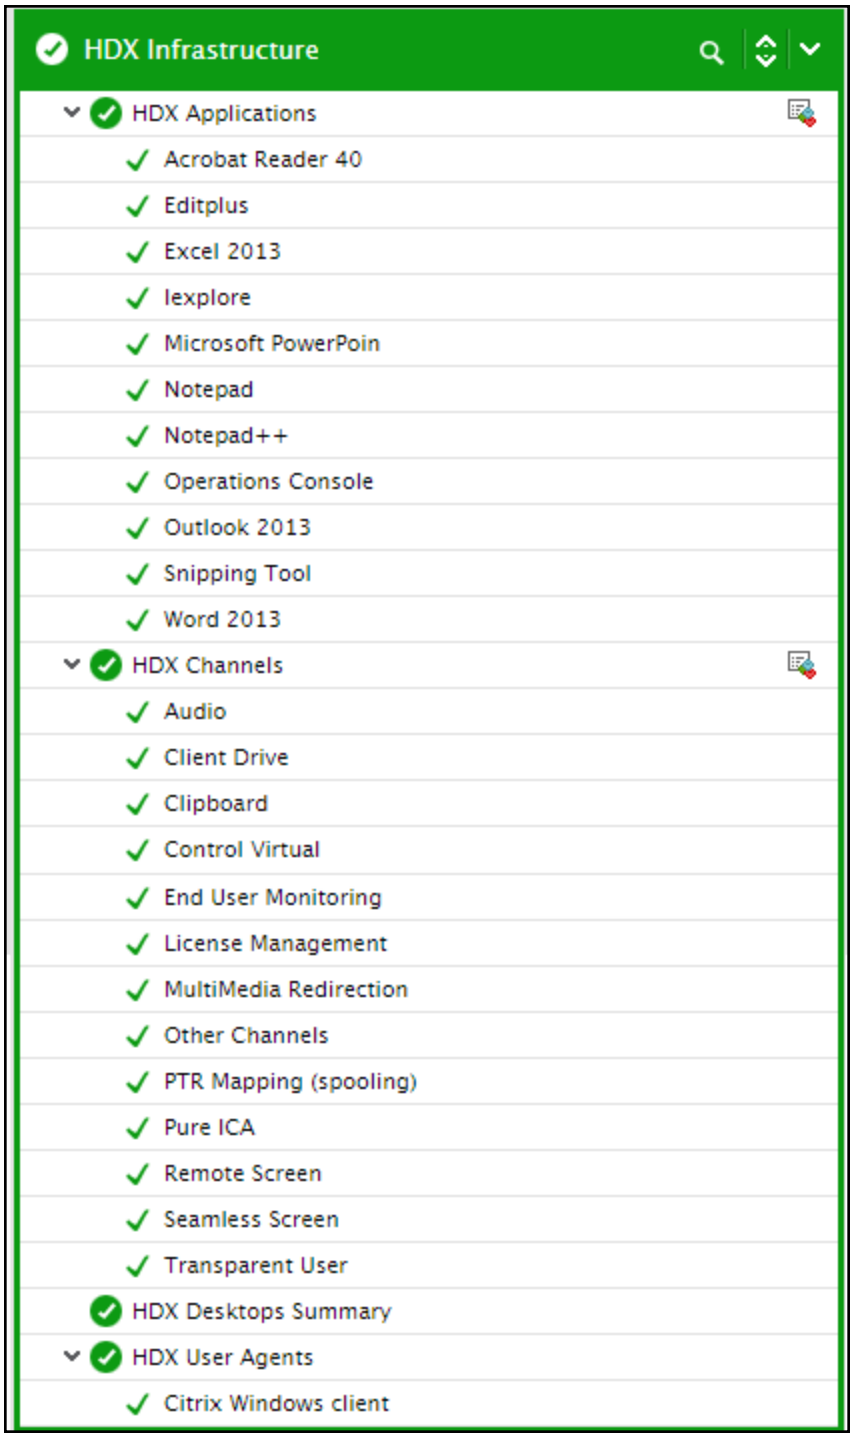

Figure 2.3: The tests mapped to the HDX Infrstructure layer

2.3.1 HDX Applications Test

How much time an application takes to be launched via ICA affects not just the user experience with that application, but also user productivity, as key business transactions may be delayed if business-critical

applications launch slowly. It is hence good practice to periodically run the **HDX Applications** test and check on the launch time of individual applications. This test auto-discovers applications that are operating on XenApp, reports the number of instances of each application that have been launched currently, and the time every application took to launch; this way, the test sheds light on slow-launching applications. Using the detailed diagnosis of this test, you can also figure out how many sessions have been impacted by such applications and how bad the impact is.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every application running on XenApp

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**,**INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Application launches:	Indicates the number of instances of this application that are currently launched.	Number	Compare the value of this measure across applications to identify the most popular application. Use the detailed diagnosis of this measure to know in how many sessions this application was launched, the average launch time of the application, and the average duration of the sessions.
Session launches:	Indicates the number of sessions in which this application was launched.	Number	
Launch duration:	Indicates the average time it took to launch this application.	Msecs	A high value for this measure indicates that the application is slow in launching. You may want to compare the value of this measure across applications to isolate the slowest application.

Use the detailed diagnosis of the *Application launches* measure to know in how many sessions this application was launched, the average launch time of the application, and the average duration of the sessions. In the process, you can quickly identify those applications that are slow in launching.

Details of Application lauches

APPLICATION NAME	APPLICATION LAUNCH COUNT	SESSION LAUNCH COUNT	SESSION DURATION(MSECS)
Jun 01, 2015 10:47:09			
Notepad	1	1	0

Figure 2.4: The detailed diagnosis of the Application launches measure of the HDX Applications test

2.3.2 HDX Channels Test

ICA traffic is composed of virtual channels. A virtual channel is a bidirectional, error-free connection used for the exchange of generalized packet data between a Citrix host for online delivery (XenApp or XenDesktop) and the Citrix Receiver online plug-in. Connections for sound, graphics, client drive mapping, printing, and end user experience monitoring are a few examples of the virtual channels. One of the important motives behind ICA traffic monitoring is to understand how much bandwidth is being consumed by each channel and in the

process identify the most bandwidth-intensive channel. The **HDX Channels** test helps administrators with this. This test monitors the ICA traffic handled by each channel and leads administrators to that channel which is consuming bandwidth excessively. The insights provided by this test help administrators in understanding how the NetScaler appliance needs to be fine-tuned to optimize bandwidth consumption.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every virtual channel

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME, INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Avg bandwidth:	Indicates the rate at which data is transferred over this virtual channel.	Kbps	Ideally, the value of this measure should be low. A high value indicates excessive bandwidth usage by a virtual channel. Compare the value of this measure across channels to know which channel is consuming bandwidth excessively

2.3.3 HDX Desktops Summary

Periodically, administrators need to correlate the session load on their virtual desktop infrastructure with their network throughput to understand whether/not adequate network resources are available to handle the load.

This is exactly what the **HDX Desktops Summary** test helps administrators achieve. This test reports the total number of sessions to the VDI infrastructure and the rate at which ICA traffic generated by those sessions is processed over the network. In the process, the test reveals whether/not sufficient bandwidth is available for handling the current and the anticipated session load. The test additionally reports the number and names of unique users to the virtual desktops, thereby indicating which users are contributing to the session load.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the virtual desktop infrastructure monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**,**INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Avg bandwidth:	Indicates the average rate at which data is transferred over the ICA sessions.	Kbps	Ideally, the value of this measure should be low. A high value is indicative of excessive bandwidth usage.
Unique users:	Indicates the number of distinct users to the virtual desktops.	Number	Use the detailed diagnosis of this measure to know the unique users to the VDI infrastructure.
Total sessions:	Indicates the total number of sessions to the virtual desktops.	Number	This is a good indicator of the total load on the VDI infrastructure.

Note:

This test will report the value `0` for all measures if no desktop sessions exist.

Using the detailed diagnosis of the Unique users measure, you can identify the unique users to the VDI infrastructure.

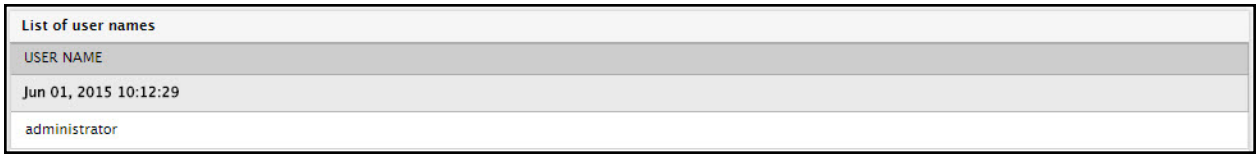

Figure 2.5: The detailed diagnosis of the Unique users measure of the HDX Desktop Summary test

2.3.4 HDX User Agents Test

Users can connect to applications/virtual desktops using may client devices – the ICA client, web browsers, mobile phones, tablets, etc. Each such client device is called a user agent. Users using certain types of client devices/user agents may be engaged in bandwidth-intensive communication over ICA, scarring the experience of other users to the XenApp/XenDesktop environment. To capture such problem conditions quickly, it would be best to know what types of client devices users are connecting from and how much bandwidth each of these device types are currently consuming. This is exactly what the **HDX User Agents** test reveals! This test auto-discovers the types of client devices that are interacting with the applications/virtual desktops, and reports the bandwidth usage of each device type. This way, the test leads administrators to those device types that are consuming bandwidth excessively.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each type of client device/user agent users are connecting from

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Avg bandwidth:	Indicates the rate at which data is transferred over the ICA sessions lauched from this device type.	Mbps	Ideally, the value of this measure should be low. A high value indicates excessive bandwidth usage by users connecting from the device type. Compare the value of this measure across device types to know users using which type of device are engaged in bandwidth-intensive communication over ICA.

2.4 The HDX Users Layer

The tests mapped to this layer track user sessions to applications on XenApp, measure the experience of each user, and reveal those users whose experience is sub-par. The tests also point to areas where users may be spending more time than normal, and thus lead administrators to the probable cause of the poor user experience. In addition, the layer reports the bandwidth usage of ICA channels and the load on NetScaler gateways in use in the environment.

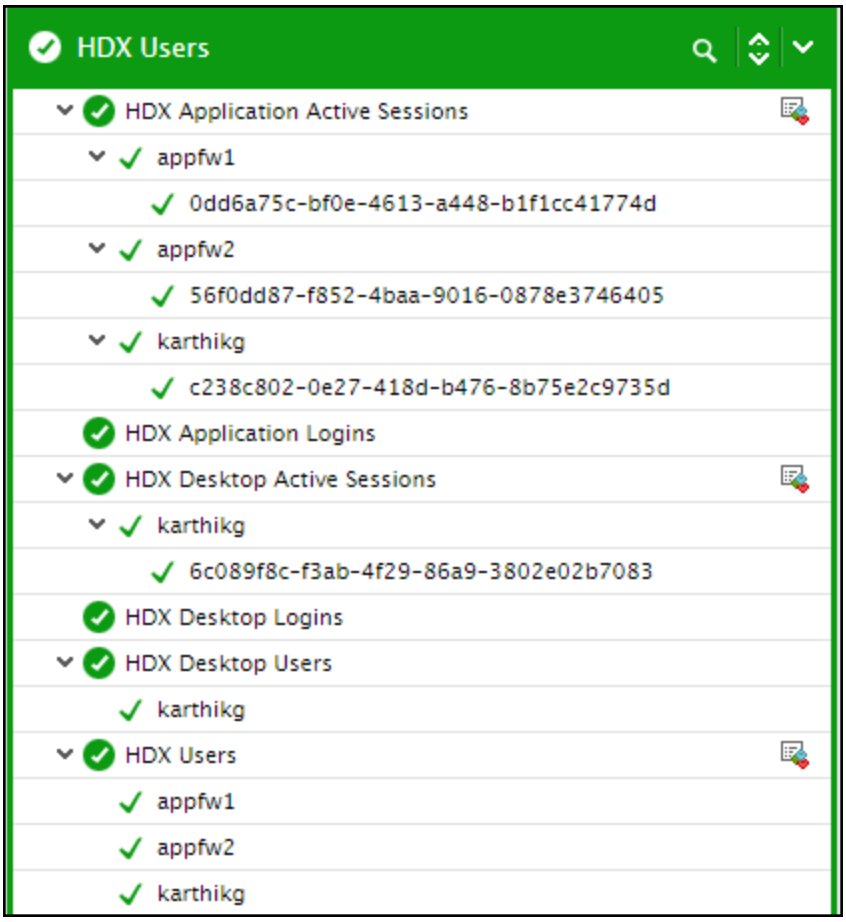

Figure 2.6: The tests mapped to the HDX Users layer

2.4.1 HDX Application User Logins Test

Use this test to determine the current session load on the XenApp servers in your environment. The detailed diagnosis of this test also reveals who launched the sessions, when, from which client, to which server, the session duration, and the overall quality of each user's session. This way, administrators can quickly identify which application user's experience is below-par and what is causing it – a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end? The test also tracks session logouts, and points to abnormal logouts.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the monitored environment

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics**

from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes.

2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions:	Indicates the number of user sessions that are currently active across all XenApp servers.	Number	This is a good indicator of the session load on the XenApp servers. To know which users are contributing to the load, use the detailed diagnosis of this measure. The detailed diagnosis additionally reveals when each user logged in, from which client, to which server, the session duration, and the quality of the user's session. This way,

Measurement	Description	Measurement Unit	Interpretation
			administrators can quickly identify which application user's experience is below-par and what is causing it - a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end?
New logins:	Indicates the number of new logins to the XenApp servers.	Number	A consistent zero value could indicate a connection issue. If this measure reports a non-zero value, use the detailed diagnosis of the measure to know which user logged into which XenApp server, when.
Percent new logins:	Indicates the percentage of current sessions that logged in during the last measurement period.	Percent	
Sessions logging out:	Indicates the number of sessions that logged out.	Number	If all the current sessions suddenly log out, it indicates a problem condition that requires investigation. The detailed diagnosis of this measure lists the sessions that logged out. The quality of each user's session is also revealed. This way, administrators can quickly identify which application user's experience is below-par and what is causing it – a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end?

To know which users are contributing to the load on XenApp, use the detailed diagnosis of the *Current sessions* measure. The detailed diagnosis additionally reveals when each user logged in, from which client, to which server, the session duration, and the quality of the user's session. This way, administrators can quickly identify which application user's experience is below-par and what is causing it - a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end?

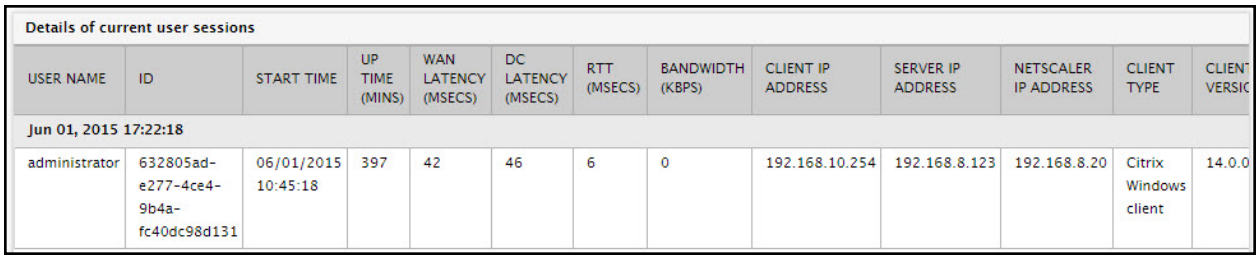

Details of current user sessions

USER NAME	ID	START TIME	UP TIME (MINS)	WAN LATENCY (MSECS)	DC LATENCY (MSECS)	RTT (MSECS)	BANDWIDTH (KBPS)	CLIENT IP ADDRESS	SERVER IP ADDRESS	NETSCALER IP ADDRESS	CLIENT TYPE	CLIENT VERSIC
Jun 01, 2015 17:22:18												
administrator	632805ad-e277-4ce4-9b4a-fc40dc98d131	06/01/2015 10:45:18	397	42	46	6	0	192.168.10.254	192.168.8.123	192.168.8.20	Citrix Windows client	14.0.0

Figure 2.7: The detailed diagnosis of the Current sessions measure of the HDX Application User Logins Test

The detailed diagnosis of the *Sessions logging out* measure lists the sessions that logged out. The quality of each user's session is also revealed. This way, administrators can quickly identify which application user's experience is below-par and what is causing it - a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end?

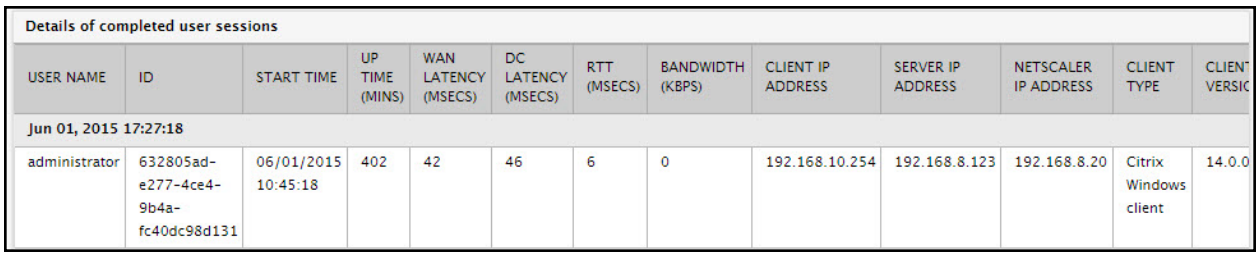

Details of completed user sessions

USER NAME	ID	START TIME	UP TIME (MINS)	WAN LATENCY (MSECS)	DC LATENCY (MSECS)	RTT (MSECS)	BANDWIDTH (KBPS)	CLIENT IP ADDRESS	SERVER IP ADDRESS	NETSCALER IP ADDRESS	CLIENT TYPE	CLIENT VERSIC
Jun 01, 2015 17:27:18												
administrator	632805ad-e277-4ce4-9b4a-fc40dc98d131	06/01/2015 10:45:18	402	42	46	6	0	192.168.10.254	192.168.8.123	192.168.8.20	Citrix Windows client	14.0.0

Figure 2.8: The detailed diagnosis of the Sessions logging out measure of the HDX Application User Logins test

2.4.2 HDX Users Test

To ensure that users are able to access applications/desktops on-demand, administrators must closely track that user's accesses, promptly detect probable access latencies, diagnose its root-cause, and take steps to avert it, well before that user notices and complains. To achieve this, administrators can use the **HDX Users** test. This test automatically discovers the users who are currently accessing applications and virtual desktops in a XenApp/VDI infrastructure, and for each user, reports the latencies that user experienced when interacting with the applications/desktops. This way, the test quickly and accurately points administrators to those users who are experiencing slowdowns, and also leads them to what is causing the slowness – the network? the NetScaler appliance? or the server hosting the applications/desktops?

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every user who is currently connected to an application/virtual desktop

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5**

minutes.

2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Application launches:	Indicates the number of applications currently launched by this user.	Number	Use the detailed diagnosis of this measure to know which applications were launched currently by the user. Compare the value of this measure across users to know which user has launched the maximum number of applications.
Bandwidth:	Indicates the rate at which data is transferred over the ICA sessions of this user.	Kbps	Ideally, the value of this measure should be low. A high value indicates excessive bandwidth usage by the user. Compare the value of this measure across users to know which user is

Measurement	Description	Measurement Unit	Interpretation
			consuming bandwidth excessively
WAN latency:	Indicates the average latency experienced by this user due to problems with the client side network.	msecs	A high value for this measure indicates that the client side network is slow. If the value of the Round trip time – RTT measure is abnormally high for a user, you can compare the value of this measure with that of the DC latency, Client jitter, and Server jitter measures of that user to know what is causing the slowness – is it the client side network? or the server side network?
DC latency:	Indicates the average latency experienced by this user due to problems with the server side network.	msecs	A high value for this measure indicates that the server side network is slow. If the value of the *Round trip time – RTT* measure is abnormally high for a user, you can compare the value of this measure with that of the *WAN latency, Client side device delay, Server side device delay*, and *Host delay* measures of that user to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the applications/desktops?
Round trip time - RTT:	Indicates the screen lag experienced by this user while interacting with an application/desktop.	msecs	A high value for this measure is indicative of the poor quality of that user's experience with applications. To know the reason for this below-par UX, compare the value of the *WAN latency, DC latency, Client jitter*, and *Server jitter* measures of that user.
Client smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this user smoothed over the	msecs	TCP implementations attempt to predict future round-trip times by sampling the

Measurement	Description	Measurement Unit	Interpretation
	client side connection.		behavior of packets sent over a connection and averaging those samples into a ‘‘smoothed’’ round-trip time estimate, SRTT. When a packet is sent over a TCP connection, the sender times how long it takes for it to be acknowledged, producing a sequence, S, of round-trip time samples: s1, s2, s3…. With each new sample, si, the new SRTT is computed from the formula: $SRTTi+1 = (\alpha \times SRTTi) + (1 - \alpha) \times si$ Here, SRTTi is the current estimate of the round-trip time, SRTTi+1 is the new computed value, and α is a constant between 0 and 1 that controls how rapidly the SRTT adapts to change. The retransmission time-out (RTOi), the amount of time the sender will wait for a given packet to be acknowledged, is computed from SRTTi. The formula is: $RTOi = \beta \times SRTTi$ Here, β is a constant, greater than 1, chosen such that there is an acceptably small probability that the round-trip time for the packet will exceed RTOi.
Server smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this user, smoothed over the server side connection.	msecs	
Client jitter:	Indicates the client side jitter.	msecs	Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

Measurement	Description	Measurement Unit	Interpretation
			A high value for these measures therefore is indicative of a long time gap between ICA packets. To know where the delay is longer – whether on the client side or on the server side - compare the value of the Client jitter measure with that of the Server jitter measure. Also, if the value of the *Round trip time – RTT* measure is abnormally high for a user, then you can compare the values of these measures with that of the *WAN latency* and *DC latency* measures to know what is causing the problem – the client side network? or the server side network?
Server jitter:	Indicates the server side jitter.	msecs	

Use the detailed diagnosis of the *Application launches* measure to know which applications were launched currently by the user.

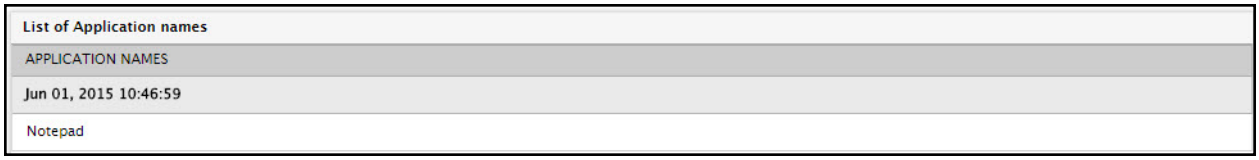

Figure 2.9: The detailed diagnosis of the Application launches measure of the HDX Application Users test

2.4.3 HDX Application Active Sessions Test

In order to ensure that the user experience with applications deployed on XenApp remains 'superlative' at all times, administrators should be able to proactively detect potential slowdowns when accessing applications, precisely pinpoint the user session affected by the slowdown, accurately isolate the root-cause of such slowness, and rapidly initiate measures to eliminate the root-cause. The **HDX Application Active Sessions** test facilitates all the above, and thus assures users of uninterrupted application access! For every user session that is currently active on the XenApp servers in an environment, this test monitors the ICA traffic generated by that session, and promptly reports latencies. This way, the test accurately points to the exact session where a user experienced slowness. Also, in the event of abnormally high latencies, the test also leads administrators to where that user's session is bottlenecked - the network? the NetScaler appliance? Or the server hosting the applications?

This test is disabled by default. To enable the test, first select the **Enable/Disable** option from the **Tests** menu in the **Agents** tile. Then, pick **Citrix NetScaler HDX Insight** as the **Component type**, **Performance** as the **Test type**, select this test from the **DISABLED TESTS** list and click **<** to enable it.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every active session of each user who is currently connected to the XenApp environment

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Bandwidth:	Indicates the rate at which data is transferred over this ICA session.	Kbps	Ideally, the value of this measure should be low. A high value indicates excessive bandwidth usage by the session. Compare the value of this measure across sessions to know which session is consuming bandwidth excessively. You can also use the detailed diagnosis of this measure to view the complete details of the session in question. The user who launched the session, when the session was launched, the session duration, the client from which the session was initiated, the virtual desktop that was accessed during the session, and the server on which this virtual desktop exists, are revealed.
WAN latency:	Indicates the average latency experienced by this user session due to problems with the client side network.	Msecs	A high value for this measure indicates that the client side network is slow. If the value of the *Round trip time – RTT* measure is abnormally high for a session, you can compare the value of this measure with that of the *DC latency, Client side device delay, Server side device delay, Host delay, Client L7 latency*, and *Server L7 latency* measures of that user session to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the applications?
DC latency:	Indicates the average latency experienced by this session due to problems with the server	Msecs	A high value for this measure indicates that the server side network is slow. If the value of the *Round trip time – RTT*

Measurement	Description	Measurement Unit	Interpretation
	side network.		measure is abnormally high for a session, you can compare the value of this measure with that of the *WAN latency, Client side device delay, Server side device delay, Host delay, Client L7 latency*, and *Server L7 latency* measures of that session to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?
Client side device delay:	Indicates the average latency experienced by this session, which was caused by the NetScaler appliance when ICA traffic flowed from client network to server network.	Msecs	A high value for these measures indicates a processing bottleneck with the NetScaler appliance. If the value of the *Round trip time – RTT* measure is abnormally high for a session, you can compare the value of these measures with that of the *WAN latency, DC latency, Host delay, Client L7 latency*, and *Server L7 latency* measures of that user to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?
Server side device delay:	Indicates the average latency experienced by this session, which was caused by the NetScaler appliance when ICA traffic flowed from server network to client network.	Msecs	
Host delay:	Indicates the delay that this session experienced when waiting for the host to process the packets.	Msecs	A high value for this measure indicates a processing bottleneck with the server hosting the applications. If the value of the *Round trip time – RTT* measure is abnormally high for a session, you can compare the value of this measure with that of the *WAN latency, DC latency, Host delay, Client L7 latency*, and *Server L7 latency* measures to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting

Measurement	Description	Measurement Unit	Interpretation
			the desktops?
Round trip time - RTT:	Indicates the screen lag experienced by this session while interacting with applications hosted on XenApp.	Msecs	A high value for this measure is indicative of the poor quality of a user's experience with applications on XenApp. To know the reason for this below-par UX, compare the value of the *WAN latency, DC latency, Client side device delay, Server side device delay, Host delay, Client L7 latency*, and *Server L7 latency* measures of that user.
Client smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this session smoothed over the client side connection.	MSecs	TCP implementations attempt to predict future round-trip times by sampling the behavior of packets sent over a connection and averaging those samples into a ''smoothed'' round-trip time estimate, SRTT. When a packet is sent over a TCP connection, the sender times how long it takes for it to be acknowledged, producing a sequence, S, of round-trip time samples: s1, s2, s3.... With each new sample, si, the new SRTT is computed from the formula: $SRTTi+1 = (\alpha \times SRTTi) + (1 - \alpha) \times si$ Here, SRTTi is the current estimate of the round-trip time, SRTTi+1 is the new computed value, and α is a constant between 0 and 1 that controls how rapidly the SRTT adapts to change. The retransmission time-out (RTOi), the amount of time the sender will wait for a given packet to be acknowledged, is computed from SRTTi. The formula is: $RTOi = \beta \times SRTTi$ Here, β is a constant, greater than 1, chosen such that there is an acceptably

Measurement	Description	Measurement Unit	Interpretation
			small probability that the round-trip time for the packet will exceed RTOi.
Server smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this session, smoothed over the server side connection.	MSecs	
Client side 0 window count:	Indicates how many times in this session the client advertised a zero TCP window during the last measurement period.	Number	TCP Zero Window is when the Window size in a machine remains at zero for a specified amount of time. TCP Window size is the amount of information that a machine can receive during a TCP session and still be able to process the data. Think of it like a TCP receive buffer. When a machine initiates a TCP connection to a server, it will let the server know how much data it can receive by the Window Size. In many Windows machines, this value is around 64512 bytes. As the TCP session is initiated and the server begins sending data, the client will decrement it's Window Size as this buffer fills. At the same time, the client is processing the data in the buffer, and is emptying it, making room for more data. Through TCP ACK frames, the client informs the server of how much room is in this buffer. If the TCP Window Size goes down to 0, the client will not be able to receive any more data until it processes and opens the buffer up again. The machine (client/server) alerting the Zero Window will not receive any more data from the host. This is why, ideally, the value of these measures should be 0. A non-zero value warrants an immediate investigation to determine the reason for

Measurement	Description	Measurement Unit	Interpretation
			the Zero Window. It could be that the client/server was running too many processes at that moment, and its processor is maxed. Or it could be that there is an error in the TCP receiver, like a Windows registry misconfiguration. Try to determine what the client was doing when the TCP Zero Window happened.
Server side 0 window count:	Indicates how many times in this session the server advertised a zero TCP window during the last measurement period.	Number	
Client side retransmit timeout – RTO:	Indicates how many times during the last measurement period the retransmit timeout got invoked in this session on the client side connection.	Number	An RTO occurs when the sender is missing too many acknowledgments and decides to take a time out and stop sending altogether. After some amount of time, usually at least one second, the sender cautiously starts sending again, testing the waters with just one packet at first, then two packets, and so on. As a result, an RTO causes, at minimum, a one-second delay on your network. A low value is hence desired for these measures.
Server side retransmit timeout - RTO:	Indicates how many times during the last measurement period the retransmit timeout got invoked in this session on the server side connection.	Number	
Client side retransmits:	Indicates the number of packets retransmitted on the client side connection during the last measurement period.	Number	Ideally, the value of these measures should be 0.
Server side retransmits:	Indicates the number of packets retransmitted on the server side connection during the last measurement period.	Number	
Client L7 latency:	Indicates the latency measured at the seventh layer (application layer) of	Msecs	The OSI, or Open System Interconnection, model defines a networking framework to implement protocols in seven layers.

Measurement	Description	Measurement Unit	Interpretation
	the OSI model using ICA probes and responses sent between Receiver and the Host, on client side pcb.		The seventh layer supports application and end- user processes. Communication partners are identified, quality of service is identified, user authentication and privacy are considered, and any constraints on data syntax are identified at this layer. This layer provides application services for file transfers, e- mail, and other network software services.
Server L7 latency:	Indicates the latency measured at the seventh layer (application layer) of the OSI model using ICA probes and responses sent between Receiver and the Host, on server side pcb.	Msecs	A high value for this measure indicates a processing bottleneck at the application layer. If the value of the Round trip time – RTT measure is abnormally high for a session, you can compare the value of these measures with that of the WAN latency, DC latency, and Host delay to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?

You can also use the detailed diagnosis of the *Bandwidth* measure to view the complete details of the session in question. The user who launched the session, when the session was launched, the session duration, the client from which the session was initiated, the server with which the session connected, the type of client device that was used, and the client version are revealed.

Details of users

USER NAME	ID	START TIME	UP TIME (MINS)	CLIENT SIDE CB	SERVER SIDE CB	CLIENT IP ADDRESS	SERVER IP ADDRESS	NETSCALER IP ADDRESS	CLIENT TYPE	CLIENT VERSION
Jun 01, 2015 17:20:43										
administrator	632805ad-e277-4ce4-9b4a-fc40dc98d131	06/01/2015 10:43:43	397	No	No	192.168.10.254	192.168.8.123	192.168.8.20	Citrix Windows client	14.0.0.91

Figure 2.10: The detailed diagnosis of the Bandwidth measure of the HDX Application Active Sessions test

2.4.4 HDX Desktop Logins Test

Use this test to determine the current session load on the virtual desktops in your infrastructure. The detailed diagnosis of this test also reveals who launched the sessions, when, from which client, to which server, the session duration, and the overall quality of each user's session. This way, administrators can quickly identify

which user's experience with virtual desktops is below-par and what is causing it – a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end? The test also tracks session logouts, and points to abnormal logouts.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the virtual desktop infrastructure monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions:	Indicates the number of user sessions that are currently active across all virtual desktops.	Number	This is a good indicator of the session load on the virtual desktops. To know which users are contributing to the load, use the detailed diagnosis of this measure. The detailed diagnosis additionally reveals when each user logged in, from which client, to which server, the session duration, and the quality of the user's session. This way, administrators can quickly identify which user's experience with virtual desktops is below-par and what is causing it – a flaky network connection? a bandwidth-intensive communication over the network? or an authentication delay with the server hosting desktops?
New logins:	Indicates the number of new logins to the virtual desktops.	Number	A consistent zero value could indicate a connection issue. If this measure reports a non-zero value, use the detailed diagnosis of the measure to know which user logged into which virtual desktop, when.
Percent new logins:	Indicates the percentage of current sessions that logged in during the last measurement period.	Percent	
Sessions logging out:	Indicates the number of sessions that logged out.	Number	If all the current sessions suddenly log out, it indicates a problem condition that requires investigation. The detailed diagnosis of this measure lists the sessions that logged out. The quality of each user's session is also revealed. This way, administrators can

Measurement	Description	Measurement Unit	Interpretation
			quickly identify which user's experience with virtual desktops is below-par and what is causing it – a flaky network connection? a bandwidth- intensive communication over the network? or a delay at the server-end?

To know which users are contributing to the load on the desktops, use the detailed diagnosis of the *Current sessions* measure. The detailed diagnosis additionally reveals when each user logged in, from which client, to which server, the session duration, and the quality of the user's session. This way, administrators can quickly identify which user's experience with virtual desktops is below-par and what is causing it – a flaky network connection? a bandwidth-intensive communication over the network? or an authentication delay with the server hosting desktops?

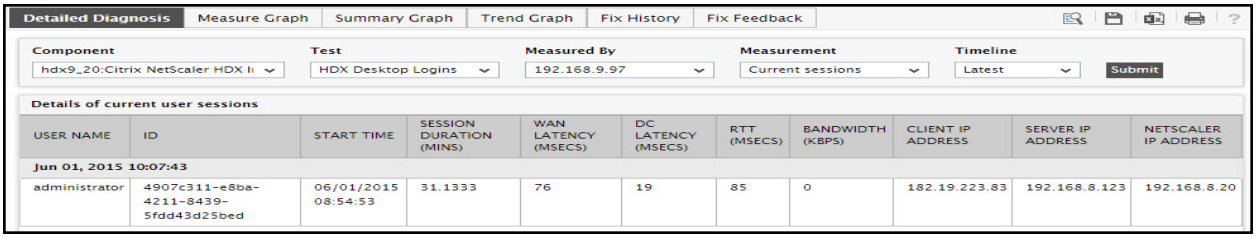

Figure 2.11: The detailed diagnosis of the Current sessions measure of the HDX Desktop Logins test

The detailed diagnosis of the *Sessions logging out* measure lists the sessions that logged out. The quality of each user's session is also revealed. This way, administrators can quickly identify which user's experience with virtual desktops is below-par and what is causing it – a flaky network connection? a bandwidth-intensive communication over the network? or a delay at the server-end?

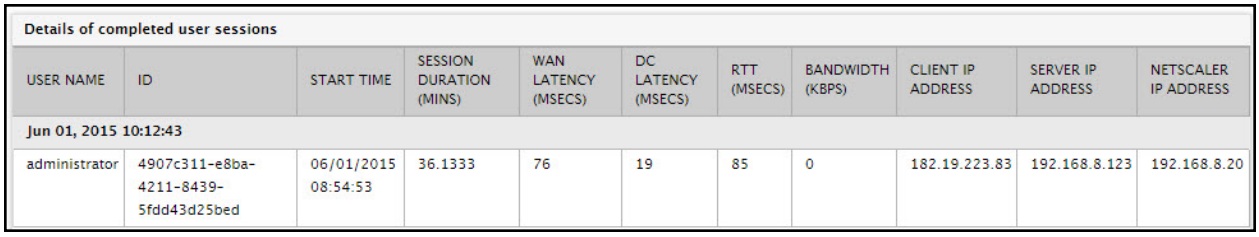

Figure 2.12: The detailed diagnosis of the Sessions logging out measure of the HDX Desktop Logins test

2.4.5 HDX Desktop Users Test

To ensure a high quality virtual desktop experience for a user, administrators must closely track that user's desktop accesses, promptly detect probable access latencies, diagnose its root-cause, and take steps to avert it, well before that user notices and complains. To achieve this, administrators can use the **HDX**

Desktop Users test. This test automatically discovers the users who are currently logged into the VDI infrastructure, and for each user, reports the latencies that user experienced when interacting with virtual desktops. This way, the test quickly and accurately points administrators to those users who are experiencing slowdowns, and also leads them to what is causing the slowness – the network? the NetScaler appliance? or the server hosting the desktops?

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every user who is currently connected to the VDI infrastructure

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Desktop launches:	Indicates the number of desktop sessions currently launched by this user.	Number	Use the detailed diagnosis of this measure to know which desktops are currently accessed by the user. Compare the value of this measure across users to know which user has launched the maximum desktop sessions, thus overloading the infrastructure.
Bandwidth:	Indicates the rate at which data is transferred over the ICA sessions of this user.	Kbps	Ideally, the value of this measure should be low. A high value indicates excessive bandwidth usage by the user. Compare the value of this measure across users to know which user is consuming bandwidth excessively
WAN latency:	Indicates the average latency experienced by this user due to problems with the client side network.	msecs	A high value for this measure indicates that the client side network is slow. If the value of the *Round trip time – RTT* measure is abnormally high for any user, you can compare the value of this measure with that of the DC latency measure of that user to know what is causing the slowness – is it the client side network? or the server side network?
DC latency:	Indicates the average latency experienced by this user due to problems with the server side network.	msecs	A high value for this measure indicates that the server side network is slow. If the value of the *Round trip time – RTT* measure is abnormally high for any user, you can compare the value of this measure with that of the WAN latency measure of that user to know what is

Measurement	Description	Measurement Unit	Interpretation
			causing the slowness – is it the client side network? or the server side network?
Round trip time - RTT:	Indicates the screen lag experienced by this user while interacting with desktops.	msecs	A high value for this measure is indicative of the poor quality of that user's experience with virtual desktops. To know the reason for this below-par UX, compare the value of the WAN latency and DC latency measures of that user.
Client smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this user smoothed over the client side connection.	msecs	TCP implementations attempt to predict future round-trip times by sampling the behavior of packets sent over a connection and averaging those samples into a ''smoothed'' round-trip time estimate, SRTT. When a packet is sent over a TCP connection, the sender times how long it takes for it to be acknowledged, producing a sequence, S, of round-trip time samples: s1, s2, s3.... With each new sample, si, the new SRTT is computed from the formula: $SRTTi+1 = (\alpha \times SRTTi) + (1 - \alpha) \times si$ Here, SRTTi is the current estimate of the round-trip time, SRTTi+1 is the new computed value, and α is a constant between 0 and 1 that controls how rapidly the SRTT adapts to change. The retransmission time-out (RTOi), the amount of time the sender will wait for a given packet to be acknowledged, is computed from SRTTi. The formula is: $RTOi = \beta \times SRTTi$ Here, β is a constant, greater than 1,

Measurement	Description	Measurement Unit	Interpretation
			chosen such that there is an acceptably small probability that the round-trip time for the packet will exceed RTOi.
Server smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this user, smoothed over the server side connection.	msecs	
Session duration:	Indicates the total duration of all sessions of this user.	Mins	Compare the value of this measure across users to know which user has interacted with the virtual desktops for the longest time.

Use the detailed diagnosis of the *Desktop launches* measure to know which desktops are currently accessed by the user.

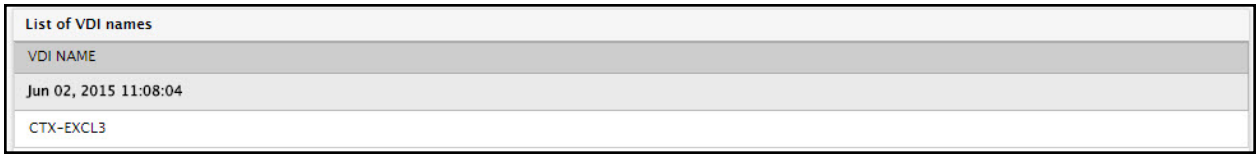

Figure 2.13: The detailed diagnosis of the Desktop launches measure of the HDX Desktop Users test

2.4.6 HDX Desktop Active Sessions Test

The key to ensuring a more than satisfactory user experience with virtual desktops lies in proactively detecting potential slowdowns when accessing virtual desktops, precisely pinpointing the user session affected by the slowdown, accurately isolating the root-cause of such slowness, and rapidly initiating measures to eliminate the root-cause. The **HDX Desktop Active Sessions** test facilitates all the above, and thus assures users of uninterrupted desktop access! For every user session that is currently active on desktops, this test monitors the ICA traffic generated by that session, and promptly reports latencies. This way, the test accurately points to the exact session where a user experienced slowness. Also, in the event of abnormally high latencies, the test also leads administrators to where that user's session is bottlenecked - the network? the NetScaler appliance? Or the server hosting the virtual desktop?

This test is disabled by default. To enable the test, first select the **Enable/Disable** option from the **Tests** menu in the **Agents** tile. Then, pick **Citrix NetScaler HDX Insight** as the **Component type**, **Performance** as the **Test type**, select this test from the **DISABLED TESTS** list and click **<** to enable it.

Target of the test : Citrix NetScaler HDX Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every active session of each user who is currently connected to the VDI infrastructure

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from HDX Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME,INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to HDX Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to HDX Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, HDX Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Bandwidth:	Indicates the rate at which data is transferred over this ICA session.	Kbps	Ideally, the value of this measure should be low.

Measurement	Description	Measurement Unit	Interpretation
			A high value indicates excessive bandwidth usage by the session. Compare the value of this measure across sessions to know which session is consuming bandwidth excessively. You can also use the detailed diagnosis of this measure to view the complete details of the session in question. The user who launched the session, when the session was launched, the session duration, the client from which the session was initiated, the virtual desktop that was accessed during the session, and the server on which this virtual desktop exists, are revealed.
WAN latency:	Indicates the average latency experienced by this user session due to problems with the client side network.	msecs	A high value for this measure indicates that the client side network is slow. If the value of the *Round trip time – RTT* measure is abnormally high for any session, you can compare the value of this measure with that of the *DC latency, Client side device delay, Server side device delay* , and *Host delay* measures of that session to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?
DC latency:	Indicates the average latency experienced by this session due to problems with the server side network.	msecs	A high value for this measure indicates that the server side network is slow. If the value of the *Round trip time – RTT* measure is abnormally high for any session, you can compare the value of this measure with that of the *WAN latency, Client side device delay, Server*

Measurement	Description	Measurement Unit	Interpretation
			side device delay, and *Host delay* measures of that session to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?
Client side device delay:	Indicates the average latency experienced by this session, which was caused by the NetScaler appliance when ICA traffic flowed from client network to server network.	Millisecs	A high value for these measures indicates a processing bottleneck with the NetScaler appliance. If the value of the *Round trip time – RTT* measure is abnormally high for any session, you can compare the value of these measures with that of the *WAN latency, DC latency*, and *Host delay* measures of that session to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?
Server side device delay:	Indicates the average latency experienced by this session, which was caused by the NetScaler appliance when ICA traffic flowed from server network to client network.	Millisecs	
Host delay:	Indicates the average delay in ICA traffic experienced by this session, which was caused by the server network.	Msecs	A high value for this measure indicates a processing bottleneck with the server hosting the desktops. If the value of the *Round trip time – RTT* measure is abnormally high for any session, you can compare the value of this measure with that of the *WAN latency, DC latency, Client side device delay, Server side device delay*, and *Host delay* measures to know what is causing the slowness – is it the client side network? the server side network? the NetScaler appliance? or the server hosting the desktops?
Round trip time -	Indicates the screen lag	Msecs	A high value for this measure is

Measurement	Description	Measurement Unit	Interpretation
RTT:	experienced by this session while interacting with desktops.		indicative of the poor quality of a user's experience with virtual desktops. To know the reason for this below-par UX, compare the value of the *WAN latency, DC latency, Client side device delay, Server side device delay*, and *Host delay* measures of that user session.
Client smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this session smoothed over the client side connection.	Msecs	TCP implementations attempt to predict future round-trip times by sampling the behavior of packets sent over a connection and averaging those samples into a ''smoothed'' round-trip time estimate, SRTT. When a packet is sent over a TCP connection, the sender times how long it takes for it to be acknowledged, producing a sequence, S, of round-trip time samples: s1, s2, s3.... With each new sample, si, the new SRTT is computed from the formula: $SRTTi+1 = (\alpha \times SRTTi) + (1 - \alpha) \times si$ Here, SRTTi is the current estimate of the round-trip time, SRTTi+1 is the new computed value, and α is a constant between 0 and 1 that controls how rapidly the SRTT adapts to change. The retransmission time-out (RTOi), the amount of time the sender will wait for a given packet to be acknowledged, is computed from SRTTi. The formula is: $RTOi = \beta \times SRTTi$ Here, β is a constant, greater than 1, chosen such that there is an acceptably small probability that the round-trip time for the packet will exceed RTOi.
Server smoothed round trip time - SRTT:	Indicates the RTT (round-trip time or screen lag time) of this session, smoothed over the server side connection.	msecs	

Measurement	Description	Measurement Unit	Interpretation
Client side 0 window count:	Indicates how many times in this session the client advertised a zero TCP window during the last measurement period.	Number	TCP Zero Window is when the Window size in a machine remains at zero for a specified amount of time. TCP Window size is the amount of information that a machine can receive during a TCP session and still be able to process the data. Think of it like a TCP receive buffer. When a machine initiates a TCP connection to a server, it will let the server know how much data it can receive by the Window Size. In many Windows machines, this value is around 64512 bytes. As the TCP session is initiated and the server begins sending data, the client will decrement it's Window Size as this buffer fills. At the same time, the client is processing the data in the buffer, and is emptying it, making room for more data. Through TCP ACK frames, the client informs the server of how much room is in this buffer. If the TCP Window Size goes down to 0, the client will not be able to receive any more data until it processes and opens the buffer up again. The machine (client/server) alerting the Zero Window will not receive any more data from the host. This is why, ideally, the value of these measures should be 0. A non-zero value warrants an immediate investigation to determine the reason for the Zero Window. It could be that the client/server was running too many processes at that moment, and its processor is maxed. Or it could be that there is an error in the TCP receiver, like a Windows registry misconfiguration.

Measurement	Description	Measurement Unit	Interpretation
			Try to determine what the client/server was doing when the TCP Zero Window happened.
Server side 0 window count:	Indicates how many times in this session the server advertised a zero TCP window during the last measurement period.	Number	
Client retransmit timeout - RTO:	Indicates how many times during the last measurement period the retransmit timeout got invoked in this session on the client side connection.	Number	An RTO occurs when the sender is missing too many acknowledgments and decides to take a time out and stop sending altogether. After some amount of time, usually at least one second, the sender cautiously starts sending again, testing the waters with just one packet at first, then two packets, and so on. As a result, an RTO causes, at minimum, a one-second delay on your network. A low value is hence desired for these measures.
Server retransmit timeout - RTO:	Indicates how many times during the last measurement period the retransmit timeout got invoked in this session on the server side connection.	Number	

You can also use the detailed diagnosis of the **`Bandwidth`** measure to view the complete details of the session in question. The user who launched the session, when the session was launched, the session duration, the client from which the session was initiated, the virtual desktop that was accessed during the session, and the server on which this virtual desktop exists, are revealed.

Details of users

USER NAME	ID	START TIME	UP TIME (MINS)	CLIENT SIDE CB	SERVER SIDE CB	CLIENT IP ADDRESS	SERVER IP ADDRESS	NETSCALER IP ADDRESS	CLIENT TYPE	CLIENT VERSION
Jun 01, 2015 17:20:43										
administrator	632805ad-e277-4ce4-9b4a-fc40dc98d131	06/01/2015 10:43:43	397	No	No	192.168.10.254	192.168.8.123	192.168.8.20	Citrix Windows client	14.0.0.91

Figure 2.14: The detailed diagnosis of the Bandwidth measure of the HDX Desktop Active Sessions test

3

Monitoring Citrix NetScaler Web Insight

Web Insight provides visibility into web applications and allows IT administrators to monitor all web applications being served by NetScaler ADCs. Web Insight captures data about web traffic that flows between the clients and the servers, generates AppFlow records by doing deep inspection of the data, and presents the records as visual reports.

eG Enterprise integrates with Web Insight and collects all the web traffic-related performance data captured by Web Insight. This easy integration capability, when coupled with eG's built-in web monitoring capabilities, transforms eG Enterprise into a comprehensive monitoring solution that provides unlimited visibility into the performance and problems related to web applications.

To visually represent the data collected from Web Insight in the eG monitoring console, eG Enterprise provides a specialized *Citrix NetScaler Web Insight* monitoring model.

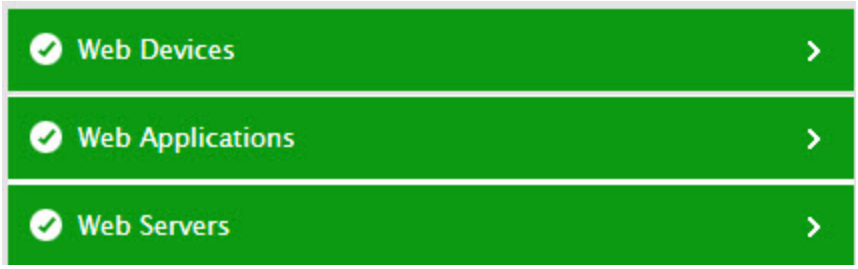

Figure 3.1: Layer model of Citrix NetScaler Web Insight

Note:

eG Enterprise cannot monitor **Citrix NetScaler Insight v10.5**.

To run these tests, the eG agent will have to be deployed on a remote Windows/Unix host in the environment and should be configured with the credentials of a user with *read-only* privileges to Web Insight.

Using the metrics collected, the following performance queries can be answered:

- Is any web server overloaded with requests? If so, which one is it?
- Are all web servers processing requests quickly, or is any web server experiencing a processing bottlenck?
- Is any web server's network very latent?
- Which is the most popular web application?
- Is any web application consuming too much bandwidth?

- Is any web application responding slowly to requests? If so, what is the root-cause of the poor responsiveness - is it owing to the poor processing power of the web server? a latent server network? or a slow client network?
- Is request load uniformly balanced across all NetScaler appliances in use in the environment?
- Is any NetScaler handling bandwidth-intensive web traffic? If so, which HTTP request methods to and response status messages from that NetScaler are hogging the bandwidth resources?
- Is any client overloading the web servers managed by NetScaler ADC?
- Is any client rendering pages slowly?
- Is any client's network slow?
- Which client operating system is consuming bandwidth abnormally?
- Which user agent is a bandwidth hog?
- Which URL request generated a very slow response? Why? - is it because of a delay in page loading? Or page rendering?

3.1 The Web Servers Layer

The test mapped to this layer tracks requests to web servers managed by NetScaler and measures how well/badly each server processes the requests it receives.

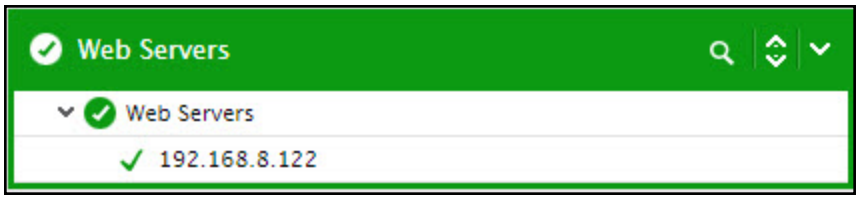

Figure 3.2: The test mapped to the Web Servers layer

3.1.1 Web Servers Test

If users accessing a web server complain of slowness, administrators must be able to quickly figure out what is causing the slowness – is it because of a processing bottleneck with the web server? or is it owing to a latent server network? The **Web Servers** test accurately points administrators to the source of the slowness! This test tracks requests to each web server managed by NetScaler ADC and reports the time every server takes to process the requests. The test thus sends out proactive alerts to administrators if it finds that any web server is responding very slowly to client requests. Additionally, the test also indicates if the slowdown experienced by the user can be attributed to a latent server-side network. This way, the test helps administrators identify slow servers and rapidly isolate the reason for the slowness, so that the problem can be fixed quickly and normalcy restored in no time.

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every web server managed by NetScaler ADC

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received by this web server.	Number	This is a good indicator of the load on the web server. Compare the value of this measure across web servers to know which server is receiving the maximum

Measurement	Description	Measurement Unit	Interpretation
			number of requests. If a single server appears to be servicing a significantly large number of requests than the rest, it could imply that the server is overloaded. This in turn indicates that a faulty/ineffective load- balancing algorithm is in use. Use the detailed diagnosis of this measure to know which NetScaler is managing this web server, so that the load-balancing logic of that NetScaler appliance can be reconfigured to avoid overloads.
Avg bandwidth:	Indicates the total amount of data received by this web server.	KB	Compare the value of this measure across web servers to know which server is consuming bandwidth excessively.
Server processing time:	Indicates the elapsed time, from when the server starts to receive the first byte of a request from the NetScaler appliance until the NetScaler appliance receives the first byte to response.	msecs	A high value for this measure indicates that the web server is processing requests slowly. Compare the value of this measure across web servers to isolate the slowest web server. In the event that a user complains of slowness, you can compare the value of this measure with that of the Server network latency measure to determine what is causing the slowness – the poor processing power of the web server? or a latent server network?
Server network latency:	Indicates the average latency caused by the server network.	msecs	A high value for this measure indicates that the server network is latent. Compare the value of this measure across web servers to know which server's network is the slowest.

Measurement	Description	Measurement Unit	Interpretation
			In the event that a user complains of slowness, you can compare the value of this measure with that of the Server processing time measure to determine what is causing the slowness – the poor processing power of the web server? or a latent server network?
Cache hits:	Indicates the number of requests to this web server serviced by the cache.	Number	If the value of this measure is the same as that of the Hits measure, it implies that the web server has serviced all requests to it using the cache. This is indicative of optimal cache size and usage. On the other hand, if the value of this measure is much lower than that of the Hits measure, it could indicate improper cache sizing and ineffective cache usage.
Cache miss:	Indicates the number of requests to this web server that were not serviced by the cache.	Number	Ideally, the value of this measure should be 0 or at least, very low. If the value is the same as that of the Hits measure, it could indicate improper cache sizing and ineffective cache usage.
Cache hit ratio:	Indicates the percentage of requests to this web server that were serviced by the cache.	Percent	Ideally, the value of this measure should be > 80%. A low hit ratio on the other hand indicates that a majority of web requests were serviced by the origin server and not the cache server. This can significantly increase request processing time and related overheads.
Cache bypass:	Indicates the number of requests that were serviced by the origin server, because the cache server was bypassed.	Number	

Measurement	Description	Measurement Unit	Interpretation
Cache hits bandwidth consumed:	Indicates the bandwidth consumed when requests to this web server were serviced by the cache server.	KB	The difference between the value of the Avg bandwidth measure and this measure for a web server will reveal the bandwidth that may have been saved by request caching. Where cache is well-sized and used optimally, this difference will be high. You can compare this difference across web servers to know which web server is making the most effective use of the cache.
Cache misses bandwidth consumed:	Indicates the bandwidth consumed when the cache server could not serve the requests to this web server.	KB	The difference between the value of this measure and the value of the Cache hits bandwidth consumed measure for a web server will reveal how much bandwidth was saved by cache hits.
Cache bypass bandwidth consumed:	Indicates the bandwidth consumed when the cache server was bypassed and the request was served from this origin server.	KB	If the difference between the value of this measure and that of the Cache hits bandwidth consumed measure results in a 'positive' integer, it indicates that cache usage has saved considerable bandwidth.

Use the detailed diagnosis of the *Hits* measure to know which NetScaler is managing this web server, so that the load-balancing logic of that NetScaler appliance can be reconfigured to avoid overloads.

Shows the Server and NetScaler details

SERVER IP ADDRESS	NETSCALER IP ADDRESS	NETSCALER HOSTNAME
Jun 01, 2015 13:32:33		
192.168.8.122	192.168.8.20	egnetscaler

Figure 3.3: The detailed diagnosis of the Hits measure of the Web Servers test

3.2 The Web Applications Layer

Using the test mapped to this layer, requests to the web applications managed by NetScaler can be tracked, the responsiveness of each web application to these requests can be measured, processing bottlenecks identified, and the source of the bottleneck isolated.

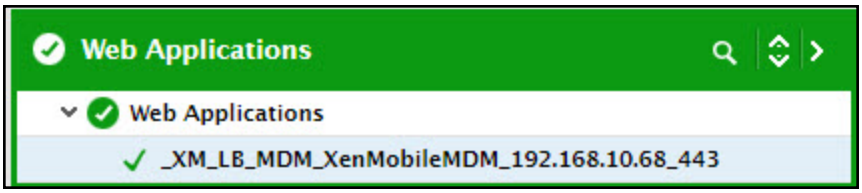

Figure 3.4: The test mapped to the Web Applications layer

3.2.1 Web Applications Test

One of the key value propositions of the NetScaler ADC is its ability to optimize user experience by making web applications run faster and by delivering them to users more quickly. If users continue to complain of slowness when accessing applications despite the usage of NetScaler ADC, it will beat the very purpose of the NetScaler ADC, forcing enterprises to hunt for alternative delivery controllers. To avoid such an eventuality, administrators should be able to detect a current/potential slowdown in a web application well before users notice, isolate the source of the slowness, and address it rapidly. This is where the **Web Applications** Test helps! This test monitors how responsive each web application delivered by NetScaler ADC is to user requests and promptly notifies administrators if any application's responsiveness dips. Upon receiving such intimations, administrators can once again zoom into the test to figure out the reason for the slowness- is it because the web server hosting the application is unable to process requests quickly? is it owing to a latent server-side network? or is it due to a latent client-side network?

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every web application managed by NetScaler ADC

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received by this web application.	Number	This is a good indicator of the load on the web application. Compare the value of this measure across web applications to know which application is the most popular. Use the detailed diagnosis of this measure to know which web server is hosting the application and which NetScaler is managing it.
Bandwidth:	Indicates the total amount of data received by this web application.	KB	Compare the value of this measure across web applications to know which web application is consuming bandwidth excessively.
Response time:	Indicates the elapsed time between the end of an enquiry and the beginning of a response from this application.	msecs	A high value for this measure indicates that the application is processing requests slowly. If this measure reports an abnormally high value for an application, compare the value of the *Server processing time, Client network latency,* and *Server*

Measurement	Description	Measurement Unit	Interpretation
			network latency measures of that application to determine the reason for the slowness.
Server processing time:	Indicates the elapsed time, from when the server starts to receive the first byte of a request from the NetScaler appliance until the NetScaler appliance receives the first byte to response.	msecs	A high value for this measure indicates that the web server is processing requests slowly. If the *Response time* measure reports an abnormally high value for a web application, then compare the value of this measure with that of the *Client network latency* and *Server network latency* measures of that application to determine what is causing the slowness – the poor processing power of the web server? a latent server network? or a slow client network?
Server network latency:	Indicates the average latency caused by the server network.	msecs	A high value for this measure indicates that the server network is latent. If the *Response time* measure reports an abnormally high value for a web application, then compare the value of this measure with that of the *Client network latency* and *Server processing time* measures of that application to determine what is causing the slowness – the poor processing power of the web server? a latent server network? or a slow client network?
Client network latency:	Indicates the average latency caused by the client network.	msecs	A high value for this measure indicates that the client network is latent. If the *Response time* measure reports an abnormally high value for a web application, then compare the value of this measure with that of the *Server network latency* and *Server processing*

Measurement	Description	Measurement Unit	Interpretation
			time measures of that application to determine what is causing the slowness – the poor processing power of the web server? a latent server network? or a slow client network?

Use the detailed diagnosis of the *Hits* measure to know which web server is hosting the web application and which NetScaler is managing it.

Shows the Application and Server details

APPLICATION NAME	SERVER IP ADDRESS	NETSCALER IP ADDRESS	NETSCALER HOSTNAME
Jun 01, 2015 13:32:06			
_XM_LB_MDM_XenMobileMDM_192.168.10.68_443	192.168.8.122	192.168.8.20	egnetscaler

Figure 3.5: The detailed diagnosis of the Hits measure of the Web Applications test

3.3 The Web Devices Layer

The tests mapped to this layer help detect issues and diagnose the root-cause of issues in the web traffic generated by:

- Each NetScaler appliance used in an environment;
- Each HTTP Request Method and Response status;
- Each client type (i.e., user agent), client, and client operating system;
- Each URL

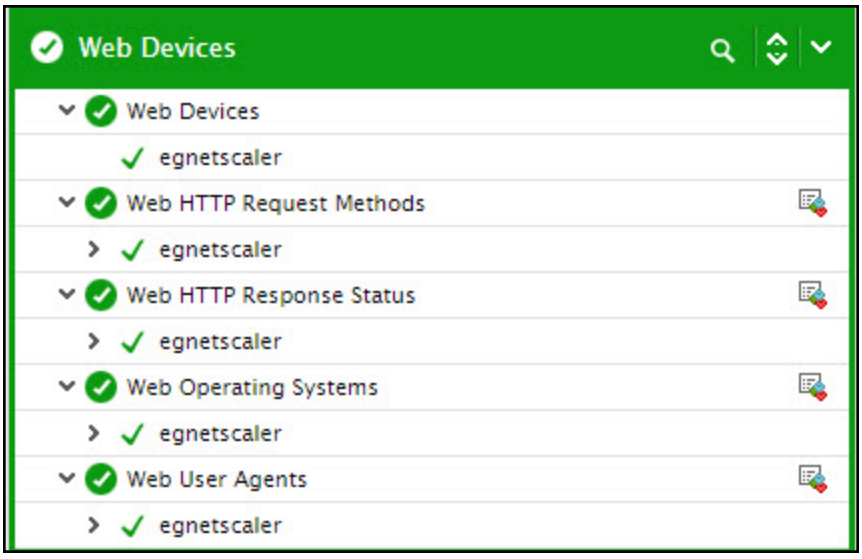

Figure 3.6: The tests mapped to the Web Devices layer

3.3.1 Web Devices Test

Where multiple NetScaler appliances have been configured, administrators can use this test to understand whether/not request load is uniformly distributed across the appliances. Irregularities in load-balancing will thus come to light. In addition, this test also captures the bandwidth used by the web traffic to and from each appliance, thus leading administrators to those appliances that are consuming bandwidth excessively.

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every NetScaler appliance configured in the target environment

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.

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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received by this appliance.	Number	This is a good indicator of the load on the appliance. Compare the value of this measure across apliances to know which NetScaler appliance is receiving the maximum number of requests. If a single NetScaler appears to be servicing a significantly large number of requests than the rest, it could imply that the NetScaler is overloaded. This in turn indicates that a faulty/ineffective load-balancing algorithm is in use. Use the detailed diagnosis of this measure to know which the IP address of the NetScaler.
Bandwidth:	Indicates the total amount of data received by this appliance.	KB	Compare the value of this measure across appliances to know which appliance is consuming bandwidth excessively.

3.3.2 Web HTTP Request Methods Test

Using the **Web Devices** test, administrators can identify which NetScaler appliance is handling bandwidth-intensive web traffic. To further investigate this anomaly, administrators may want to drill down to the individual request methods used in the HTTP requests to the appliance and understand how much bandwidth each method consumes. These method-level insights can help isolate the exact method that could be eating into the bandwidth resources available to the NetScaler. The **Web HTTP Request Methods** test provides these useful method-level insights. This test auto-discovers the HTTP request methods in use and reports the number of requests received and bandwidth used per method. In the process, the test points to the most popular and the most bandwidth-intensive HTTP request methods.

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every HTTP request method in use

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

- The eG manager license should allow the detailed diagnosis capability
- Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received by this NetScaler appliance using this HTTP request method.	Number	Compare the value of this measure across methods to know which method is the most popular. Use the detailed diagnosis of this measure to know the IP address and host name of the NetScaler ADC.
Bandwidth:	Indicates the total amount of data received by this appliance using this HTTP request method.	KB	Compare the value of this measure across methods to know which HTTP request method is consuming bandwidth excessively.

Use the detailed diagnosis of the *Hits* measure to know the IP address and host name of the NetScaler ADC.

Shows the HTTP request method and NetScaler details

METHOD NAME	NETSCALER IP ADDRESS	NETSCALER HOSTNAME
Jun 01, 2015 13:32:05		
GET	192.168.8.20	egnetscaler

Figure 3.7: The detailed diagnosis of the Hits measure of the HTTP Request Methods test

3.3.3 Web HTTP Response Status Test

Like the HTTP request methods, one/more HTTP response status messages may also consume considerable bandwidth, thus increasing the bandwidth usage of the NetScaler appliances managing the web traffic. To identify such bandwith-intensive response statuses, use this test.

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every HTTP response status

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of responses of this status sent by this NetScaler appliance.	Number	Compare the value of this measure across statuses to know which response is received often. Use the detailed diagnosis of this measure to know the IP address and

Measurement	Description	Measurement Unit	Interpretation
			host name of the NetScaler ADC.
Bandwidth:	Indicates the total size of the responses of this status sent by this appliance.	KB	Compare the value of this measure across statuses to know which type of responses is consuming bandwidth excessively.
Render time:	Indicates the elapsed time, from when the browser starts to receive the first byte of a response of this type until either all page content has been rendered or the page load action has timed out.	Msecs	Compare the value of this measure across response types to know which type of response is delaying page rendering.

Use the detailed diagnosis of the *Hits* measure to know the IP address and host name of the NetScaler ADC.

Shows the HTTP response status and NetScaler details

RESPONSE STATUS	NETSCALER IP ADDRESS	NETSCALER HOSTNAME
Jun 01, 2015 13:32:07		
Found	192.168.8.20	egnetscaler

Figure 3.8: The detailed diagnosis of the Hits measure of the HTTP Response Status test

3.3.4 Web Operating Systems Test

When measuring user experience with web applications, it would also help administrators to know which client operating systems are widely used for accessing the applications and whether bandwidth usage is impacted by the operating system in use. This way, administrators will be able to identify bandwidth-intensive operating systems and may firmly recommend that users not use clients running such incompatible operating systems. The **Web Operating Systems** test enables administrators make such recommendations. This test auto-discovers the operating systems used by clients and reports the requests received and bandwidth consumed per operating system. In the process, the test points to the most popular and the most bandwidth-intensive operating systems.

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every client operating system

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received from clients running this operating system.	Number	Compare the value of this measure across operating systems to identify the most popular operating system. Use the detailed diagnosis of this measure to know the IP address and

Measurement	Description	Measurement Unit	Interpretation
			host name of the NetScaler ADC.
Bandwidth:	Indicates the total amount of data received from clients running this operating system.	KB	Compare the value of this measure across operating systems to know which operating system has been consistently consuming more bandwidth than the rest.
Render time:	Indicates the elapsed time, from when clients running this operating system start to receive the first byte of a response until either all page content has been rendered or the page load action has timed out.	msecs	Compare the value of this measure across operating systems to know which OS is seeing the maximum page rendering time.

Use the detailed diagnosis of the *Hits* measure to know the IP address and host name of the NetScaler ADC in question.

Shows the Client OS name and NetScaler details

CLIENT OS NAME	NETSCALER IP ADDRESS	NETSCALER HOSTNAME
Jun 01, 2015 13:32:10		
Windows XP	192.168.8.20	egnetscaler

Figure 3.9: The detailed diagnosis of the Hits measure of the Web Operating Systems test

3.3.5 Web User Agents Test

Users can connect to web applications using many client devices –web browsers, mobile phones, tablets, etc. Each such client device is called a user agent. Users using certain types of client devices/user agents may be engaged in bandwidth-intensive communication over the web, scarring the experience of other users to the web applications. To capture such problem conditions quickly, it would be best to know what types of client devices users are connecting from and how much bandwidth each of these device types are currently consuming. This is exactly what the **Web User Agents** test reveals! This test auto-discovers the types of client devices that are used for connecting to web applications, and reports the requests received from and bandwidth used by each device type. This way, the test leads administrators to those device types that are popular amidst users and those that are consuming bandwidth excessively.

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every device type/user agent

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
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 enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received from this type of client.	Number	Compare the value of this measure across client types to identify that client that is used by a wide cross-section of application users.

Measurement	Description	Measurement Unit	Interpretation
			Use the detailed diagnosis of this measure to know the IP address and host name of the NetScaler ADC.
Bandwidth:	Indicates the total amount of data received from clients of this type.	KB	Compare the value of this measure across clients to know which type of client has been consistently consuming more bandwidth than the rest.
Render time:	Indicates the elapsed time, from when the browser on this device starts to receive the first byte of a response until either all page content has been rendered or the page load action has timed out.	msecs	Compare the value of this measure across client devices to know which type of device is seeing the maximum page rendering time.

Use the detailed diagnosis of the *Hits* measure to know the IP address and host name of the NetScaler ADC.

Shows the User Agent name and NetScaler details

USER AGENT NAME	NETSCALER IP ADDRESS	NETSCALER HOSTNAME
Jun 01, 2015 13:32:39		
Internet Explorer	192.168.8.20	egnetscaler

Figure 3.10: The detailed diagnosis of the Hits measure of the Web User Agents test

3.3.6 Web URLs Test

If traffic to a web application is found to be bandwidth-intensive, administrators may instantly want to know which specific URL in that application is actually hogging the bandwidth resources. Administrators may later want to block access to these URLs, so as to conserve bandwidth. Likewise, if user experience with a web application deteriorates, administrators may want to swiftly check the responsiveness of each URL requested by the user, to identify which URL is adversely impacting the user experience and why. The **Web URLs** test helps administrators with this. This test automatically discovers the URLs accessed by users and reports the number of requests received and bandwidth used per URL. This way, the test pinpoints those URLs that are accessed frequently and the ones that generate bandwidth-intensive web traffic. Additionally, the test also reports the page load time and render time of every URL, so that, when a slowdown occurs, administrators can instantly identify the URL that resulted in a slow response and where the slowdown occurred – when the requested page was loaded? Or when it was rendered by the client?

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every URL that fulfills the condition configured using the **MIN HITS FOR URL** parameter

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
5. **EXCLUDE PATTERNS** – You can exclude specific URLs/patterns of URLs from monitoring. By default, this test ignores the following URL patterns when monitoring: **.png,*.bmp,*.jpg,*.jepg,*.jfif,*.giff,*.js,*.css*. If required, you can remove one/more patterns from the default list or append more URLs/URL patterns to this list. For instance, if you want the test to ignore all URLs that end with *.asp* , then, your specification will be: **.png,*.bmp,*.jpg,*.jepg,*.jfif,*.giff,*.js,*.css,*.asp*. Likewise, if you want the test to ignore the URL, http://www.xyz.com/web/sitemap.html , your specification will be: **.png,*.bmp,*.jpg,*.jepg,*.jfif,*.giff,*.js,*.css,http://www.xyz.com/web/sitemap.html*
6. **MIN HITS FOR URL** – Using this parameter, you can exclude the less-accessed URLs from monitoring. By default, the test ignores all URLs with a hit count less than **30**. This is why, the default value of this parameter is set to **30**. You can change this value if you want more or less number of URLs monitored.
7. **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**.
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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits:	Indicates the number of requests received for this URL.	Number	Compare the value of this measure across URLs to identify the URL that is most popular. Use the detailed diagnosis of this measure to know the IP address of the NetScaler ADC that received requests for this URL, the host name of the NetScaler, and the IP address of the web server that processed requests for this URL.
Load time:	Indicates the elapsed time, from when the browser starts to receive the first byte of a response until the user starts to interact with the page. At this stage, some of the page content might not yet have been loaded.	msecs	A high value for this measure is a cause for concern as it indicates that the requested web page is taking too long to load. In the event of a slowdown, you may want to compare the value of this measure with that of the Render time measure to accurately determine the reason for the slowness – is it because of a delay in page loading? Or page rendering?
Render time:	Indicates the elapsed time, from when the browser starts to receive the first byte of a response until either all page content has been rendered or the page load action has timed out.	msecs	A high value for this measure indicates that the client is having problems rendering the requested pages. In the event of a slowdown, you may want to compare the value of this measure with that of the Load time measure to accurately determine the reason for the slowness – is it because of a delay in page loading? Or page

Measurement	Description	Measurement Unit	Interpretation
			rendering?
Bandwidth:	Indicates the total amount of data received from this URL.	KB	Compare the value of this measure across URLs to know the web traffic to which URL has been consistently consuming more bandwidth than the rest.
Cache hits:	Indicates the number of requests to this URL that was serviced by the cache.	Number	If the value of this measure is the same as that of the *Requests* measure, it implies that all requests to the URL were serviced by the cache server. This is indicative of optimal cache size and usage. On the other hand, if the value of this measure is much lower than that of the *Requests* measure, it could indicate improper cache sizing and ineffective cache usage.
Cache miss:	Indicates the number of requests to this URL not serviced by the cache.	Number	Ideally, the value of this measure should be 0 or at least, very low. If the value is the same as that of the *Requests* measure, it could indicate improper cache sizing and ineffective cache usage.
Cache hit ratio:	Indicates the percentage of requests to this URL that were serviced by the cache.	Percent	Ideally, the value of this measure should be > 80%. A low hit ratio on the other hand indicates that a majority of web requests were serviced by the origin server and not the cache server. This can significantly increase request processing time and related overheads.
Cache bypass:	Indicates the number of requests to the URL that were serviced by the origin server, because the cache server was bypassed.	Number	

Measurement	Description	Measurement Unit	Interpretation
Cache hits bandwidth consumed:	Indicates the bandwidth consumed when requests to this URL were serviced by the cache server.	KB	The difference between the value of the *Bandwidth* measure and this measure for a URL will reveal the bandwidth that may have been saved by request caching. Where cache is well-sized and used optimally, this difference will be high.
Cache misses bandwidth consumed:	Indicates the bandwidth consumed when the cache server could not serve the requests to this URL.	KB	The difference between the value of this measure and the value of the *Cache hits bandwidth consumed* measure for a web server will reveal how much bandwidth was saved by cache hits.
Cache bypass bandwidth consumed:	Indicates the bandwidth consumed when the cache server was bypassed and the requests to this URL were served from the origin server.	KB	If the difference between the value of this measure and that of the *Cache hits bandwidth consumed* measure results in a 'positive' integer, it indicates that cache usage has saved considerable bandwidth.

3.3.7 Web Clients Test

When encountered by a request overload on a web server, administrators must quickly identify the client that could have contributed to that load. Also, when users connecting from a specific client complain of slowness, administrators should be able to swiftly zero-in on the root-cause of that slowness, so that the problem can be resolved before user productivity is impacted. The **Web Clients** test helps in both accounts! This test tracks requests from every client connecting to the web servers managed by a NetScaler appliance, and reports the requests count, bandwidth usage, page rendering time, and the network latency for each client. From these metrics, administrators can easily infer which client is imposing the maximum load on the web servers. Moreover, if help desk receives frequent complaints of slowness from users connecting from a particular client, then administrators can use the metrics reported by this test to isolate the source of the delay – is it because the client is unable to render the response pages quickly? is it because of a latent client network? or is it owing to a contention for bandwidth resources?

Target of the test : Citrix NetScaler Web Insight

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every client that fulfills the condition configured using the **SHOW BY** and **SHOW ABOVE LIMIT** parameters

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed. **It is recommended that you set the TEST PERIOD to 5 minutes. This is because, the Nitro API using which the eG agent collects metrics from Web Insight, is capable of capturing only the performance data related to the last 5 minutes**.
2. **HOST** - The host for which the test is to be configured.
3. **INSIGHT USERNAME**, **INSIGHT PASSWORD**, and **CONFIRM PASSWORD** - To connect to Web Insight and collect the metrics it captures, the eG agent needs to be configured with the credentials of a user with **read-only** permissions to Web Insight. Type the name of this user against **INSIGHT USERNAME** and the password of this user against **INSIGHT PASSWORD**. Then, confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
4. **SSL** – By default, Web Insight is not SSL-enabled. This is why, this flag is set to **No** by default. If it is SSL-enabled, then change this flag to **Yes**.
5. **SHOW BY** and **SHOW ABOVE LIMIT** – Not all clients interacting with a web server may be worth monitoring. Some clients may be seldom used and hence may not require monitoring; some other clients may rarely encounter any slowdowns, and hence can be ignored when monitoring. Administrators may want to exclude such clients from monitoring, so that the strain on the eG agent is reduced. From the **SHOW BY** drop-down, you can pick the criterion based on which a client's 'monitoring worth' is to be evaluated. The options here are: **Render time**, **Requests**, and *All*. If the *Render time* option is chosen from the **SHOW BY** list, then the test will monitor only those clients for which the *Render time* measure of this test registers a value equal to or above the value specified against **SHOW ABOVE LIMIT**. If the *Requests* option is chosen from the **SHOW BY** list, then the test will monitor only those clients for which the *Requests* measure of this test registers a value equal to or above the value specified against **SHOW ABOVE LIMIT**. By default, the **SHOW ABOVE LIMIT** is set to **30**; this implies that:
   - If **SHOW BY** is set to *Render time*, then this test will monitor only those clients that report a value of 30 msecs or above for the *Render time* measure;
   - If **SHOW BY** is set to *Requests*, then this test will monitor only those clients that send at least 30 requests to the web servers;

   On the other hand, if you want to monitor all clients that communicate with the web server, regardless of the render time or request count, then select the *All* option from the **SHOW BY** list. In this case, the *show above limit* setting will be disregarded.
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.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Requests:	Indicates the number of requests received from this client.	Number	In the event of an overload condition, you can compare the value of this measure across clients to know which client is overloading the servers. Use the detailed diagnosis of this measure to know which servers this client is accessing and which NetScaler manages the web traffic generated by this client.
Render time:	Indicates the elapsed time, from when the browser starts to receive the first byte of a response until either all page content has been rendered by this client or the page load action has timed out.	msecs	A high value for this measure is a cause for concern as it indicates that the client is delaying page rendering. In the event of a slowdown, you can compare the value of this measure with that of the *Client network latency* measure to zero-in on the root-cause of the slowdown - is it because the client is unable to render the response pages quickly? or is it because of a latent client network?
Client network latency:	Indicates the latency caused by the client-side network.	msecs	A high value for this measure is a cause for concern as it indicates that a bottleneck in the client-side network. In the event of a slowdown, you can compare the value of this measure with

Measurement	Description	Measurement Unit	Interpretation
			that of the *Render time* measure to zero-in on the root-cause of the slowdown - is it because the client is unable to render the response pages quickly? or is it because of a latent client network?
Bandwidth:	Indicates the total amount of data received from this client.	KB	Compare the value of this measure across clients to know which client is hogging the bandwidth. In the event of a slowdown, you can use the value of this measure to figure out if the lack of adequate bandwidth is what is slowing down user accesses from this client.
Cache hits:	Indicates the number of requests from this client that were serviced by the cache.	Number	If the value of this measure is the same as that of the *Hits* measure, it implies that all requests from the client were serviced by the cache. This is indicative of optimal cache size and usage. On the other hand, if the value of this measure is much lower than that of the *Hits* measure, it could indicate improper cache sizing and ineffective cache usage.
Cache miss:	Indicates the number of requests from this client that were not serviced by the cache.	Number	Ideally, the value of this measure should be 0 or at least, very low. If the value is the same as that of the *Hits* measure, it could indicate improper cache sizing and ineffective cache usage.
Cache hit ratio:	Indicates the percentage of requests from this client that were serviced by the cache.	Percent	Ideally, the value of this measure should be > 80%. A low hit ratio on the other hand indicates that a majority of web requests were serviced by the origin server and not the cache server. This can significantly increase request processing time and related overheads.

Measurement	Description	Measurement Unit	Interpretation
Cache bypass:	Indicates the number of requests from this client that were serviced by the origin server, because the cache server was bypassed.	Number	
Cache hits bandwidth consumed:	Indicates the bandwidth consumed when requests from this client were serviced by the cache server.	KB	The difference between the value of the *Bandwidth* measure and this measure for a client will reveal the bandwidth that may have been saved by request caching. Where cache is well-sized and used optimally, this difference will be high.
Cache misses bandwidth consumed:	Indicates the bandwidth consumed when the cache server could not serve the requests from this client.	KB	The difference between the value of this measure and the value of the *Cache hits bandwidth consumed* measure for a client will reveal how much bandwidth was saved by cache hits.
Cache bypass bandwidth consumed:	Indicates the bandwidth consumed when the cache server was bypassed and the requests from this client were served from the origin server.	KB	If the difference between the value of this measure and that of the *Cache hits bandwidth consumed* measure results in a 'positive' integer, it indicates that cache usage has saved considerable bandwidth.

4

Conclusion

This document has clearly explained how eG Enterprise integrates with **Citrix NetScaler HDX Insight** and **Citrix NetScaler Web Insight**. For more information on eG Enterprise, please visit our web site at www.eginnovations.com or write to us at sales@eginnovations.com.