Monitoring Citrix Netscaler VPX/MPX

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

Citrix NetScaler application delivery solutions combine the features and functions of traditional data center point products - load balancing, caching, compression, SSL acceleration, and attack defense - into a single network appliance, built from the ground up to maximize the performance of mission-critical applications.

The Citrix NetScaler LB products are built on Citrix's patented Request Switching™ architecture, the industry's only wire-speed technology that handles every application request based on powerful user-defined policies. The Citrix NetScaler application-aware policy engine, AppExpert™, allows the creation of detailed policy-based decisions for individual requests, irrespective of connections. AppExpert lets administrators build sophisticated application request handling policies that enable powerful, comprehensive application-based features.

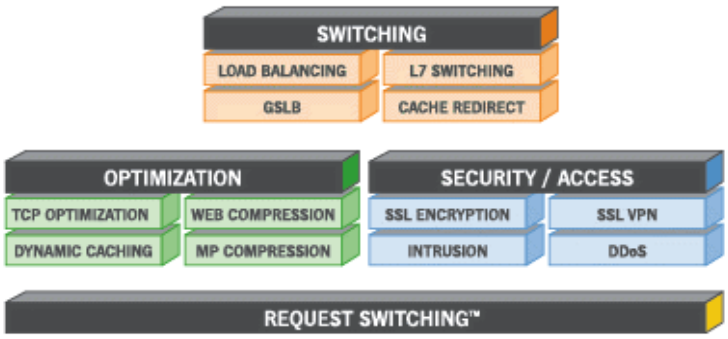

Figure 1.1: The NetScaler architecture

As business entities have begun to rely enormously on the Citrix NetScaler solutions to deliver service continuity and to ensure the secure transaction of business, the smooth functioning of the NetScaler appliance has become super-critical in Citrix infrastructures today. Round-the-clock monitoring of NetScaler products, proactive error reporting, and swift error-clearance are a must to ensure that the Citrix NetScaler is always up and running, and is well enough to attend to application requests from users. This is where the eG Enterprise lends helping hands to administrators.

Chapter 2: How to Monitor Citrix NetScaler VPX/MPX Using eG Enterprise?

eG Enterprise can monitor the Citrix NetScaler VPX/MPX appliance in an agentless manner only. All that is required for this is an eG agent deployed on any remote Windows host. This agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler appliance. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode.

To start NetScaler monitoring, manage the Netscaler component using the eG administrative interface and configure the tests to enable the eG agent to use the NITRO commands to fetch measures. The procedure for achieving this has been discussed in the following section.

2.1 Managing the Citrix NetScaler VPX/MPX Appliance

To manage the eG Manager to monitor the Citrix NetScaler VPX/MPX, do the following:

1. Log into the eG administrative interface.
2. eG Enterprise cannot automatically discover the Citrix NetScaler VPX/MPX server. You need to manually add the server using the **COMPONENTS** page (see Figure 2.1) that appears when the Infrastructure -> Components -> Add/Modify menu sequence is followed. Remember that components manually added are managed automatically.

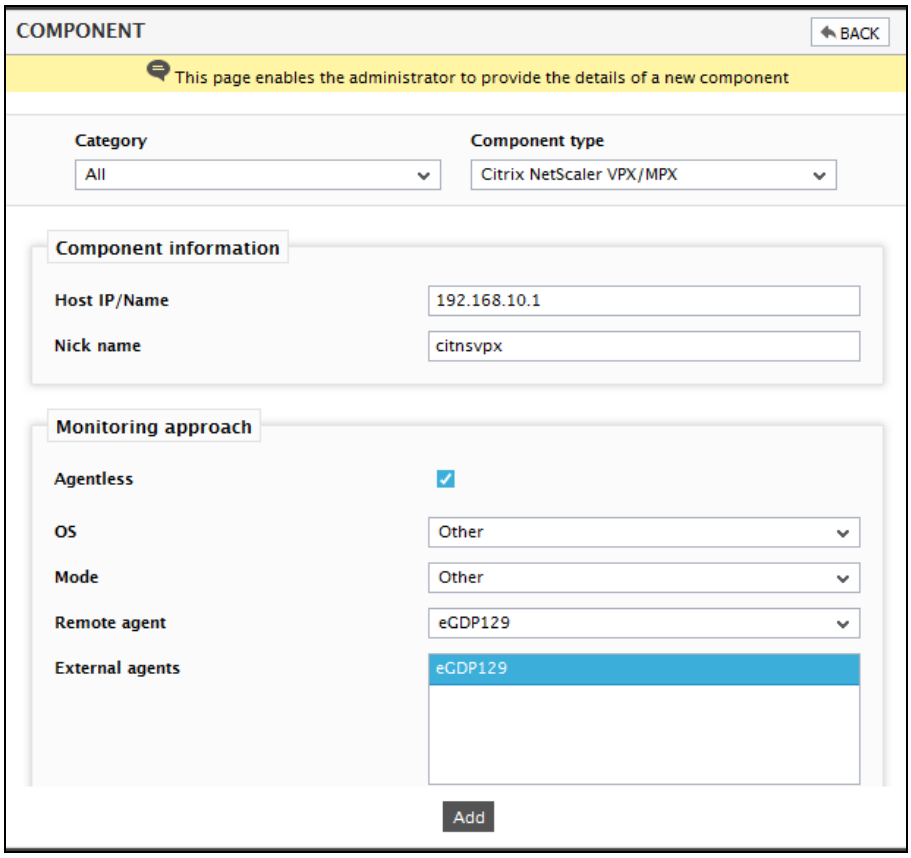

Figure 2.1: Adding the Citrix NetScaler VPX/MPX

3. Specify the **Host IP/Name** and the **Nick name** for the Citrix NetScaler VPX/MPX appliance.
4. Since the appliance can be monitored only in an agentless manner, the **Agentless** flag will be switched on by default (see Figure 2.1).
5. Select *Other* from the **OS** and **Mode** drop-down lists.
6. Pick the **Remote agent** that will monitor the appliance. Similarly, select the **External agent** that will report the network availability and responsiveness of the appliance.
7. Finally, click the **Add** button to add the appaliance for monitoring.
8. Next, try to to sign out, a list of unconfigured tests appears as shown in Figure 2.2.

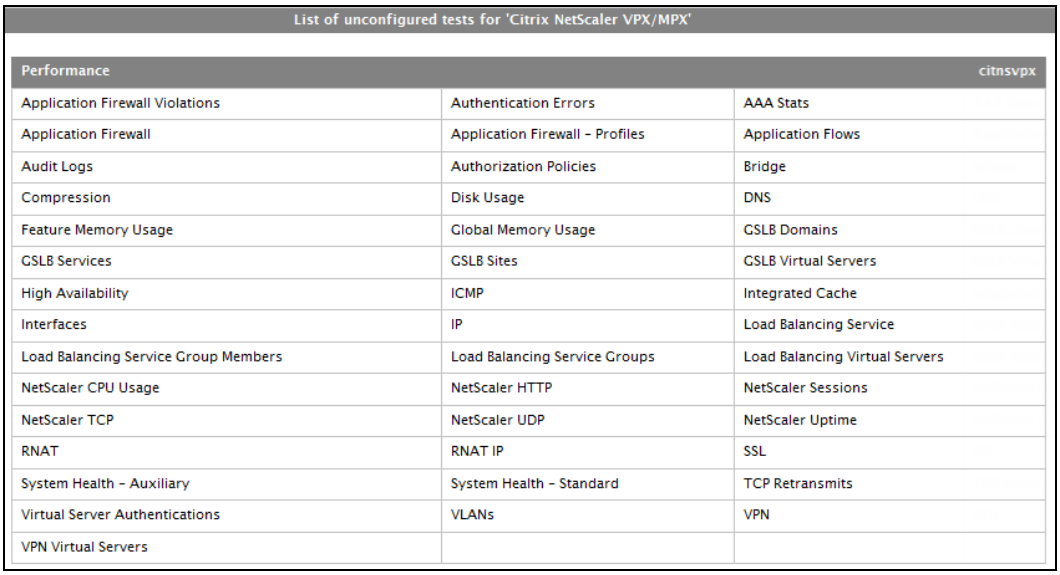

List of unconfigured tests for 'Citrix NetScaler VPX/MPX'

Performance		citnsvpx
Application Firewall Violations	Authentication Errors	AAA Stats
Application Firewall	Application Firewall - Profiles	Application Flows
Audit Logs	Authorization Policies	Bridge
Compression	Disk Usage	DNS
Feature Memory Usage	Global Memory Usage	GSLB Domains
GSLB Services	GSLB Sites	GSLB Virtual Servers
High Availability	ICMP	Integrated Cache
Interfaces	IP	Load Balancing Service
Load Balancing Service Group Members	Load Balancing Service Groups	Load Balancing Virtual Servers
NetScaler CPU Usage	NetScaler HTTP	NetScaler Sessions
NetScaler TCP	NetScaler UDP	NetScaler Uptime
RNAT	RNAT IP	SSL
System Health - Auxiliary	System Health - Standard	TCP Retransmits
Virtual Server Authentications	VLANs	VPN
VPN Virtual Servers		

Figure 2.2: List of unconfigured tests to be configured for the Citrix NetScaler VPX

9. Click on the **Application Firewall Violations** test in Figure 2.2 to configure it. To know how to configure the test, refer to Section **3.6.4**.
10. Once all the tests are configured, signout of the eG administrative interface.

Chapter 3: Monitoring Citrix NetScaler VPX/MPX

Citrix NetScaler VPX/MPX is an all-in-one service and application delivery solution that accelerates application performance, increases application availability and improves application security.

For your mission-critical applications to operate at peak capacity, you need to ensure that the NetScaler VPX/MPX in your environment functions without a glitch! To enable administrators to ensure the continuous availability and problem-free execution of the NetScaler solution, the eG Enterprise Suite offers a dedicated *Citrix NetScaler VPX/MPX* model.

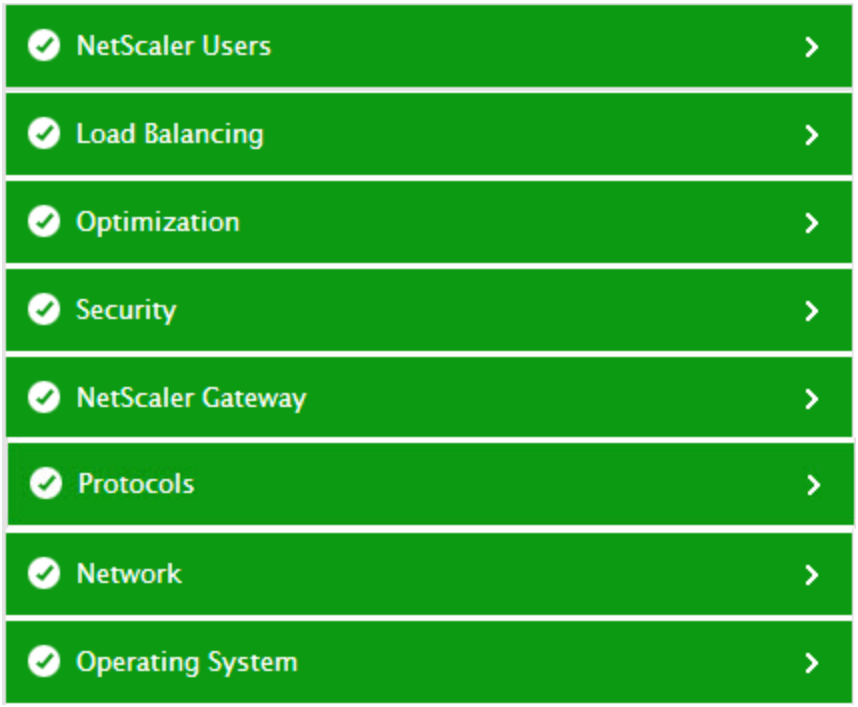

Figure 3.1: The Citrix NetScaler VPX/MPX monitoring model

Each layer of Figure 3.1 is mapped to tests that initiate *Nitro API* calls on the NetScaler appliance, in a non-intrusive, 'agentless' manner, and pull out a wealth of performance information related to the appliance. These metrics reveal the following:

- Is the NetScaler device available over the network?
- Was the NetScaler rebooted recently?
- Is NetScaler consuming CPU resources excessively?
- Is the NetScaler running out of disk space?

- Does the NetScaler have enough free global memory or has memory been over-allocated to the features? Which feature is utilizing the allocated memory unwisely?
- Is the NetScaler appliance been subjected to high levels of voltage/temperature?
- Are the fans supported by the appliance operating at an optimal speed? If not, which fan is operating at an abnormal speed?
- Which CPU core of the NetScaler is currently experiencing abnormally high or low voltage?
- Which CPU core's fan is operating at an abnormal speed?
- Are the system fans operating at optimal speeds?
- The temperature of which CPU core is very high right now?
- Which voltage rail - +3.3V, 5V, or 12V - inside a power supply is currently conducting a very high or low voltage of current than permitted? Is this abnormality inside the main power supply itself itself or only in the standby unit?
- Is any power supply currently in an abnormal state?
- Did any bridge table collisions occur recently?
- Were too many loops detected on the bridge?
- Were any interfaces muted?
- How much network traffic has been generated by the RNAT sessions that are active on the NetScaler? Which client IP address is responsible for generating the maximum load?
- Were too many packets dropped by any VLAN? If so, which VLAN is it?
- Were any large/invalid HTTP requests received and similar responses sent by the NetScaler?
- Was the ICMP rate threshold violated?
- Are too many TCP requests queued in the Surge Queue?
- Have too many TCP retransmissions occurred recently?
- Have any IP address lookups failed on the NetScaler?
- Is any virtual server down? Which service/service group will be impacted by this?
- Is any service overloaded?
- Is the HA node in a High Availability setup currently up and running?
- Is the monitored node the primary of secondary node of an HA setup?
- Did command propagation time out too often between the primary and secondary nodes?

- Did the primary and secondary nodes in a High Availability setup fail to synchronize?
- Is any load balancing virtual server in an unhealthy state currently?
- Was the spill over threshold of any virtual server violated?
- Were any audit log messages not sent to the SYSLOG server? What is causing problems in tranmission - is it because of NAT/NSB allocation failures? is it because memory allocations of the Access Gateway context structure failed? is it due to too many port allocation failures?
- Did any user authentications fail recently?
- Was any Appflow data not transmitted or ignored during transmission to IPv4 data collectors configured on the appliance?
- How well is the NetScaler Application Firewall functioning? Were any requests to the firewall aborted before completion? How many and what type of security check violations were captured by the firewall?
- Is any authentication virtual server down?
- Is any authentication virtual server experiencing processing delays?
- Is the SSL engine down?
- Is the number of Front-End SSL session reuse misses high?
- Did too many Back-End SSL session multiplexing attempts fail?
- Were any UDP packets received on an unknown NetScaler port?
- Were any UDP packets received with a UDP checksum error?
- Was the UDP rate threshold violated?
- Is any NetScaler service overloaded? If so, which one is it?
- Has the NetScaler retransmitted too many packets?
- Did the NetScaler respond to all DNS queries that were received? Were any queries found to be invalid? Were any DNS requests refused? Were any invalid responses sent?
- Did the VPN login page fail to appear too often?
- Were too many STA connection failures, SOCKS client errors, and ICA license failures detected on the NetScaler?

- How is the hit ratio of the Integrated Cache? Optimal or poor? If poor, what could be the reason for the same? Is it because the cache does not have sufficient memory to hold many objects? Should the cache memory be resized?
- How efficient is the current compression algorithm? Does the current compression ratio result in significant bandwidth savings?
- Does compression occur too frequently? Should the quantumSize be reconfigured to reduce the frequency of compression?

The sections that follow will discuss each layer of Figure 3.1 elaborately.

3.1 The Operating System Layer

The tests mapped to this layer report how well the appliance uses the CPU, disk space, and global memory resources available to it, points you to the memory-hungry features on the appliance, and alerts you to potential hardware failures on the appliance.

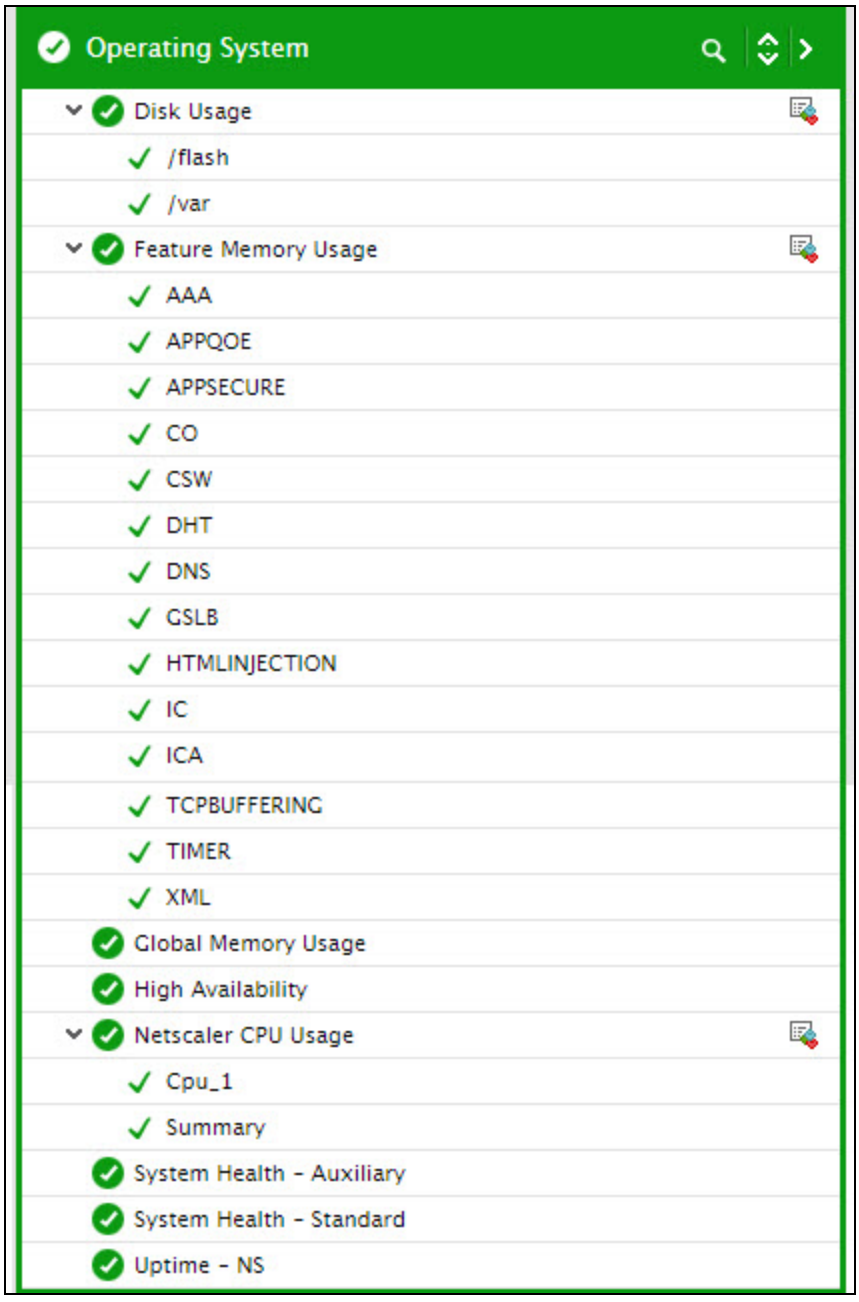

Figure 3.2: The tests mapped to the Operating System layer

3.1.1 System Health – Standard Test

Using this test, you can determine the following:

- Which CPU core of the NetScaler is currently experiencing abnormally high or low voltage?
- Which CPU core's fan is operating at an abnormal speed?

- Are the system fans operating at optimal speeds?
- The temperature of which CPU core is very high right now?
- Which voltage rail - +3.3V, 5V, or 12V - inside a power supply is currently conducting a very high or low voltage of current than permitted? Is this abnormality inside the main power supply itself or only in the standby unit?
- Is any power supply currently in an abnormal state?

Target of the test : A NetScaler VPX/MPX device

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler device being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
CPU 0 core voltage	Indicates the current voltage of CPU core 0 of	Volts	The CPU core voltage must be within the range of 1.08V to 1.65V.

Measurement	Description	Measurement Unit	Interpretation
	the NetScaler device.		
CPU 1 core voltage	Indicates the current voltage of CPU core 1 of the NetScaler device.	Volts	
CPU Fan 0 speed	Indicates the current speed of the CPU fan 0.	Rpm	The speed of the fan must be within the range of 4000 to 6000 rpm.
CPU Fan 1 speed	Indicates the current speed of the CPU fan 1.	Rpm	
CPU 0 temperature	Indicates the current temperature of CPU core 0.	Celsius	An abnormal temperature may cause severe damage to the CPU.
CPU 1 temperature	Indicates the current temperature of CPU core 1.	Celsius	
Intel CPU Vtt Power	Indicates the termination voltage (Vtt) used to interface the MCH with the CPU die(s).	Volts	Very high VTT voltage can cause irreparable damage to the CPU.
Main 3.3v power supply	Indicates the current voltage output of the 3.3v voltage rail inside the main power supply unit.	Volts	Ideally, the value of this measure should be well within normal limits.
Standby 3.3v power supply	Indicates the current voltage output of the 3.3v voltage rail inside the standby power supply unit.	Volts	Ideally, the value of this measure should be well within normal limits.
+5v power supply	Indicates the current voltage output of the positive 5v power supply of this NetScaler device.	Volts	
-5v power supply	Indicates the current voltage output of the negative 5v power supply of the NetScaler device.	Volts	

Measurement	Description	Measurement Unit	Interpretation
Standby 5v power supply	Indicates the current voltage output of 5v voltage rail inside the standby power supply unit.	Volts	
+12v power supply	Indicates the current voltage output of the positive 12v power supply of the NetScaler device.	Volts	
-12v power supply	Indicates the current voltage output of the negative 12v power supply of this NetScaler device.	Volts	
Battery power supply	Indicates the onboard battery power of this NetScaler device.	Volts	
Internal temperature	Indicates the internal temperature of the health monitoring chip of this NetScaler device.	Celsius	Low temperature levels are ideal for uninterrupted functioning of the device.
System Fan speed	Indicates the current speed of the system fan.	Rpm	
System Fan 1 speed	Indicates the current speed of the system fan 1.	Rpm	
System Fan 2 speed	Indicates the current speed of the system fan 2.	Rpm	
Power supply 1 status	Indicates the current state (whether normal/not) of power supply 1.		The values reported by this measure and their corresponding numeric equivalents are described in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>Abnormal</td></tr>
<tr><td>-4</td><td>Not Supported</td></tr>
</table>
If the power supply 1 is not present in the target NetScaler device, then, this measure will not appear in the eG monitoring console.
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether the power supply 1 is normal or not. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.
</td></tr>
<tr><td>Power supply 2 status</td><td>Indicates the current state (whether normal/not) of power supply 2.</td><td></td><td>
The values reported by this measure and their corresponding numeric equivalents are described in the table below:
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Normal</td></tr>
<tr><td>1</td><td>Abnormal</td></tr>
<tr><td>-4</td><td>Not Supported</td></tr>
</table>
If the power supply 2 is not present in the target NetScaler device, then, this measure will not appear in the eG monitoring console.
Note:
By default, this measure reports the above-mentioned Measure Values
</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>while indicating whether the power supply 2 is normal or not. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Power supply 3 status</td><td>Indicates the current state (whether normal/not) of power supply 3.</td><td></td><td>The values reported by this measure and their corresponding numeric equivalents are described in the table below:
<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>0</td><td>Normal</td></tr><tr><td>1</td><td>Abnormal</td></tr><tr><td>-4</td><td>Not Supported</td></tr></table>
If the power supply 3 is not present in the target NetScaler device, then, this measure will not appear in the eG monitoring console.
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether the power supply 3 is normal or not. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Power supply 4 status</td><td>Indicates the current state (whether normal/not) of power supply 4.</td><td></td><td>The values reported by this measure and their corresponding numeric equivalents are described in the table below:</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>0</td><td>Normal</td></tr><tr><td>1</td><td>Abnormal</td></tr><tr><td>-4</td><td>Not Supported</td></tr></table> If the power supply 4 is not present in the target NetScaler device, then, this measure will not appear in the eG monitoring console. **Note:** By default, this measure reports the above-mentioned **Measure Values** while indicating whether the power supply 4 is normal or not. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.

3.1.2 NetScaler CPU Usage Test

This test reports how well the CPU resources are utilized by the NetScaler device and also reports the number of CPU cores available in the NetScaler device.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each processor supported by the NetScaler device.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username	To monitor a NetScaler device, the eG agent should be configured with the credentials

Parameter	Description
and NetScaler Password	of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
CPU usage	Indicates the percentage of CPU resources used by this processor.	Percent	**For the 'Summary' descriptor, this value indicates the overall CPU usage of the NetScaler device as a whole.** A value close to 100% indicates a CPU bottleneck on the NetScaler device.
Management CPU usage	Indicates the current management CPU usage of the NetScaler device.	Percent	**This measure is available only for the 'Summary' descriptor.** Ideally, the value of this measure should be low.
Packet CPU usage	Indicates the current packet CPU usage of the NetScaler device.	Percent	**This measure is available only for the 'Summary' descriptor.** Ideally, the value of this measure should be low.
CPU cores	Indicates the number of CPU cores currently supported by the NetScaler device.	Number	**This measure is available only for the 'Summary' descriptor.**

3.1.3 Disk Usage Test

This test monitors the space usage of every disk partition on the NetScaler device.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each disk partition supported by the NetScaler device.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total capacity	Indicates the total capacity of this disk partition.	MB	
Used space	Indicates the amount of space in this disk partition that is currently in use.	MB	Ideally, this value should be low.
Free space	Indicates the amount of	MB	Ideally, this value should be high.

Measurement	Description	Measurement Unit	Interpretation
	space in this disk partition that is currently unused/free.		
Percent usage	Indicates the percentage of space usage on this disk partition.	Percent	A value close to 100% may indicate a potential problem situation where applications executing on this NetScaler device may not be able to write data to the disk partition(s) with very high usage.

3.1.4 Global Memory Usage Test

This test reports how well the NetScaler device uses the memory available to it, tracks shared memory usage by the device, and promptly alerts administrators to abnormal memory usage by the device.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why,

Parameter	Description
	the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total memory	Indicates the total memory capacity of the NetScaler device.	MB	
Total memory available	Indicates the total memory available for use in the packet engine (PE) of this NetScaler device.	MB	
Memory allocated	Indicates the total amount of memory that is currently allocated to the NetScaler features.	MB	
Percent of memory allocated	Indicates the percentage of memory that is allocated to the NetScaler features.	MB	
Memory in use	Indicates the amount of memory that is currently utilized by the NetScaler device.	MB	A low value is desired for this measure. A consistent increase in this value could be indicative of a potential memory contention.
Percent of memory in use	Indicates the percentage of memory that is currently in use in this NetScaler device.	Percent	
Free memory	Indicates the amount of memory that is currently unused in the NetScaler device.	MB	A high value is desired for this measure.
Total shared memory	Indicates the amount of memory that is allocated for sharing by this	MB	

Measurement	Description	Measurement Unit	Interpretation
	NetScaler device.		
Shared memory in use	Indicates the amount of shared memory that is currently in use in this NetScaler device.	MB	Ideally, this value should be low.
Percent of shared memory in use	Indicates the percentage of shared memory that is currently in use in this NetScaler device.	Percent	Ideally, this value should be low.

3.1.5 Feature Memory Usage Test

Fractions of the global memory are allocated to each feature that is operational on the NetScaler appliance. Using this test, you can track the memory usage of each feature and accurately identify those features that are memory-intensive. Memory allocations to features can be fine-tuned based on the usage results reported by this test.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each feature configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device.

Parameter	Description
	Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
NetScaler memory available for use	Indicates the amount of NetScaler memory currently allocated for this feature.	MB	
NetScaler memory used	Indicates the percentage of allocated memory currently used by this feature.	Percent	A value close to 100% is a cause for concern as it indicates excessive memory consumption by a feature. You can compare the value of this measure across features to identify which feature memory-hungry. Memory allocation to such features can be increased if found necessary.
Memory allocation failure	Indicates the amount of memory that failed during allocation to this feature.	MB	Allocation failures can occur owing to a memory crunch.

3.1.6 System Health – Auxiliary Test

This test captures the temperature, voltage, and fan speed of the hardware of the NetScaler device when it connects to the health monitoring chip through pins. Besides alerting administrators to abnormal temperature, voltage, and fan speed readings, this test also helps you figure out which pin registered these abnormalities.

Target of the test : A NetScaler VPX/MPX device

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler device being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Voltage 0	Indicates the voltage of the device connected to the health monitoring chip through pin 0.	Volts	Ideally, the value of this measure should be well within the range of the device that is connected to the health monitoring chip. A sudden/consistent increase/decrease in this value could warrant an investigation. Observe the variations in the value of this measure over time and compare the values across pins to isolate the pin through which the device was connecting when it most often registered the maximum voltage.
Voltage 1	Indicates the voltage of the device connected to the health monitoring chip through pin 1.	Volts	
Voltage 2	Indicates the voltage of the device connected to the health monitoring chip through pin 2.	Volts	
Voltage 3	Indicates the voltage of the device connected to the health monitoring chip through pin 3.	Volts	

Measurement	Description	Measurement Unit	Interpretation
Voltage 4	Indicates the voltage of the device connected to the health monitoring chip through pin 4.	Volts	Ideally, the value of this measure should be well within the range of the device that is connected to the health monitoring chip. A sudden/consistent increase/decrease in this value could warrant an investigation. Observe the variations in the value of this measure over time and compare the values across pins to isolate the pin through which the device was connecting when it most often registered the maximum voltage.
Voltage 5	Indicates the voltage of the device connected to the health monitoring chip through pin 5.	MB	
Voltage 6	Indicates the voltage of the device connected to the health monitoring chip through pin 6.	MB	
Voltage 7	Indicates the voltage of the device connected to the health monitoring chip through pin 7.	MB	
Fan 0 speed	Indicates the speed of fan 0 if associated pin is connected to the health monitoring chip.	Rpm	Ideally, the speed of the fans must be within normal limits. High fan speed can very often be attributed to high temperatures. If a fan speed is found to be abnormal, then, you may want to check the Temperature measures reported by the test to figure out whether any device connecting to the health monitoring chip has recorded a very high temperature at around the same time. This way, you can also isolate the pin through which that device was connecting at the time of the problem.
Fan 1 speed	Indicates the speed of fan 1 if associated pin is connected to the health monitoring chip.	Rpm	
Fan 2 speed	Indicates the speed of fan 2 if associated pin is connected to the health monitoring chip.	Rpm	
Fan 3 speed	Indicates the speed of fan 3 if associated pin is connected to the health monitoring chip.	Rpm	
Temperature 0	Indicates the temperature of the device that is connected to the health monitoring chip through	Celsius	An abnormal temperature may cause damage to the device. So the temperature of the device must be well within normal limits.

Measurement	Description	Measurement Unit	Interpretation
	pin 0.		
Temperature 1	Indicates the temperature of the device that is connected to the health monitoring chip through pin 1.	Celsius	
Temperature 2	Indicates the temperature of the device that is connected to the health monitoring chip through pin 2.	Celsius	
Temperature 3	Indicates the temperature of the device that is connected to the health monitoring chip through pin 3.	Celsius	

3.1.7 NetScaler Uptime Test

In most production environments, it is essential to monitor the uptime of critical servers/devices in the infrastructure. By tracking the uptime of each of the servers/device, administrators can determine what percentage of time a server/device has been up. Comparing this value with service level targets, administrators can determine the most trouble-prone areas of the infrastructure.

In some environments, administrators may schedule periodic reboots of their servers/devices. By knowing that a specific server/device has been up for an unusually long time, an administrator may come to know that the scheduled reboot task is not working on a server/device.

The **NetScaler Uptime** test included in the eG agent monitors the uptime of the NetScaler device.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
ReportManagerTime	By default, this flag is set to **Yes**, indicating that, by default, the detailed diagnosis of this test, if enabled, will report the shutdown and reboot times of the device in the manager's time zone. If this flag is set to **No**, then the shutdown and reboot times are shown in the time zone of the system where the agent is running(i.e., the system being managed for agent-based monitoring, and the system on which the remote agent is running - for agentless monitoring).
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

Parameter	Description
	• 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
Has the NetScaler been rebooted?	Indicates whether the server has been rebooted during the last measurement period or not.	Boolean	If this measure shows 1, it means that the server was rebooted during the last measurement period. By checking the time periods when this metric changes from 0 to 1, an administrator can determine the times when this server was rebooted.
Uptime of the NetScaler during the last measure period	Indicates the time period that the system has been up since the last time this test ran.	Secs	If the server has not been rebooted during the last measurement period and the agent has been running continuously, this value will be equal to the measurement period. If the server was rebooted during the last measurement period, this value will be less than the measurement period of the test. For example, if the measurement period is 300 secs, and if the server was rebooted 120 secs back, this metric will report a value of 120 seconds. The accuracy of this metric is dependent on the measurement period – the smaller the measurement period, greater the accuracy.
Total uptime of the NetScaler	Indicates the total time that the server has been up since its last reboot.		This measure displays the number of years, months, days, hours, minutes and seconds since the last reboot. Administrators may wish to be alerted if a server has been running without a reboot for a very long period. Setting a

Measurement	Description	Measurement Unit	Interpretation
			threshold for this metric allows administrators to determine such conditions.

3.1.8 High Availability Test

A high availability (HA) deployment of two Citrix NetScaler appliances can provide uninterrupted operation in any transaction. With one appliance configured as the primary node and the other as the secondary node, the primary node accepts connections and manages servers while the secondary node monitors the primary. If, for any reason, the primary node is unable to accept connections, the secondary node takes over.

The secondary node monitors the primary by sending periodic messages (often called heartbeat messages or health checks) to determine whether the primary node is accepting connections. If a health check fails, the secondary node retries the connection for a specified period, after which it determines that the primary node is not functioning normally. The secondary node then takes over for the primary (a process called failover).

When monitoring a NetScaler appliance, you may first want to know whether the target NetScaler has been deployed in an HA setup or not. If so, you may then want to time and again evaluate the effectiveness of the HA configuration, by checking the following:

- Is the target NetScaler appliance the primary node or the secondary node of the HA setup?
- What is the current state of the monitored NetScaler appliance? Did failover occur when the primary node went down?
- Are both nodes able to communicate freely with each other via heartbeat packets?
- Is the configuration of both the primary and secondary managers in sync?

The **High Availability** test does all the above, and more!

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is HA enabled?	Indicates whether/not the target NetScaler appliance is enabled for High Availability.		This measure reports the value Yes if 'High availability' is enabled for the monitored node, and returns the value No if 'High availability' is not configured for the node - i.e., if the node is not part of a 'High Availability' setup. The values reported by this measure and their numeric equivalents are as shown in the table: **Numeric Value** / **Measure Value**: 0 – Yes; 1 – No **Note:**

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>By default, this measure reports the above-mentioned Measure Values while indicating the high availability state of the current node. However, in the graph of this, the HA states will be represented using the corresponding numeric equivalents i.e., 0 or 1.</td></tr>
<tr><td>HA node state</td><td>Indicates the current state of the highly available node.</td><td></td><td>The values that this measure can report and their numeric equivalents are as shown in the table below:
<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>0</td><td>Up</td></tr><tr><td>1</td><td>Init</td></tr><tr><td>2</td><td>Disabled</td></tr></table>Note:
By default, this measure reports the above-mentioned Measure Values while indicating the current state of the highly available node. However, in the graph of this measure, the state will be represented using the corresponding numeric equivalents only - i.e., 0 to 2.</td></tr>
<tr><td>HA role of this node</td><td>Indicates the current role of this NetScaler device in a high availability setup.</td><td></td><td>The values that this measure can report and their corresponding numeric equivalents are shown in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Primary</td><td>0</td></tr><tr><td>Secondary</td><td>1</td></tr><tr><td>Stay Sec-ondary</td><td>2</td></tr></table></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			Note: By default, this measure reports the above-mentioned **Measure Value**s while indicating whether/not this NetScaler device is the master. However, in the graph of this measure, the states will be represented using the corresponding numeric equivalents only - i.e., 0 to 2.
Heartbeats received	Indicates the number of heartbeat packets received from the peer node - i.e., from the secondary NetScaler devices in a high availability setup - during the last measurement period.	Number	The heartbeat messages are UDP packets sent to port 3003 of the other node in an HA pair. These heartbeats are sent at regular intervals (default is 200 milliseconds) to determine the state of the peer node.
Heartbeats sent	Indicates the number of heartbeat packets sent to the peer node by the master NetScaler device during the last measurement period.	Number	
Propagation timed out	Indicates the number of times the command propagation from the primary to the secondary nodes timed out during the last measurement period.	Number	Command propagation is a feature of the NetScaler appliance that ensures that the commands run on the primary NetScaler appliance of the high availability setup are automatically run on the secondary NetScaler appliance. When you run a command on the primary appliance, this feature ensures that the command runs on the secondary appliance before it runs on the primary appliance. If command execution fails on the secondary or times out when executing

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>on the secondary, it may cause a non-sync between the configuration of the primary and the secondary.</td></tr>
<tr><td>Synch failure</td><td>Indicates the number of times the configuration of the primary and secondary nodes failed to synchronize during the last measurement period.</td><td>Number</td><td>Synchronization is a process of duplicating the configuration of the primary node on the secondary node. The purpose of synchronization is to ensure that there is no loss of configuration information between the primary and the secondary nodes, regardless of the number of failovers that occur.
Synchronization is triggered by either of the following circumstances:
a. The secondary node in an HA setup comes up after a restart.
b. The primary node becomes secondary after a failover.
A synchronization failure results in mismatched configuration. It can be caused by a mismatch in the Remote Procedural Call (RPC) password on the two nodes forming the high availability pair.</td></tr>
<tr><td>Has the master role been changed?</td><td>Indicates whether/not the state of the monitored appliance changed since the last measurement period.</td><td></td><td>The values that this measure can report and their corresponding numeric equivalents are shown in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>No</td><td>0</td></tr><tr><td>Yes</td><td>1</td></tr></table>
If the state of the appliance changes</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			from primary to secondary or secondary to primary, then the value of this measure will be *Yes*. If the state of the appliance remains unchanged since the previous measurement period, then the value of this measure will be *No*. **Note:** By default, this measure reports the above-mentioned **Measure Value**s while indicating whether/not the state of the master is changed. However, in the graph of this measure, the states will be represented using the corresponding numeric equivalents only.

3.1.9 SSL Certificates Test

In the Citrix NetScaler VPX/MPX appliance, SSL certificates are used to establish the secure connection between the users and the web applications that are accessed by the users via the appliance. The SSL certificates are important to maintain the confidentiality of data and also organization's reputation and integrity. The SSL certificates are small data files that digitally bind a cryptographic key to organization's details. With the SSL certificates, data is encrypted prior to being transmitted via Internet, and the encrypted data can be decrypted only by the application server to which you actually send it. This ensures that the information you transmit will not be stolen. Typically, the SSL certificates are prepared with a specific validity time beyond which the connections will be no longer secure. If the certificates are suddenly expires, the users will no longer be able to access the applications and encounter the applications with the expired SSL certificate. To avoid this, administrators should proactively identify certificates nearing expiry and renew the certificates before expiry. This is where the **SSL Certificates** test helps administrators!

This test monitors all the SSL certificates that have been configured for the Citrix NetScaler VPX/MPX appliance. For each SSL certificate, this test captures the expiry date of the SSL certificates, computes how long each certificate will remain valid, and proactively alerts administrators if any certificate is nearing expiry. In addition, this test also reports the current status of each certificate and checks whether the expiry monitor for each SSL certificate has been enabled or not.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every SSL certificate on the NetScaler VPX/MPX being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements of the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Days to expire</td><td>Indicates the number of days from the current day for which this SSL certificate will be valid.</td><td>Number</td><td>A high value is preferred for this measure. A low value of this measure indicates that the SSL certificate is going to be expired soon and you should update the certificate before it expires.

The detailed diagnosis of this measure reveals the key file name, the format of the certificate file and the notification period beyond which the alert will be generated.</td></tr>
<tr><td>Status</td><td>Indicates the current status of this SSL certificate.</td><td></td><td>The values that this measure can report and their numeric equivalents are listed in the table below:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Valid</td><td>0</td></tr><tr><td>Expired</td><td>1</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values to indicate the state of the SSL certificate. However, in the graph of this measure the SSL certificate state will be represented using the corresponding numeric equivalents only - i.e., 0 or 1.</td></tr>
<tr><td>Expiry monitor</td><td>Indicates whether/not the Expiry Monitor has been enabled for this SSL certificate.</td><td></td><td>The values that this measure can report and their numeric equivalents are listed in the table below:</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Enabled</td><td>0</td></tr><tr><td>Disabled</td><td>1</td></tr></table> **Note:** By default, this measure reports the above-mentioned **Measure Value**s to indicate whether the Expiry Monitor has been enabled for each SSL certificate. However, in the graph of this measure the SSL certificate state will be represented using the corresponding numeric equivalents only - i.e., 0 or 1.

The detailed diagnosis of the *Days to expire* measure reveals the file name of the SSL certificate, the key file name, the format of the certificate file and also displays the notification period beyond which the alert will be generated.

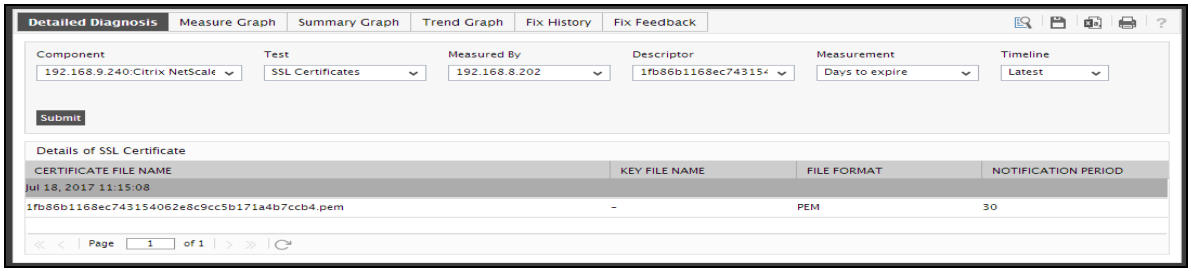

Figure 3.3: The detailed diagnosis of the Days to expire measure

3.2 The Network Layer

The tests mapped to this layer report the availability of the NetScaler appliance over the network, point you to unavailable NICs and NIC-related errors, and monitor the RNAT sessions and the VLAN traffic on the appliance.

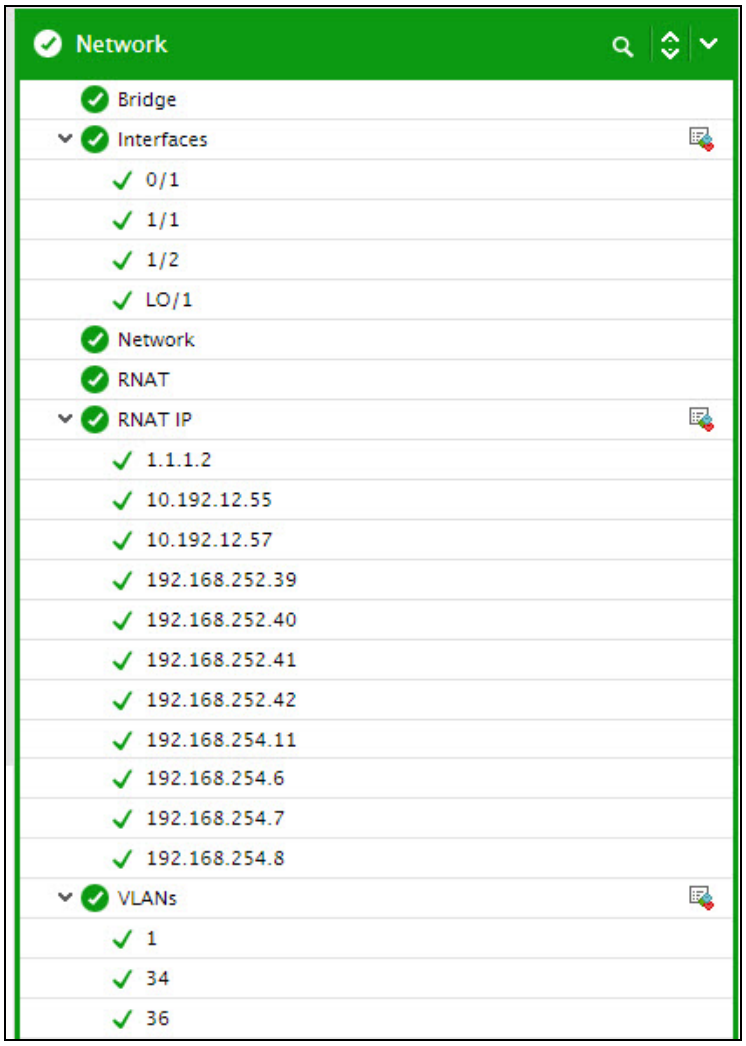

Figure 3.4: The tests mapped to the Network layer

As the **Network** test in Figure 3.4 has already been dealt with in the *Monitoring Unix and Windows Servers* document, the sections that follow will deal with the remaining tests only.

3.2.1 Bridge Test

Bridging is a technique used by the NetScaler appliance for forwarding packets between network interfaces in the VLANs configured on it. Layer 2 traffic is bridged within a port-based VLAN. By observing the movement of the bridged packets over time, administrators can understand the load on the appliance, and be proactively alerted to issues such as bridge loops and packet collisions. The **Bridge** Test enables administrators to do just that.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Bridged Mbits	Indicates the rate at which Mbits forwarded to the NetScaler interfaces during the last measurement period.	Mbps	This is a good indicator of the load on the network interfaces.
Bridged packets	Indicates the number of packets forwarded to the NetScaler interfaces during the last measurement period.	Number	
Collisions	Indicates the number of table collisions that occurred in the bridge during the last	Number	When two or more interfaces attempt to transmit a packet over the bridge at the same time, a collision occurs. Packet collisions

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		can result in the loss of packet integrity or can impede the performance of the NetScaler appliance. Ideally therefore, the value of this measure should be 0.
Loops	Indicates the number of loops that occurred in the bridge during the last measurement period.	Number	A loop occurs when there is more than one Layer 2 path between two interfaces in a VLAN. Typically, when a link aggregate channel configured on a NetScaler appliance is removed, the network interfaces bound to it induce network looks. Looping creates broadcast radiation - i.e., the accumulation of broadcast and multicast traffic on an appliance. Extreme amounts of broadcast traffic constitute a broadcast storm. A broadcast storm can consume sufficient resources so as to render the appliance unable to transport normal traffic. Ideally therefore, the value of this measure should be 0.
Interfaces muted	Indicate the number of times the interfaces were muted during the last measurement table.	Number	An interface is said to be muted if it stops transmitting and receiving packets. This could be owing to dropped packets or because there were too many MAC moves on that interface or due to a suspected configuration issue.

3.2.2 Interfaces Test

Network interfaces in the NetScaler appliance are numbered in <slot>/<port> notation. For each such network interface, this test reports the following:

- The current state of the interface;
- The duration for which the link to the interface was up;
- The bandwidth utilized by the interface and the nature of traffic handled by the interface - i.e., multicast packets, NetScaler packets, LACPDUs, etc.
- Issues experienced by the interface during transmission/reception of packets - eg., packet drops, discards, errors, stalls, hangs, and more!

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the each network interface configured on the NetScaler appliance being monitored.

Measures are also reported for an additional **Summary** descriptor.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Parameter	Description
Show Up Interface Only	If this flag is set to **Yes**, then only the NetScaler interfaces that are operational - i.e. whose MIB-II operStatus variable has a value "up" - are monitored. If this flag is set to **No**, all NetScaler interfaces that have an adminStatus of "up" will be monitored. By default, this flag is set to **No**, indicating that by default the test will monitor the interfaces that are up/down.
Show Default Down Interface	A certain number of network interfaces in the Netscaler infrastructure will by default be in down state and not operational. When this test runs for the first time, it will unnecessarily alert administrators to the down state of these interfaces. To avoid such unwanted alerts, by default the Show Default Down Interface flag is set to **No**, indicating that the test will not monitor those interfaces that have been shutdown by default. If this flag is set to **Yes**, then the test will notify administrators of the down state of such interfaces as well.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Interface state</td><td>Indicates the current status of this network interface.</td><td></td><td>This measure is not reported for the Summary descriptor.

This measure reports a value Up if the Interface is up and a value Down to indicate that the Interface is down. The numeric equivalents of these measure values have been discussed in the table below:

<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Up</td></tr>
<tr><td>1</td><td>Down</td></tr>
</table>

Note:

By default, this measure reports the above-mentioned Measure Values to indicate whether/not the NetScaler Interface is Up or Down. However, in the graph of this measure, these values will be represented using their</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>corresponding numeric equivalents only.</td></tr>
<tr><td>Link state</td><td>Indicates the current status of the logical NetScaler Interface i.e., the link.</td><td></td><td>This measure is not reported for the Summary descriptor.
This measure reports a value Up if the link is up and a value Down to indicate that the link is down. The numeric equivalents of these measure values have been discussed in the table below:
<table><tr><th>Numeric Value</th><th>Measure Value</th></tr><tr><td>0</td><td>Up</td></tr><tr><td>1</td><td>Down</td></tr></table>
Note:
By default, this measure reports the Measure Values (listed in the table above) to indicate whether the link is Up or not. However, in the graph of this measure, these values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Link uptime</td><td>Indicates the time duration for which this link was Up.</td><td>Mins</td><td>This measure is not reported for the Summary descriptor.
A high value is desired for this measure.</td></tr>
<tr><td>Link downtime</td><td>Indicates the time duration for which the link was Down.</td><td>Mins</td><td>This measure is not reported for the Summary descriptor.
Ideally, the value of this measure should be 0.</td></tr>
<tr><td>Data received</td><td>Indicates the amount of data received by this</td><td>MB</td><td>These measures serve as good indicators of the bandwidth utilized by</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	NetScaler interface over a network link during the last measurement period. For the Summary descriptor, this measure reports the total data received by the VPX instance across all its interfaces.		a network interface. Compare the values of these measures across interfaces to isolate that interface which is handling bandwidth-intensive traffic.
Data transmitted	Indicates the amount of data transmitted by this NetScaler interface over a network link during the last measurement period. For the Summary descriptor, this measure reports the total data transmitted by the VPX instance across all its interfaces.	MB	
Packets received	Indicates the number of packets received by this NetScaler interface during the last measurement period. For the Summary descriptor, this measure reports the total number of packets received by all interfaces of the VPX instance.	Number	
Packets transmitted:	Indicates the number of packets transmitted by this NetScaler interface during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
	For the Summary descriptor, this measure reports the total number of packets transmitted by all interfaces of the VPX instance.		
Multicast packets	Indicates the number of packets received on this NetScaler interface that were destined for multiple hosts during the last measurement period. For the Summary descriptor, this measure reports the total number of multicast packets received by the VPX instance across all its interfaces.	Number	
NetScaler packets	Indicates the number of packets received during the last measurement period on this NetScaler interface in which the destination MAC address is either the address of one of the NetScaler interfaces or the VMAC address configured by the user. For the Summary descriptor, this measure reports the total number of NetScaler packets received by the VPX instance across all its interfaces.	Number	
LACPDUs received	Indicates the number of the Link Aggregation	Number	LACP (802.3ad) is a control protocol that configures multiple ports into a

Measurement	Description	Measurement Unit	Interpretation
	Control Protocol Data Units (LACPDUs) that were received on the selected port during the last measurement period. For the Summary descriptor, this measure reports the total number of LACPDUs received by the VPX instance via all its interfaces.		single high-speed link. An aggregated link is referred to as a channel. On this aggregated link, the LACP-enabled appliances exchange LACP data units (LACPDU). After LACP is enabled on a port, the port sends LACPDUs to notify the remote system of its system LACP priority, system MAC address, port LACP priority, port number, and operational key. Upon receipt of an LACPDU, the remote system compares the received information with the information received on other ports to determine the ports that can operate as selected ports. This allows the two systems to reach agreement on the states of the related ports.
LACPDUs transmitted	Indicates the number of the Link Aggregation Control Protocol Data Units (LACPDUs) that were transmitted by the selected port during the last measurement period. For the Summary descriptor, this measure reports the total number of LACPDUs transmitted by the VPX instance via all its interfaces.	Number	
Dropped packets received	Indicates the number of error-free outbound packets that were dropped when received by this NetScaler interface during the last measurement period. For the Summary descriptor, this measure reports the total number of outbound packets that were dropped during reception by the	Number	Ideally, the value of this measure should be zero. The packets may be dropped due to a VLAN mismatch, an oversized packet or a disabled network interface card.

Measurement	Description	Measurement Unit	Interpretation
	VPX instance across all its interfaces.		
Dropped packets transmitted	Indicates the number of error-free outbound packets that were dropped during transmission by this NetScaler interface during the last measurement period. For the Summary descriptor, this measure reports the total number of outbound packets that were dropped during transmission by the VPX instance across all its interfaces.	Number	
Inbound packets discarded	Indicates the number of error-free inbound packets that were discarded by this NetScaler interface during the last measurement period. For the Summary descriptor, this measure reports the total number of inbound packets that were discarded by the VPX instance across all its interfaces.	Number	Ideally, the value of these measures should be zero. These measures may report a non-zero value when an interface runs short of resources – e.g., NIC buffers.
Outbound packets discarded	Indicates the number of error-free outbound packets that were discarded by this NetScaler interface during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
	For the Summary descriptor, this measure reports the total number of outbound packets that were discarded by the VPX instance across all its interfaces.		
Error packets received	Indicates the number of packets that were received with errors by this NetScaler interface during the last measurement period. For the Summary descriptor, this measure reports the total number of packets that were received with errors by the VPX instance across all its interfaces.	Number	Ideally, the value of this measure should be zero.
Error packets transmitted	Indicates the number of packets that were transmitted with errors by this NetScaler interface during the last measurement period. For the Summary descriptor, this measure reports the total number of packets that were transmitted with errors by the VPX instance across all its interfaces.	Number	
MAC moves registered	Indicates the number of MAC moves that were registered between ports during the last	Number	

Measurement	Description	Measurement Unit	Interpretation
	measurement period. For the Summary descriptor, this measure reports the total number of MAC moves registered by the VPX instance across all its interfaces.		
NIC hangs	Indicates the number of times the Network Interface Card (NIC) hanged. For the Summary descriptor, this measure reports the total number of times the interfaces of this NetScaler VPX instance hanged.	Number	Ideally, the value of this measure should be zero. A non-zero value is reported when the NetScaler detects an error on the transmission or reception path of the NIC.
Stalls status	Indicates the number of system detected stalls that occurred during transmission or reception of packets on this Network Interface Card during the last measurement period. For the Summary descriptor, this measure reports the total number of the VPX instance acrsystem detected stalls that occurred on the VPX instance across all its interfaces.	Number	When the status is not updated within 0.8 seconds by the NIC hardware, the NIC is said to be in a status stall state.
Stalls received	Indicates the number of system registered stalls that occurred during reception of packets on this NetScaler Interface.	Number	When the link is UP for more than 10 minutes and packets are transmitted, but no packets are received for 16 seconds, the NIC is said to be in a receive stall state. This commonly

Measurement	Description	Measurement Unit	Interpretation
	For the Summary descriptor, this measure reports the total number of system registered stalls that occurred on the VPX instance during reception across all its interfaces.		occurs in lab environments when no packets, including spanning tree, are received on the wire.
Stalls transmitted	Indicates the number of system detected stalls that occurred during transmission of packets on this Network Interface Card. For the Summary descriptor, this measure reports the total number of system registered stalls that occurred on the VPX instance during packet transmission across all its interfaces.	Number	When a packet posted for transmission is not transmitted in 4 seconds, the NIC is said to be in a transmit stall state.
Error disables	Indicates the number of times this NetScaler Interface has been disabled. For the Summary descriptor, this measure reports the total number of times the network interfaces of the target VPX instance were disabled.	Number	Ideally, the value of this measure should be zero. The NetScaler interface would be disabled due to errors such as the stall in the transmission or reception of packets.
Duplux mismatches	Indicates the number of duplux mismatches that were detected on this NetScaler Interface.	Number	Ideally, the value of this measure should be zero.

Measurement	Description	Measurement Unit	Interpretation
	For the Summary descriptor, this measure reports the total number of duplux mismatches that were detected on the interfaces of the target VPX instance.		
Link re-initializations	Indicates the number of times the network link has been re-initialized. For the Summary descriptor, this measure reports the total number of time network links were re-initialized on the target VPX instance across all its interfaces.	Number	Ideally, the value of this measure should be zero.
NIC muted	Indicates the number of times this NetScaler interface was muted i.e., this NetScaler interface stopped transmitting or receiving the packets. For the Summary descriptor, this measure reports the total number of times the interfaces on the VPX instance were muted.	Number	Ideally, the value of this measure should be zero. The NetScalrer interface may be muted when there were too many dropped packets or when there were too many MAC moves on this NetScaler interface which would occur due to a suspected configuration issue.
Input throughput	Indicates the input throughput of this interface during the last measurement period. For the Summary descriptor, this measure reports the total input	Mbps	Compare the value of this measure across interfaces to know which interface has the lowest input throughput.

Measurement	Description	Measurement Unit	Interpretation
	throughput of the VPX instance across all its interfaces.		
Output throughput	Indicates the output throughput of this interface during the last measurement period. For the Summary descriptor, this measure reports the total output throughput of the VPX instance across all its interfaces.	Mbps	Compare the value of this measure across interfaces to know which interface has the lowest output throughput.
License allowed maximum throughput	Indicates the maximum throughput that the NetScaler license allowed.	Mbps	**This measure is reported for the Summary descriptor only.**
License throughput	Indicates the percentage of maximum throughput that the NetScaler VPX instance is using.	Percent	**This measure is reported for the Summary descriptor only.** The value of this measure is computed using the following formula: (Input throughput / License allowed maximum throughput)*100 A value close to 100% is a cause for concern, as it implies that the VPX instance is rapidly exhausting its licensed throughput limit.

3.2.3 RNAT Test

With 'Reverse Network Address Translation' (RNAT), the NetScaler replaces the source IP addresses in the packets generated by hosts in the configured subnet with the configured, NAT IP address(es). Typically, RNAT is used to allow servers configured with private non-routable IP addresses to initiate connections to the Internet.

Using this test, you can measure the load imposed by the RNAT sessions to the NetScaler appliance.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Active sessions	Indicates the number of current RNAT active sessions	Number	This is a good indicator of the RNAT session load on the NetScaler appliance.
Request for connections sent	Indicates the number of requests for connections that were sent during the RNAT sessions since the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the number of bytes that were received during RNAT sessions since the last measurement period.	MB	These measures are good indicators of the load imposed by the RNAT sessions on the NetScaler appliance.
Data transmitted	Indicates the number of bytes that were transmitted during RNAT sessions since the last measurement period.	MB	
Packets received	Indicates the number of packets that were received during RNAT sessions since the last measurement period.	Number	These measures are good indicators of the load imposed by RNAT sessions on the NetScaler appliance.
Packets transmitted	Indicates the number of packets that were transmitted during RNAT sessions since the last measurement period.	Number	

3.2.4 RNAT IP Test

With 'Reverse Network Address Translation' (RNAT), the NetScaler replaces the source IP addresses in the packets generated by hosts in the configured subnet with the configured, NAT IP address(es). Typically, RNAT is used to allow servers configured with private non-routable IP addresses to initiate connections to the Internet.

In the event of an overload, the RNAT test will enable you to determine whether/not the RNAT sessions are imposing excessive load on the NetScaler appliance. If so, then, you can use the RNAT IP test to isolate the client IP address that is generating maximum load.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each client IP address that initiates RNAT sessions to the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Active sessions	Indicates the number of current RNAT active sessions originating from this client IP address.	Number	This is a good indicator of the RNAT session load generated by a client IP address on the NetScaler appliance. Compare the value of this measure across IP addresses to identify the client that is generating the maximum session load.
Request for connections sent	Indicates the number of requests for connections that were sent during the RNAT sessions that originated from this client IP address since the last measurement period.	Number	
Data received	Indicates the number of bytes that were received	MB	These measures are good indicators of the load imposed by RNAT

Measurement	Description	Measurement Unit	Interpretation
	during RNAT sessions that originated from this client IP address since the last measurement period.		sessions started by a client IP address on the NetScaler appliance. By comparing the value of these measures across IP addresses, you can quickly and accurately isolate the IP address that is responsible for the maximum data load.
Data transmitted	Indicates the number of bytes that were transmitted during RNAT sessions that originated from this client IP address since the last measurement period.	MB	
Packets received	Indicates the number of packets that were received during RNAT sessions that originated from this client IP address since the last measurement period.	Number	These measures are good indicators of the load imposed by RNAT sessions started by a client IP address on the NetScaler appliance. By comparing the value of these measures across IP addresses, you can quickly and accurately isolate the IP address that is responsible for the maximum packet load.
Packets transmitted	Indicates the number of packets that were sent during RNAT sessions that originated from this client IP address since the last measurement period.	Number	

3.2.5 VLANs Test

A NetScaler appliance supports Layer 2 port and IEEE 802.1q tagged VLANs. VLAN configurations are useful when you need to restrict traffic to certain groups of stations. You can configure a network interface as a part of multiple VLANs by using IEEE 802.1q tagging. You can configure VLANs and bind them to IP subnets. The NetScaler then performs IP forwarding between these VLANs (if it is configured as the default router for the hosts on these subnets).

The NetScaler supports the following types of VLANs:

- **Port-Based VLAN**: The membership of a port-based VLAN is defined by a set of network interfaces that share a common, exclusive Layer 2 broadcast domain. You can configure multiple port-based VLANs. By default, all network interfaces on the NetScaler are members of VLAN 1. If

you apply 802.1q tagging to the port, the network interface belongs to a port-based VLAN. Layer 2 traffic is bridged within a port-based VLAN, and Layer 2 broadcasts are sent to all members of the VLAN if Layer 2 mode is enabled. When you add an untagged network interface as a member of a new VLAN, it is removed from its current VLAN.

- **Default VLAN**: By default, the network interfaces on the NetScaler are included in a single, port-based VLAN as untagged network interfaces. This VLAN is the default VLAN. It has a VLAN ID (VID) of 1. This VLAN exists permanently. It cannot be deleted, and its VID cannot be changed. When you add a network interface to a to a different VLAN as an untagged member, the network interface is automatically removed from the default VLAN. If you unbind a network interface from its current port-based VLAN, it is added to the default VLAN again.
- **Tagged VLANs**: 802.1q tagging (defined in the IEEE 802.1q standard) allows a networking device (such as the NetScaler) to add information to a frame at Layer 2 to identify the VLAN membership of the frame. Tagging allows network environments to have VLANs that span multiple devices. A device that receives the packet reads the tag and recognizes the VLAN to which the frame belongs.

By continuously tracking the traffic over the VLANs, you can quickly identify the VLAN that is handling the maximum traffic. The **VLANs** test does just that.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each VLAN configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface

Parameter	Description
	Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the amount of data received over this VLAN during the last measurement period.	MB	Compare the value of these measures across VLANs to identify the VLAN that is experiencing the maximum traffic.
Data transmitted	Indicates the amount of data transmitted over this VLAN during the last measurement period.	MB	
Packets received	Indicates the number of packets received over this VLAN during the last measurement period.	Number	Compare the value of these measures across VLANs to identify the VLAN that is experiencing the maximum traffic.
Packets transmitted	Indicates the number of packets that were transmitted over this VLAN since the last measurement period.	Number	
Broadcast packets sent and received	Indicates the total broadcast packets that were sent and received on this VLAN during the last measurement period.	Number	
Packets dropped	Indicates the number of inbound packets that were dropped upon reception by this VLAN during the last measurement period.	Number	Ideally, the value of this measure should be zero.

3.3 The Protocols Layer

This layer monitors the traffic flowing into and out of the NetScaler appliance, protocol-wise.

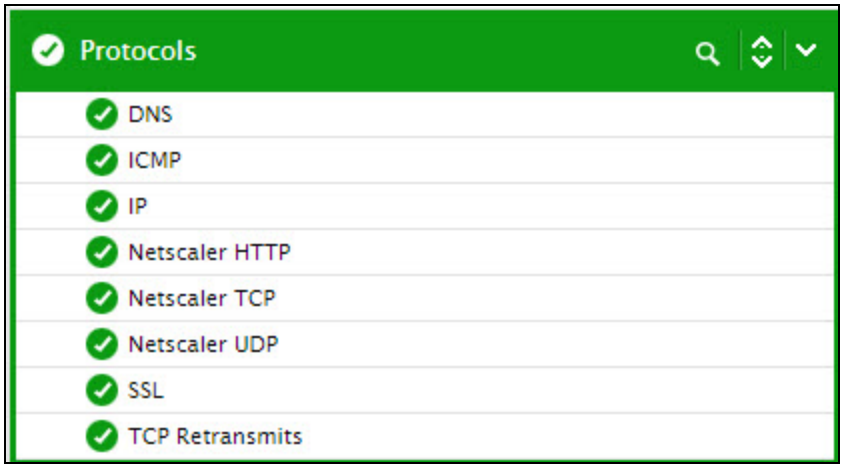

Figure 3.5: The tests mapped to the Protocols layer

3.3.1 IP Test

This test monitors the IP (Internet Protocol) traffic on the NetScaler appliance, measures the current IP load on the NetScaler device, and promptly captures IP-related anomalies.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.

Parameter	Description
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the amount of IP data received during the last measurement period.	MB	
Data transmitted	Indicates the amount of IP data transmitted during the last measurement period.	MB	
Packets transmitted	Indicates the IP packets transmitted during the last measurement period.	Number	
Packets received	Indicates the number of IP packets received during the last measurement period.	Number	
Routed IP packets	Indicates the number of IP packets that were routed during the last measurement period.	Number	
Routed IP data	Indicates the amount of IP data that was routed during the last measurement period.	Mbits	
IP fragments received	Indicates the number of IP packet fragments received during the last measurement period.	Number	The Internet Protocol (IP) implements fragmentation, so that packets may be formed that can pass through a link with a smaller maximum

Measurement	Description	Measurement Unit	Interpretation
			transmission unit (MTU) than the original datagram size. IP fragments received by the NetScaler appliance will have to be reassembled by the appliance before they are forwarded to the virtual server they are meant for. IP fragmentation and reassembly operations can cause the NetScaler appliance to incur heavy overheads. To alleviate this, you can enable the Path MTU discovery algorithm of the Netcaler apliance for dynamically learning the MTU (Maximum Transmission Unit) of any Internet path. The discovered Path MTU is then used by the TCP/IP stack on the NetScaler to create packets of this size (or smaller). Path MTU Discovery is disabled by default.
IP packets successful reassembly	Indicates the number of IP packets that were successfully reassembled by the NetScaler appliance during the last measurement period.	Number	IP fragments received by the NetScaler appliance will have to be reassembled by the appliance before they are forwarded to the virtual server they are meant for. Ideally therefore, the value of this measure should be high. A low value indicates too many reassembly failures. This can occur when there is a checksum failure, an identification field mismatch, or when one of the fragments is missing.
IP packets attempts to reassembly	Indicates the number of fragmented IP packets that the NetScaler attempted to reassemble during the last measurement period.	Number	
IP address lookups	Indicates the number of IP address lookups	Number	When an IP packet is received on a non-

Measurement	Description	Measurement Unit	Interpretation
	performed by the NetScaler during the last measurement period.		established session, the NetScaler checks if the destination IP address is one of the NetScaler owned IP addresses. The NetScaler-owned IP Addresses—NetScaler IP Address (NSIP), Virtual IP Addresses (VIPs), Subnet IP Addresses (SNIPs), Mapped IP Addresses (MIPs), and Global Server Load Balancing Site IP Addresses (GSLBIPs)—exist only on the NetScaler appliance. The NSIP uniquely identifies the NetScaler on your network, and it provides access to the appliance. The IP address lookup failure occurs when the destination IP address of the packet does not match any of the NetScaler owned IP addresses. Naturally therefore, the value of IP address lookup failures measure should be very low.
IP address lookup failures:	Indicates the number of IP address lookups that failed during the last measurement period.	Number	
Percent of IP address lookup success	Indicates the percentage of IP address lookups that were successfully performed by the NetScaler during the last measurement period.	Percent	A high value is desired for this measure.
Percent of IP address lookup failure	Indicates the percentage of IP address lookups that failed during the last measurement period.	Percent	The IP address lookup failure occurs when the destination IP address of the packet does not match any of the NetScaler owned IP addresses. The value of this measure is computed using the following formula: (IP address lookup failures/IP address lookups) * 100.0 A low value is is desired for this measure.

Measurement	Description	Measurement Unit	Interpretation
UDP fragments forwarded	Indicates the number of UDP fragments that were forwarded to the client or server during the last measurement period.	Number	
TCP fragments forwarded	Indicates the number of TCP fragments that were forwarded to the client or server during the last measurement period.	Number	
Fragmented packets created	Indicates the number of UDP fragments that were forwarded to the client or server during the last measurement period.	Number	
Bad IP checksum	Indicates the number of IP packets that were received with IP checksum error during the last measurement period.	Number	
Unsuccessful reassambly	Indicates the number of received IP packets that could not be reassembled during the measurement period.	Number	IP fragments received by the NetScaler appliance will have to be reassembled by the appliance before they are forwarded to the virtual server they are meant for. Ideally therefore, the value of this measure should be low. A high value indicates too many reassembly failures. This can occur when there is a checksum failure, an identification field mismatch, or when one of the fragments is missing.
Reassembled data too big	Indicates the number of IP packets whose data length exceeds the Ethernet packet data length of 1500 bytes after being reassembled by the	Number	

Measurement	Description	Measurement Unit	Interpretation
	NetScaler during the last measurement period.		
Zero fragment length received	Indicates the number of IP packets that were received with a fragment length of 0 bytes during the last measurement period.	Number	
Duplicate fragments received	Indicates the number of duplicate IP fragments received during the last measurement period.	Number	This can occur when the acknowledgement was not received within the expected time.
Out of order fragments received	Indicates the number of IP fragments that were received in out of order condition during the last measurement period.	Number	When a datagram is fragmented, each fragment becomes its own datagram and is routed to the NetScaler independently of any other datagrams. This is why, the original datagram often arrives at the NetScaler out of order.
Unknown destination received	Indicates the number of IP packets received with the destination IP address not reachable or not owned by the NetScaler during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Bad transport	Indicates the number of packets received in which the protocol specified in the IP header is unknown to the NetScaler during the last measurement period.	Number	Ideally, the value of this measure should be 0.
VIP down	Indicates the number of IP packets received by the NetScaler when the Virtual IP (VIP) is down during the last measurement period.		A VIP is a public IP address to which a client sends requests. The NetScaler receives these request packets when all the services bound to the VIP are down or the VIP is manually disabled.

Measurement	Description	Measurement Unit	Interpretation
Fix header failure	Indicates the number of received packets with errors in one or more fields of the IP header during the last measurement period.	Number	Ideally, this value should be 0.
Time-to-live expired during transit	Indicates the number of packets for which the time-to-live (TTL) expired during transit during the last measurement period.	Number	These packets are dropped.
Max Non-TCP clients	Indicates the total number of times during the last measurement period NetScaler tried to open a new connection to a service having maximum number of allowed open client connections.	Number	
Unknown services	Indicates the number of IP packets received on a port or service that is not configured for the NetScaler during the last measurement period.	Number	
Land-attack packets received	Indicates the number of Land-attack packets received by the NetScaler during the last measurement period.	Number	The Land-attack packets are spoofed packets that are designed to attack systems. A Land Attack consists of a stream of TCP SYN packets that have the source IP address and TCP port number set to the same value as the destination address and port number (i.e., that of the attacked host).
Invalid IP header size packets received	Indicates the number of IP packets received with an invalid header size during the last measurement period.	Number	The IP header size may be termed as invalid due to an invalid data length in the header or when the value in the length field and the actual data length does not match.

Measurement	Description	Measurement Unit	Interpretation
Invalid IP packet size received	Indicates the number of IP packets received by the NetScaler with invalid packet size during the last measurement period.	Number	
Truncated IP packet received	Indicates the number of truncated IP packets received during the last measurement period.	Number	An overflow in the routers along the path can truncate the IP packets.
Truncated Non-IP packets received	Indicates the number of truncated non-IP packets received during the last measurement period.	Number	
Zero next hop	Indicates the number of IP packets received with a 0 value in the next hop field during the last measurement period.	Number	These packets are dropped.
Packets with length greaterthen 1514 received	Indicates the number of IP packets received with a length greater than the normal MTU (maximum transmission unit) of 1514 bytes during the last measurement period.	Number	
Packets with bad MAC sent	Indicates the number of IP packets transmitted with a bad MAC address during the last measurement period.	Number	

3.3.2 NetScaler HTTP Test

This test monitors HTTP connections handled by the NetScaler device, and reveals whether all HTTP requests have been responded to, and whether any incomplete requests/responses have been received/sent by the NetScaler.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each HTTP on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total rate of requests	Indicates the rate at which HTTP requests were received.	Number/Sec	These measures are effective indicators of the workload on the NetScaler device.
Total requests	Indicates the total number of HTTP requests (inclusive of HTTP v1.0 requests and HTTP v1.1 requests) that were received during the last measurement period.	Number	You can even compare the values of these measures to determine whether/not all requests have been responded to. Request processing bottlenecks can thus be isolated.
Total responses	Indicates the total number	Number	

Measurement	Description	Measurement Unit	Interpretation
	of HTTP responses (inclusive of HTTP v1.0 responses and HTTP v1.1 responses) that were received during the last measurement period.		
Total rate of responses	Indicates the rate at which HTTP responses were received.	Number/Sec	
GETs requests	Indicates the number of requests that were received using the GET method during the last measurement period.	Number	
Percent of GETs requests	Indicates the percentage of HTTP requests received using the GET method during the last measurement period.	Percent	
POSTs requests	Indicates the number of requests that were received using the POST method during the last measurement period.	Number	
Percent of POSTs requests	Indicates the percentage of HTTP requests received using the POST method during the last measurement period.	Percent	
Other methods requests	Indicates the number of HTTP requests that were received using methods other than GET and POST during the last measurement period.	Number	Some of the well-defined HTTP methods are HEAD, PUT, DELETE, OPTIONS, and TRACE. User-defined methods are also counted.
Percent of Other methods requests	Indicates the percentage	Percent	

Measurement	Description	Measurement Unit	Interpretation
	of HTTP requests received using methods other than the GET and POST methods during the last measurement period.		
HTTP/1.0 requests	Indicates the number of HTTP v 1.0 requests received by the NetScaler device during the last measurement period.	Number	Since HTTP 1.0 connections are not capable of providing information about the client's ability to accept compressed data, which is one of the features of the NetScaler devices, it is important to be able to monitor the number of HTTP 1.0 connections relative to the total connections. By comparing the values of these two measures, you can determine whether/not all HTTP/1.0 requests have been responded to by the NetScaler appliance.
HTTP/1.0 responses	Indicates the number of HTTP v 1.0 responses sent by the NetScaler device during the last measurement period.		
HTTP/1.1 requests	Indicates the number of HTTP v 1.1 requests received by the NetScaler device during the last measurement period.	Number	By comparing the values of these two measures, you can determine whether/not all HTTP/1.1 requests have been responded to by the NetScaler appliance.
HTTP/1.1 responses	Indicates the number of HTTP v 1.1 responses sent by the NetScaler device during the last measurement period.	Number	
Content-length requests	Indicates the number of HTTP requests with the Content-length field of the HTTP header set during the last measurement period.	Number	NetScaler compression module treats server response as a data stream. It accumulates server response packets until certain conditions are met. Whence the conditions are met, a batch job would be triggered to compress the accumulated data and the compressed data would be output

Measurement	Description	Measurement Unit	Interpretation
			subsequently. For a server response that only triggers a single batch job, NetScaler would output a **Content-length** response. This is same for both HTTP 1.0 and 1.1 client types. When an HTTP client is reading a response message from a server or when an HTTP server is reading a request from a client, the client/server (as the case may be) needs to know when it has reached the end of the message. This is particularly important with persistent (keep alive) connections, because a connection can only be re-used by another HTTP transaction after the request message has been read and the response message has been fully received. One of the ways by which the HTTP client/server indicates the end of a request/response message is by using the **Content-Length Header**. The length of the content after the request/response headers can be specified in bytes with the **Content-Length** header. Using the values of these measures, you can figure out how many large requests were processed and responded to by the NetScaler device, and thus assess the workload on the device.
Content-length responses	Indicates the number of HTTP responses that were sent with the Content-length field of the HTTP header set during the last measurement period.	Number	
Percent of content-length responses:	Indicates the percentage of responses sent out by the NetScaler device with the Content-length field of the HTTP response header set.	Percent	
Chunked requests	Indicates the number of HTTP requests with the Transfer-Encoding field of	Number	When an HTTP client is reading a response message from a server or

Measurement	Description	Measurement Unit	Interpretation
	the HTTP header set to chunked during the last measurement period.		when an HTTP server is reading a request from a client, the client/server (as the case may be) needs to know when it has reached the end of the message. This is particularly important with persistent (keep alive) connections, because a connection can only be re-used by another HTTP transaction after the request message has been read and the response message has been fully received. One of the ways by which the HTTP client/server indicates the end of a request/response message is by using **Chunked Encoding**. Chunked encoding allows HTTP messages to be broken up into several parts. Chunking is most often used by the server for responses, but clients can also chunk large requests. If the Transfer-Encoding field of an HTTP request/response is set to chunked, then the client/server (as the case may be) starts sending the HTTP request/response before knowing its total length. The client/server then breaks the request/response into chunks and sends them in sequence, inserting the length of each chunk before the actual data. The message ends with a chunk of size zero. Using the values of these measures, you can figure out how many large requests were processed and responded to by the NetScaler device, and thus assess the workload on the device.
Chunked responses	Indicates the number of HTTP responses that were sent with the Transfer-Encoding field of the HTTP header set to chunked during the last measurement period.	Number	
Percent of chunked responses	Indicates the percentage of HTTP responses that were sent with the Transfer-Encoding field of the HTTP header set to chunked during the last measurement period.	Percent	
Multi-part responses	Indicates the number of HTTP multi-part	Number	In multi-part responses, one or more entities are encapsulated within the

Measurement	Description	Measurement Unit	Interpretation
	responses sent during the last measurement period.		body of a single message.
Percent of multi-part responses	Indicates the percentage of HTTP multi-part responses that were sent during the last measurement period.	Percent	
FIN-terminated responses	Indicates the number of FIN-terminated responses sent during the last measurement period.	Number	In FIN-terminated responses, the server finishes sending the data and closes the connection. For large server response (or any other condition) that would trigger multiple compression batch jobs, NetScaler would output FIN terminated response without Content-Length header (RFC compliant). Because compression is performed on batch basis, NetScaler can't determine length of compressed data until the last batch is done. Using FIN-terminated response, instead of chunked encoding, maximizes NetScaler's compatibility with HTTP 1.0 client.
Percent of FIN-terminated responses	Indicates the percentage of FIN-terminated responses that were sent during the last measurement period.	Percent	
Request data received	Indicates the amount of HTTP data received during the last measurement period.	MB	
Request data transmitted	Indicates the amount of HTTP data transmitted during the last measurement period.	MB	
Response data received	Indicates the amount of data received as responses during the last measurement period.	MB	
Response data transmitted	Indicates the amount of data sent as responses	MB	

Measurement	Description	Measurement Unit	Interpretation
	during the last measurement period.		
Incomplete headers	Indicates the number of HTTP requests and responses that were received with the HTTP header spanning more than one packet, during the last measurement period.	Number	Ideally, the value of this measure should be 0. A high value for this measure may warrant an investigation. You may begin your investigation by determining the type of headers that were incomplete most often - the request headers or response headers? For this, compare the values of the *Incomplete request headers* and *Incomplete response headers* measures.
Incomplete request headers	Indicates the number of HTTP requests that were received with the HTTP header spanning more than one packet, , during the last measurement period.	Number	
Incomplete response headers	Indicates the number of HTTP responses that were sent with the HTTP header spanning more than one packet, during the last measurement period.	Number	
HTTP 500 server-busy responses:	Indicates the current status of HTTP responses received from the NetScaler appliance.	Number	Response status codes beginning with the digit "5" indicate cases in which the HTTP server - in this case. the NetScaler appliance - is aware that it has encountered an error or is otherwise incapable of performing the request.
Large/Invalid messages	Indicates the number of large or invalid requests and responses received during the last measurement period.	Number	Ideally, the values of these measures should be 0.
Large/Invalid chunk requests	Indicates the large or invalid requests received in which the Transfer-Encoding field of the HTTP	Number	Ideally, the values of these measures should be 0.

Measurement	Description	Measurement Unit	Interpretation
	header has been set to chunked.		
Large/Invalid content-length	Large or invalid requests received in which the Content-length field of the HTTP header has been set.	Number	

3.3.3 ICMP Test

Use this test to monitor the ICMP traffic on the NetScaler and to understand how well the NetScaler handles the traffic. The metrics reported by this test promptly capture ICMP rate threshold violations and thus reveal a potential ICMP overload on the NetScaler appliance. In addition, the test sends out instant alerts to administrators when ICMP-related errors are detected.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
ICMP rate threshold	Indicates the limit of ICMP packets handled every 10 milliseconds.	Pkts/Sec	This threshold is configurable. Once this threshold is violated, subsequent ICMP packets will be dropped by the appliance. You are hence advised to configure this limit based on the current and anticipated ICMP traffic to the NetScaler appliance. To view the configured threshold, use the show ns rateControl command. To set/alter this limit, use the set rateControl command.
ICMP data received	Indicates the amount of ICMP data received during the last measurement period.	MB	These measures are good indicators of the ICMP load on the NetScaler appliance.
ICMP data transmitted	Indicates the amount of ICMP data transmitted during the last measurement period.	MB	
ICMP packets received	Indicates the number of ICMP packets received during the last measurement period.	Number	
ICMP packets transmitted	Indicates the number of ICMP packets transmitted during the last measurement period.	Number	
ICMP echo received	Indicates the number of ICMP “Echo Request” and “Echo Reply” packets received during the last measurement period.	Number	The echo request is an ICMP message whose data is expected to be received back in an echo reply ("ping"). The host must respond to all echo requests with an echo reply containing the exact data received in the request message.
ICMP echo replies received:	Indicates the number of ICMP echo replies received during the last	Number	

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		
ICMP echo replies transmitted	Indicates the number of ICMP echo replies transmitted during the last measurement period.	Number	
ICMP port unreachable received	Indicates the number of times the “ICMP Port Unreachable” error message was received during the last measurement period.	Number	The ICMP Port Unreachable error is generated when there is no service running on the port. Ideally, the value of these measures should be 0.
ICMP port unreachable transmitted	Indicates the number of times the “ICMP Port Unreachable” error message was received during the last measurement period.		
Need fragmentation received	Indicates the number of times the “ICMP Fragmentation Needed” error message was received for the ICMP packets during the last measurement period.	Number	This measure tracks the ICMP Fragmentation Needed error messages received for packets that must be fragmented but Don't Fragment is specified in the header.
ICMP rate threshold exceeded	Indicates the number of times the value reported by the ICMP rate threshold measure has been violated.	Number	A high value of this measure indicates that the ICMP rate threshold has been violated often. When this happens, you must first ensure that the ICMP packets received are genuine. If they are genuine, then you must increase the current rate threshold. Note that if the rate threshold is violated, then the appliance will drop subsequent ICMP packets it receives. To assess the impact of the threshold

Measurement	Description	Measurement Unit	Interpretation
			violation, use the ICMP packets dropped measure. If the value of this measure is very high, its a clear call for a change in the rate threshold. To view the configured threshold, use the show ns rateControl command. To set/alter this limit, use the set rateControl command.
ICMP packets dropped	Indicates the number of ICMP packets that were dropped during the last measurement period because the rate threshold was violated.	Number	A high value is a cause for concern, and presents a strong case for changing the rate threshold.
Bad ICMP checksum	Indicates the number of ICMP Fragmentation Needed error messages received with an ICMP checksum error in the last measurement period.	Number	Ideally, the value of this measure should be zero.
PMTU non-first IP fragments	Indicates the number of “ICMP Fragmentation Needed” error messages received for an IP fragment other than the first one upon Path MTU Discovery during the last measurement period.	Number	NetScalers have a feature called Path MTU Discovery, which is actually a common feature on most networking devices. Path MTU Discovery allows a networking device such as the NetScaler, or routers and switches, to determine the largest packet size allowed along an arbitrary network path. This enables network traffic to flow correctly from one endpoint to another, without any of the traffic being dropped. The IP protocol has a mechanism for signaling that datagrams are too large to pass through an interface on a network path – when a datagram is received on a router or Layer 3 switch

Measurement	Description	Measurement Unit	Interpretation
			interface that is larger than the interface's MTU, the device sends an Internet Control Message Protocol (ICMP) message to the previous hop device indicating that the datagram needs to be fragmented in order to pass through that interface, as well as the MTU of the interface. The previous hop device breaks the datagram into pieces that are small enough to pass the next hop interface's MTU. While this mechanism is usually sufficient to allow traffic to continue normally, it does have some drawbacks. The increase in the number of datagrams from fragmenting means Layer 3 routing has that much more work to do in routing decisions. There is always the possibility that another device further along the network path has an interface with an even smaller MTU and requires further fragmentation, and when the destination device for these datagrams eventually receives them, all datagram fragments must be received so it can be reassembled correctly. If all fragments are not received, the entire original datagram is dropped and must be retransmitted by the sending station (and will probably be fragmented in transit again). This is why, the value of this measure should be kept at a minimum.
PMTU invalid body length received	Indicates the number of "ICMP Fragmentation Needed" error messages	Number	

Measurement	Description	Measurement Unit	Interpretation
	received for invalid body length of the packets determined by the Path MTU Discovery during the last measurement period.		
PMTU no TCP connections	Indicates the number of "ICMP Fragmentation Needed" error messages received for TCP packets during the last measurement period.	Number	The state of the connection for these packets is not maintained on the NetScaler appliance.
PMTU no UDP connections	Indicates the number of "ICMP Fragmentation Needed" error messages received for UDP packets during the last measurement period.	Number	The state of the connection for these packets is not maintained on the NetScaler appliance.
PMTU invalid TCP sequence number received	Indicates the number of "ICMP Fragmentation Needed" error messages received for the packets containing an invalid TCP address determined by the Path MTU Discovery during the last measurement period.	Number	
Invalid next MTU value received	Indicates the number of "ICMP Fragmentation Needed" error messages received for the packets in which the Maximum Transmission Unit (MTU) for the next hop was out of range during the last measurement period.	Number	The range for the MTU is 576-1500.
Next MTU greater than current MTU	Indicates the number of "ICMP Fragmentation	Number	

Measurement	Description	Measurement Unit	Interpretation
	Needed" error messages received in which the value for the next MTU was higher than the current MTU during the last measurement period.		
PMTU invalid protocol received	Indicates the number of "ICMP Fragmentation Needed" error messages received for the packets containing protocols other than the TCP and UDP protocols during the last measurement period.	Number	
PMTU IP check sum error	Indicates the number of "ICMP Fragmentation Needed" error messages received for the packets containing IP checksum errors during the last measurement period.	Number	
PMTU PCB with no link	Indicates the number of "ICMP Fragmentation Needed" error messages received on a Protocol Control Block (PCB) with no link during the last measurement period.	Number	The PCB maintains the state of the connection.
PMTU discovery not enabled	Indicates the number of "ICMP Fragmentation Needed" error messages received when the Path MTU Discovery was not enabled during the last measurement period.	Number	

3.3.4 NetScaler TCP Test

This test monitors the TCP connections to the NetScaler appliance and reports the count of connections that are in various states - i.e., open, closed, opening, closing, etc. In the process, the test holds a mirror to the TCP packet load on the appliance and helps administrators understand the nature of the TCP traffic (SYN, FIN, TIMED_WAIT). In addition, the test also periodically tracks the growth of the Surge Queue of the NetScaler device, and proactively alerts administrators to processing bottlenecks on servers managed by the device.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total client	Indicates the number of	Number	These measures include connections

Measurement	Description	Measurement Unit	Interpretation
connections	client connections on this NetScaler device.		in the Opening, Established, and Closing states.
Total server connections	Indicates the total number of server connections on this NetScaler device.	Number	
Opening client connections	Indicates the number of client connections that are currently in the Opening state.	Number	A connection is said to be OPENING if the handshakes between the server and the client are not completed.
Opening server connections	Indicates the number of server connections that are currently in the Opening state.	Number	
Established client connections:	Indicates the number of client connections that are currently in the *ESTABLISHED* state.	Number	An established connection indicates that the data transfer can occur between the NetScaler device and the client/server.
Established server connections	Indicates the number of server connections that are currently in the *ESTABLISHED* state.	Number	
Closing client connections	Indicates the number of client connections that are currently in the CLOSING state.	Number	A connection is said to be in the CLOSING state when the connection termination process has been initiated but not completed.
Closing server connections	Indicates the number of server connections that are currently in the CLOSING state.	Number	
Opened client connections	Indicates the number of client connections that were initiated on this NetScaler device since startup.	Number	
Opened server connections	Indicates the number of server connections that	Number	

Measurement	Description	Measurement Unit	Interpretation
	were initiated on this NetScaler device since startup.		
Percent of established client connections	Indicates the percentage of client connections in *ESTABLISHED* state.	Percent	
Percent of established server connections	Indicates the percentage of server connections in *ESTABLISHED* state.	Percent	
Percent of closing client connections	Indicates the percentage of client connections in *CLOSING* state.	Percent	
Percent of closing server connections	Indicates the percentage of server connections in *CLOSING* state.	Percent	
Surge queue connections	Indicates the number of connections in the surge queue of this NetScaler device.	Number	The NetScaler device can be used to limit the number of simultaneous requests that are passed on to a server. When a request is completed, additional requests are forwarded to the server. If a request arrives and the server is handling the maximum configured number of requests, the NetScaler device places the new request in a surge queue, where the request waits for its turn to be sent to the server for processing. The surge queue allows a server to run at peak capacity without the risk of having it spiral out of control because of a surge of incoming requests. The surge queue length indicates whether a server is able to keep up with its incoming workload or not. If the surge queue is consistently greater than 0, this indicates that the server is not able to keep up with the workload and

Measurement	Description	Measurement Unit	Interpretation
			additional server capacity is required. On the other hand, a periodic surge is not a cause for concern. When a surge in client requests overloads a server, server response becomes slow, and the server is unable to respond to new requests. The Surge Protection feature ensures that connections to the server occur at a rate that the server can handle. The response rate depends on how surge protection is configured. The NetScaler appliance also tracks the number of connections to the server, and uses that information to adjust the rate at which it opens new server connections.
Spare connections:	Indicates the number of spare connections ready to be used in this NetScaler device.	Number	The NetScaler does connection multiplexing between clients and physical servers. When it receives a client request to access a service on a server, the NetScaler looks for an already established connection to the server that is free. If it finds a free/spare connection, it uses that connection to establish a virtual link between the client and the server. To save time and resources in establishing another connection for a new client, the connection on the server is not closed after completing the request from the first client and is available for serving future requests.
Server active connections	Indicates the number of TCP connections to the server that are currently serving requests.	Number	

Measurement	Description	Measurement Unit	Interpretation
Client idle flushed	Indicates the number of client connections that were flushed during the last measurement period.	Number	The client connection would be flushed when the client has remained idle for a specified time.
Server idle flushed:	Indicates the number of server connections that were flushed during the last measurement period.	Number	The server connections would be flushed when there are no client requests in the queue for a specified time.
Client half opened flushed	Indicates the number of half-open client connections that were flushed during the last measurement period.	Number	A half-opened connection often refers to the TCP connection that is in the process of being established. These connections are flushed when the three-way handshake (SYN, SYN/ACK and ACK) process is not completed. The TCP protocol has a three-way handshake process for opening a connection. First, the originating endpoint (A) sends a SYN packet to the destination (B). A is now in an embryonic state (specifically, SYN_ SENT), and awaiting a response. B now updates its kernel information to indicate the incoming connection from A, and sends out a request to open a channel back (the SYN/ACK packet). At this point, B is also in an embryonic state (specifically, SYN_RCVD). Note that B was put into this state by another machine, outside of B's control. Under normal circumstances (see denial-of-service attack for deliberate failure cases), A will receive the SYN/ACK from B, update its tables (which now have enough information for A to both send and receive), and

Measurement	Description	Measurement Unit	Interpretation
			send a final ACK back to B.
Server half opened flushed	Indicates the number of half-open server connections that were flushed during the last measurement period.	Number	Once B receives this final ACK, it also has sufficient information for two-way communication, and the connection is fully open. Both endpoints are now in an established state.
Client active half closed flushed	Indicates the number of half-closed client connections that were flushed during the last measurement period.	Number	A half-closed connection refers to the connections closed by the client/server and there is no activity taking place on the connection. A half-closed connection may also be referred to as the connection through which the client/server would have stopped sending data but still data is received through the same.
Server active half closed flushed	Indicates the number of half-closed server connections that were flushed during the last measurement period.	Number	
Client passive half closed flushed	Indicates the number of passive half-closed client connections that were flushed during the last measurement period.	Number	A passive half-closed connection refers to the connections closed by the NetScaler and there is no activity taking place on the connection.
Server passive half closed flushed	Indicates the number of passive half-closed server connections that were flushed during the last measurement period.	Number	
Zombie cleanup calls	Indicates the number of times the zombie cleanup function was called during the last measurement period.	Number	Every time a connection is flushed, it is marked for cleanup. The zombie cleanup function clears all these connections at predefined intervals.
Data received	Indicates the amount of TCP data received during the last measurement period.	MB	These are good indicators of the load on the NetScaler appliance.
Data transmitted	Indicates the amount of	MB	

Measurement	Description	Measurement Unit	Interpretation
	TCP data transmitted during the last measurement period.		
Packets received:	Indicates the number of TCP packets received during the last measurement period.	Number	
Packets transmitted	Indicates the number of TCP packets transmitted during the last measurement period.	Number	
SYN packets received	Indicates the number of SYN packets received during the last measurement period.	Number	
Server probes	Indicates the number of probes from this NetScaler device to the server during the last measurement period.	Number	The NetScaler sends a SYN packet to the server to check its availability and expects a SYN_ACK packet from the server before a specified response timeout.
FIN packets from client	Indicates the number of FIN packets received from the clients during the last measurement period.	Number	
FIN packets from server	Indicates the number of FIN packets received from the server during the last measurement period.	Number	
SYN packets received in time wait state	Indicates the number of SYN packets received on connections that are in the TIME_WAIT state during the last measurement period.	Number	Packets cannot be transferred on a connection in this state.
Data received in time wait state	Indicates the amount of data received on	MB	

Measurement	Description	Measurement Unit	Interpretation
	connections that are in the TIME_WAIT state during the last measurement period.		
SYN packets held	Indicates the number of SYN packets held on this NetScaler device during the last measurement period.	Number	The SYN packets would be held when the NetScaler device is waiting for the server connection.
SYN packets flushed	Indicates the number of SYN packets flushed on this NetScaler device during the last measurement period.	Number	The SYN packets would be flushed when there is no response from the server for 3 or more seconds.
Time wait connections closed	Indicates the number of connections that were closed on this NetScaler device because the number of connections in the TIME_WAIT state exceeded the default value of 7000 during the last measurement period.	Number	

3.3.5 TCP Errors/Retransmits Test

Due to network congestion, traffic load balancing, or other unpredictable network behavior, IP packets can be lost. TCP detects these problems, and requests retransmission of lost data. Most often TCP retransmissions have a significant impact on application performance, and will hence have to be kept at a minimum. Using this test, you can determine how often packets sent to/sent by the NetScaler were retransmitted, and can promptly detect the following:

- What type of packets (Client, Server, SYN, FIN, etc.) were retransmitted the most?
- What is causing the retransmissions - is it a bad network link between the NetScaler appliance and the virtual server? is it a poor network connection between the client and the NetScaler appliance? or is it due to an improperly set timeout value for TCP connections?

- Were the retransmissions successful?
- Is any packet received with TCP checksum error?

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Client retransmissions	Indicates the number of packets that were retransmitted by a client during the last measurement period.	Number	The packets are retransmitted when the acknowledgment from the NetScaler has not reached the client.
Server retransmissions	Indicates the number of packets that were	Number	The packets are retransmitted when the acknowledgment from the

Measurement	Description	Measurement Unit	Interpretation
	retransmitted by a server during the last measurement period.		NetScaler has not reached the server.
Full packet retransmissions	Indicates the number of full packets (i.e., the packets as it was originally transmitted) retansmitted by the client or the server during the last measurement period.	Number	
Partial packet retransmissions	Indicates the number of partial packets that were retransmitted by a client or server during the last measurement period.	Number	Some packets may be lost/dropped during transmission due to network congestion or due to a possible connection failure. When retransmission takes place, the remaining packets will alone be sent by the client/server and such remaining packets are termed as partial packets.
SYN packet retries	Indicates the number of SYN (synchronize) packets resent to a server during the last measurement period.	Number	
FIN packet retries	Indicates the number of FIN packets resent to a server or client during the last measurement period.	Number	
SYN packets timeout	Indicates the number of SYN packets that were not retransmitted during the last measurement period.	Number	The SYN packets may not be retransmitted due to the timeout that occurred while establishing a connection on the NetScaler.
FIN packets timeout	Indicates the number of FIN packets that were not retransmitted even after four attempts during the last measurement period.	Number	The FIN packets may not be retransmitted due to a connection timeout that may have occurred because of not receiving the ACK packet after retransmitting the FIN packet four times.

Measurement	Description	Measurement Unit	Interpretation
TCP retransmissions	Indicates the number of TCP packets retransmitted during the last measurement period.	Number	Ideally, the value of this measure has to be low.
TCP retransmission giveup	Indicates the number of times the NetScaler terminates a connection due to non-retransmission of the packets even after seven attempts on that connection during the last measurement period.	Number	If the value of this measure is high, you may want to check what is causing the repeated transmission failures.
Fast retransmits	Indicates the number of TCP packets on which the NetScaler performs a fast retransmission in response to three duplicate acknowledgements or a partial acknowledgement during the last measurement period.	Number	Fast retransmission occurs because the NetScaler assumes that the packet is lost and retransmits the packet before its time-out.
TCP level client header insertion failure	Indicates the number of times the TCP level client header insertion failed during the last measurement period.	Number	
First retransmissions	Indicates the number of packets that were retransmitted in the first attempt by the NetScaler during the last measurement period.	Number	If a large number of packets take too long to be successfully retransmitted, you may have to figure out what is causing the repeated retransmission failures and fix it before more packet loss occurs.
Second retransmissions	Indicates the number of packets that were retransmitted in the second attempt by the NetScaler during the last	Number	

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		
Third retransmissions	Indicates the number of packets that were retransmitted in the third attempt by the NetScaler during the last measurement period.	Number	
Fourth retransmissions	Indicates the number of packets that were retransmitted in the fourth attempt by the NetScaler during the last measurement period.	Number	
Fifth retransmissions	Indicates the number of packets that were retransmitted in the fifth attempt by the NetScaler during the last measurement period.	Number	
Sixth retransmissions	Indicates the number of packets that were retransmitted in the sixth attempt by the NetScaler during the last measurement period.	Number	
Seventh retransmissions	Indicates the number of packets that were retransmitted in the seventh attempt by the NetScaler during the last measurement period.	Number	
Bad TCP checksum	Indicates the number of packets that are received with TCP checksum errors.	Number	
Data after FIN	Indicates the number of bytes received following a connection termination	Number	This error is usually caused by reordering packets during transmission.

Measurement	Description	Measurement Unit	Interpretation
	request.		
SYN in SYN_RCVD state	Indicates the number of SYN packets received on a connection that is in the SYN_RCVD state.	Number	A connection goes into the SYN_ RCVD state after receiving a SYN packet.
SYN in established state	Indicates the number of SYN packets received on a connection that is in the ESTABLISHED state.	Number	A SYN packet is not expected on an ESTABLISHED connection
SYN_SENT incorrect ACK packets	Indicates the number of incorrect ACK packets received on a connection that is in the SYN_SENT state.	Number	An incorrect ACK packet is the third packet in the three-way handshake that has an incorrect sequence number.
Reset packets received	Indicates the number of reset packets received from a client or a server.	Number	
Reset on not established	Indicates the number of reset packets received on a connection that is not in the ESTABLISHED state.	Number	
Reset out of window	Indicates the number of reset packets received on a connection that is out of the current TCP Window	Number	
Reset in time waits	Indicates the number of reset packets received on a connection that is in the TIME_WAIT state.	Number	Typically, packets cannot be transferred on a connection in the TIME_WAIT state. Therefore, the value of this measure is desired to be very low.
Client out of order packets	Indicates the number of out of order TCP packets received from a client.	Number	
Server out of order packets	Indicates the number of out of order TCP packets received from a server.	Number	

Measurement	Description	Measurement Unit	Interpretation
TCP holes on client connection	Indicates the number of TCP holes created on a client connection.	Number	TCP holes are created on the NetScaler appliance for each group of missing packets when out of order packets are received from a client.
TCP holes on server connection	Indicates the number of TCP holes created on a server connection.	Number	TCP holes are created on the NetScaler appliance for each group of missing packets when out of order packets are received from a server.
Sequence number SYN cookie rejects	Indicates the number of SYN cookie packets rejected due to an incorrect sequence number.	Number	
Signature SYN cookie rejects	Indicates the number of SYN cookie packets rejected due to an incorrect signature.	Number	
Sequence number SYN cookie drops	Indicates the number of SYN cookie packets dropped because the sequence number specified in the packets is outside the current Window.	Number	
MSS SYN cookie rejects	Indicates the number of SYN cookie packets rejected due to the incorrect maximum segment size (MSS) specified in the packets.	Number	
Any IP port allocation failures	Indicates the number of port allocations failed on a mapped IP address.	Number	These failures occur when the maximum limit of 65536 has exceeded or the mapped IP address is not configured.
IP port allocation failures	Indicates the number of port allocations that failed on a subnet IP address or virtual server IP address.	Number	These kind of failures occur when the maximum limit of 65536 has exceeded.

Measurement	Description	Measurement Unit	Interpretation
Stray packets	Indicates the number of packets received on a connection whose state is not maintained on the NetScaler appliance.	Number	
Reset packets sent	Indicates the number of reset packets that are sent to a client or a server.	Number	
Bad state connections	Indicates the number of connections that are not in a valid TCP state.	Number	
Reset threshold dropped	Indicates the number of reset packets dropped due to the violation in default threshold.	Number	If the value of this measure increases gradually/suddenly, administrators should reconfigure the default threshold value using the set rate Control command.
Packets out of window	Indicates the number of packets received which are out of the current advertised Window.	Number	
SYNs dropped	Indicates the number of SYN packets dropped due to network congestion	Number	

3.3.6 NetScaler UDP Test

This test monitors the UDP packets/data flowing into and out of the NetScaler appliance, and thus reveals the current UDP load on the appliance. In the process, the test promptly detects UDP packet limit violations and alerts administrators to it, so that potential overload conditions or probable network congestions due to high UDP traffic can be detected and averted.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current rate threshold	Indicates the limit of UDP packets handled every 10 milliseconds.	Number	This implies that within 10 milliseconds (ms) range, NetScaler can allow (receive or pass through) the set number of UDP packets. This threshold is configurable. Once this threshold is violated, subsequent UDP packets will be dropped by the appliance. You are hence advised to configure this limit based on the current and anticipated UDP traffic to the NetScaler appliance. To set/alter this limit, use the set *rateControl* command.
Data received	Indicates the amount of UDP data received during the last measurement period.	MB	These measures are good indicators of the UDP load on the NetScaler appliance.

Measurement	Description	Measurement Unit	Interpretation
Data transmitted	Indicates the amount of UDP data transmitted during the last measurement period.	MB	
Packets received	Indicates the number of UDP packets received during the last measurement period.	Number	
Packets transmitted	Indicates the number of UDP packets sent during the last measurement period.	Number	
Unknown service	Indicates the number of UDP packets that were received on a NetScaler port number that is not assigned to any service.	Number	This measure counts those UDP packets that were received, but dropped.
Bad UDP checksum	Indicates the number of packets that were received with a checksum error.	Number	
Rate threshold exceeded	Indicates the number of times the value reported by the Current rate threshold measure has been violated.	Number	A high value of this measure indicates that the **Current rate threshold** has been violated often. When this happens, you must first ensure that the UDP packets received are genuine. If they are genuine, then you must increase the current rate threshold. To set/alter this limit, use the set *rateControl* command.

3.3.7 DNS Test

You can configure the Citrix NetScaler appliance to function as an authoritative domain name server (ADNS server) for a domain. You can add the DNS resource records that belong to the domain for which the appliance is authoritative and configure resource record parameters. You can also configure the NetScaler appliance as a proxy DNS server that load balances a farm of DNS name servers that are either within your network or outside your network. You can configure the appliance as an end resolver and forwarder. You can configure DNS suffixes that enable name resolution when fully qualified domain names are not configured. The appliance also supports the DNS ANY query that retrieves all the records that belong to a domain.

Using the **DNS** test, you can monitor the DNS queries to the NetScaler appliance, and evaluate how efficiently the appliance handles these queries. DNS requests that were refused and invalid responses that were sent can thus be promptly detected, and their reasons investigated.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
DNS queries received	Indicates the number of DNS queries received during the last measurement period.	Number	
Authoritatively answered queries	Indicates the number of queries that were authoritatively answered during the last measurement period.	Number	An ADNS (Authoritative DNS) server is a DNS server that contains complete information about a zone. To configure the NetScaler as an ADNS server for a zone, you must add an ADNS service, and then configure the zone. To do so, you add valid SOA (Start of Authority) and NS records for the domain. When a client sends a DNS request, the NetScaler appliance searches the configured resource records for the domain name. You can delegate a subdomain, by adding NS records for the subdomain to the zone of the parent domain. You can then make the NetScaler authoritative for the subdomain, by adding a "glue record" for each of the subdomain name servers. If GSLB is configured, the NetScaler makes a GSLB load balancing decision based on its configuration and replies with the IP address of the selected virtual server.
Multi query requests received	Indicates the number of multi query requests received during the last measurement period.	Number	
Server queries sent	Indicates the number of server queries sent during the last measurement period.	Number	
DNS responses	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
received	DNS responses received during the last measurement period.		
Cache flushed	Indicates the number of times the cache was flushed during the last measurement period.	Number	The NetScaler can cache DNS responses (records) and can function as a DNS proxy. This enables the NetScaler to provide quick responses for repeated translations. To configure the NetScaler as a DNS proxy, you must enable caching of DNS records. You must also create a load balancing DNS virtual server, and DNS services, and then bind these services to the virtual server. The NetScaler provides two options, minimum time to live (TTL) and maximum TTL for configuring the lifetime of the cached data. The cached data times out as specified by your settings for these two options. The NetScaler checks the TTL of the DNS record coming from the server. If the TTL is less than the configured minimum TTL, it is replaced with the configured minimum TTL. If the TTL is greater than the configured maximum TTL, it is replaced with the configured maximum TTL. The NetScaler discards (flushes) a record stored in its cache when the time-to-live (TTL) value of the record reaches the configured value.
Server responses received	Indicates the number of server responses received during the last measurement period.	Number	
Cache entries flushed	Indicates the number of cache entries that were	Number	

Measurement	Description	Measurement Unit	Interpretation
	flushed during the last measurement period.		
Updated records	Indicates the number of A records that were updated during the last measurement period.	Number	You can add DNS records on the NetScaler, including address (A) records. Address (A) records are DNS records that map a domain name to an IPv4 address.
Non-existent domain queries:	Indicates the number of queries for which the records were not found in the domain during the last measurement period.	Number	If information pertaining to a requested domain does not exist, it indicates a negative response. This measure therefore reveals the count of negative responses.
Response type unsupported	Indicates the number of responses for which the requested response type was not supported during the last measurement period.	Number	Ideally, the value of this measure should be low.
Response class unsupported	Indicates the number of responses for which the response types were not supported during the last measurement period.	Number	
Query class unsupported	Indicates the number of queries for which the base query class was not supported during the last measurement period.	Number	
Invalid query format	Indicates the number of queries received with an invalid format during the last measurement period.	Number	
Invalid response format	Indicates the number of responses received with a format error during the last measurement period.	Number	Ideally, the value of this measure should be 0.

Measurement	Description	Measurement Unit	Interpretation
Stray answers	Indicates the number of stray answers received during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Responses received without answer	Indicates the number of DNS responses received without an answer during the last measurement period.	Number	Responses received without an answer are deemed as negative responses. Ideally, the value of this measure should be 0.
Responses received with invalid resource data length	Indicates the number of DNS responses received with an invalid resource data length during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Multi queries disabled	Indicates the number of multi queries that were disabled during the last measurement period.	Number	
DNS requests refused	Indicates the number of DNS requests that were refused during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Other errors	Indicates the miscellaneous errors detected during the last measurement period.	Number	Ideally, the value of this measure should be 0.

3.3.8 SSL Test

A Citrix NetScaler appliance configured for SSL acceleration transparently accelerates SSL transactions by offloading SSL processing from the server. To configure SSL offloading, you configure a virtual server to intercept and process SSL transactions, and send the decrypted traffic to the server (unless you configure end-to-end encryption, in which case the traffic is re-encrypted). Upon receiving the response from the server, the appliance completes the secure transaction with the client. From the client's perspective, the transaction seems to be directly with the server. A

NetScaler configured for SSL acceleration also performs other configured functions, such as load balancing.

The **SSL** test reveals how efficiently the NetScaler performs SSL acceleration. The metrics reported by this test provide administrators with indepth insights into the SSL session load on the appliance and the nature of SSL transactions (eg., SSLv1, SSLv2, TLSv1, etc.) that were performed during these sessions, and promptly alerts them to to issues affecting SSL acceleration such as a high number of session reuse missies and failures in multiplexing.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each authentication virtual server configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
SSL cards present	Indicates the number of SSL crypto cards currently present in this NetScaler device.	Number	A server accelerator card (also known as an SSL card) is a Peripheral Component Interconnect (PCI) card used to generate encryption keys for secure transactions on e-commerce Web sites. When a secure transaction is initiated, the Web site's server sends its certificate, which has been provided by a certifying authority, to the client machine to verify the Web site's authenticity. After this exchange, a secret key is used to encrypt all data transferred between sender and receiver so that all personal and credit card information is protected. This process can severely overload a server resulting in fewer transactions processed per second, which means fewer sales. The server accelerator card takes over this process, thus reducing the load on the server. Server accelerator cards support a number of security protocols including Secure Sockets Layer (SSL) and Secure Electronic Transaction (SET).
SSL cards up	Indicates the number of SSL cards that are currently UP in this NetScaler device.	Number	A low value for this measure indicates that many SSL cards are currently Down.
SSL engine status	Indicates the current status of the SSL engine.		The values reported by this measure and their numeric equivalents are as shown in the table:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Down</td></tr>
<tr><td>1</td><td>Up</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating the status of the SSL engine. However, in the graph of this measure, the states will be represented using the corresponding numeric equivalents - i.e., 0 or 1.
</td></tr>
<tr><td>SSL sessions</td><td>Indicates the number of current SSL sessions on this NetScaler device.</td><td>Number</td><td>This measure is a good indicator of the current SSL session load on the appliance.</td></tr>
<tr><td>SSL transactions</td><td>Indicates the number of SSL transactions performed on this NetScaler device during the last measurement period.</td><td>Number</td><td>For an SSL transaction to be initiated, and for successful completion of the SSL handshake, the server and the client should agree on an SSL protocol that both of them support. If the SSL protocol version supported by the client is not acceptable to the server, the server does not go ahead with the transaction, and an error message is displayed.</td></tr>
<tr><td>SSLv2 transactions</td><td>Indicates the number of SSLv2 transactions performed on this NetScaler device during the last measurement period.</td><td>Number</td><td></td></tr>
<tr><td>SSLv3 transactions</td><td>Indicates the number of SSLv3 transactions performed on this NetScaler device during the last measurement</td><td>Number</td><td></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	period.		
TLSv1 transactions	Indicates the number of TLSv1 transactions on this NetScaler device during the last measurement period.	Number	
Front-End SSL sessions	Indicates the number of Front-end SSL sessions on this NetScaler device during the last measurement period.	Number	In certain deployments, you might be concerned about network vulnerabilities between the NetScaler appliance and the backend servers, or you might need complete end-to-end security and interaction with certain devices that can communicate only in clear text (for example, caching devices). In such cases, you can set up an HTTP virtual server that receives data from clients that connect to it at the front end and hands the data off to a secure service, which securely transfers the data to the web server. To implement this type of configuration, you configure an HTTP virtual server on the NetScaler and bind SSL based services to the virtual server. The NetScaler receives HTTP requests from the client on the configured HTTP virtual server, encrypts the data, and sends the encrypted data to the web servers in a secure SSL session. This measure reports of the count of those SSL sessions that are front-ended by such virtual servers.
Front-End SSLv2 sessions	Indicates the number of Front-end SSLv2 sessions on this NetScaler device during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
Front-End SSL v3 sessions	Indicates the number of Front-end SSLv3 sessions on this NetScaler device during the last measurement period.	Number	
Front-End TLSv1 sessions	Indicates the number of TLSv1 sessions on this NetScaler device during the last measurement period.	Number	
Front-End new sessions	Indicates the number of new Front-end SSL sessions on this NetScaler device during the last measurement period.	Number	
Front-End SSL session reuse misses	Indicates the number of SSL session reuse misses on the NetScaler appliance since the last measurement period.	Number	For SSL transactions, establishing the initial SSL handshake requires CPU-intensive public key encryption operations. Most handshake operations are associated with the exchange of the SSL session key (client key exchange message). When a client session is idle for some time and is then resumed, the SSL handshake is typically conducted all over again. With session reuse enabled, session key exchange is avoided for session resumption requests received from the client. Session reuse is enabled on the NetScaler appliance by default. Enabling this feature reduces server load, improves response time, and increases the number of SSL transactions per second (TPS) that can be supported by the server. A server therefore, is said to be performing at peak capacity if the

Measurement	Description	Measurement Unit	Interpretation
Front-End SSL session reuse hits	Indicates the number of SSL session reuse hits on the NetScaler appliance since the last measurement period.	Number	value of the Front-End SSL session reuse misses measure is low and the value of the Front-End SSL session reuse hits measure is high.
Front-End SSLv1 client authentications	Indicates the number of client authentications performed through the Front-end SSLv2 transactions on this NetScaler device during the last measurement period.	Number	
Front-End SSLv3 client authentications	Indicates the number of client authentications performed through the Front-end SSLv3 trensactions on this NetScaler device during the last measurement period.	Number	
Front-End TLSv1 client authentications	Indicates the number of client authentications performed through the Front-end TLSv1 transactions on this NetScaler device during the last measurement period.	Number	
Back-End SSL sessions	Indicates the number of Back-end SSL sessions through which transactions were performed on the virtual server by this NetScaler device during the last measurement period.	Number	In certain deployments, you might be concerned about network vulnerabilities between the NetScaler appliance and the backend servers, or you might need complete end-to-end security and interaction with certain devices that can communicate only in clear text (for example, caching

Measurement	Description	Measurement Unit	Interpretation
			devices). In such cases, you can set up an HTTP virtual server that receives data from clients that connect to it at the front end and hands the data off to a secure service, which securely transfers the data to the web server. To implement this type of configuration, you configure an HTTP virtual server on the NetScaler and bind SSL based services to the virtual server. The NetScaler receives HTTP requests from the client on the configured HTTP virtual server, encrypts the data, and sends the encrypted data to the web servers in a secure SSL session. This measure reports the count of those SSL sessions between the front-end HTTP virtual server and the backend web servers.
Back-End SSLv3 sessions	Indicates the number of Back-end SSLv3 sessions through which transactions were performed on the virtual server by this NetScaler device during the last measurement period.	Number	
Back-End TLSv1 sessions	Indicates the number of Back-end TLSv1 sessions through which transactions were performed on the virtual server by this NetScaler device during the last measurement period.	Number	
Back-End SSL sessions multiplex	Indicates the number of Back-end SSL session	Number	You can configure the back-end SSL

Measurement	Description	Measurement Unit	Interpretation
attempts	multiplexing attempts made by this NetScaler device to access the virtual servers during the last measurement period.		transactions so that the NetScaler appliance uses SSL session multiplexing to reuse existing SSL sessions with the back-end web servers, thus avoiding CPU-intensive key exchange (full handshake) operations. This reduces the overall number of SSL sessions on the server, and therefore accelerates the SSL transaction while maintaining end-to-end security. This is why, a large number of Backend SSL sessions multiplex attempts successes is desired. On the other hand, too many Backend SSL sessions multiplex attempts failures could imply that SSL sessions could not be reused. This in turn can result in increased full handshakes, probable session overloads on the backend web servers, and consequently, slower SSL transaction processing.
Back-End SSL sessions multiplex attempts successes	Indicates the number of Back-end SSL session multiplexing attempts that were successfully made by this NetScaler device during the last measurement period.	Number	
Back-End SSL sessions multiplex attempts failures:	Indicates the number of failed Back-end SSL session multiplexing attempts made by this NetScaler device during the last measurement period.	Number	
Back-End SSLv3 client authentications	Indicates the number of client authentications performed by the virtual server through SSLv3 sessions during the last measurement period.	Number	
Back-End TLSv1 client authentications	Indicates the number of client authentications performed by the virtual server through TLSv1 sessions during the last measurement period.	Number	
Data decrypted	Indicates the amount of	MB	

Measurement	Description	Measurement Unit	Interpretation
	data decypted on this NetScaler device during the last measurement period.		
Data encrypted	Indicates the amount of data encrypted on this NetScaler device during the last measurement period.	MB	

3.4 The NetScaler Gateway Layer

Using the tests mapped to this layer, administrators can figure out:

- Bottlenecks in the transfer of log information from the NetScaler appliance to the SYSLOG server;
- Errors in VPN sessions
- Overloaded VPN virtual servers
- Processing slowdowns on the authentication virtual server
- Users with the maximum number of open ICA connections
- Validity of SSL Certificates

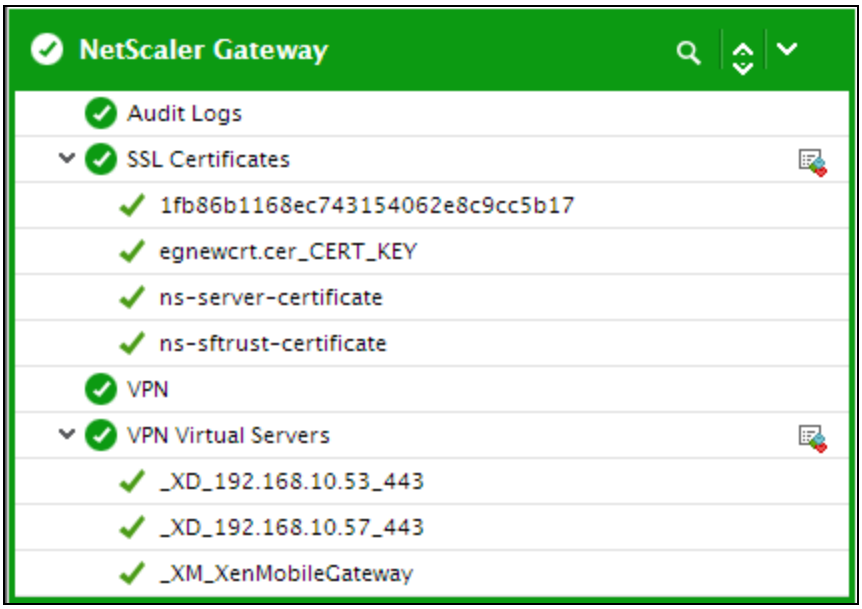

Figure 3.6: The tests mapped to the NetScaler Gateway layer

3.4.1 Audit Logs Test

Auditing is a methodical examination or review of a condition or situation. The Audit Logging feature enables you to log the Citrix NetScaler states and status information collected by various modules in the kernel and in the user-level daemons. For audit logging, you have the options to configure SYSLOG, the native NSLOG protocol, or both.

SYSLOG is a standard protocol for logging. It has two components - the SYSLOG auditing module, which runs on the NetScaler appliance, and the SYSLOG server, which can run on the underlying FreeBSD operating system (OS) of the NetScaler appliance or on a remote system. SYSLOG uses user data protocol (UDP) for the transfer of data.

When you run a SYSLOG server, it connects to the NetScaler appliance. The NetScaler appliance then starts sending all the log information to the SYSLOG server, and the server can filter the log entries before storing them in a log file. A SYSLOG server can receive log information from more than one NetScaler appliance and a NetScaler appliance can send log information to more than one SYSLOG server or NSLOG server.

Using this test, you can monitor the transfer of log information from the NetScaler appliance to the SYSLOG server, so that you can instantly spot bottlenecks in data transfer and identify the probable causes for the same - is it because of NAT/NSB allocation failures? is it because memory allocations of the Access Gateway context structure failed? is it due to too many port allocation failures?

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each load balancing virtual server configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Logs sent to the syslog servers	Indicates the number of Syslog messages sent to the Syslog server during the last measurement period.	Number	
Logs not sent to the syslog servers	Indicates the number of Syslog messages that were not sent to the Syslog server during the last measurement period.	Number	
Log messages generated	Indicates the number of Syslog messages that were about to be sent to	Number	If the value of this measure is a lot higher than the value of the *Logs not sent to the syslog servers* measure,

Measurement	Description	Measurement Unit	Interpretation
	the Syslog server during the last measurement period.		it could indicate bottlenecks in message transmission. Further investigation is hence recommended.
NAT allocation failed	Indicates the number of NAT allocations that failed during the last measurement period.	Number	
NSB allocation failed	Indicates the number of NetScaler Buffer (NSB) allocations that failed during the last measurement period.	Number	
Memory allocation failed	Indicates the failures in memory allocation of the Access Gateway context structure during the last measurement period.	Number	When an Access Gateway session is established, the NetScaler appliance creates an internal context structure, which identifies the user and the IP address from which the user has logged in.
Port allocation failed	Indicates the number of times the NetScaler failed to allocate a port when sending a syslog message to the syslog server during the last measurement period.	Number	These measures serve as effective indicators of data/packet load on a virtual server.
NAT lookup failed	Indicates the number of NAT lookups that failed during the last measurement period.	Number	
Context not found	Indicates the failures in finding the context structure for an Access Gateway session during attempts to send session-specific audit messages during the last	Number	During an Access Gateway session, audit messages related to the session are queued up in the auditlog buffer for transmission to the audit log server(s). If the session is terminated before the messages are sent, the context

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		structure allocated at session creation is removed. This structure is required for sending the queued auditlog messages. If it is not found, then this counter is incremented.
NSB chain allocation failed	Indicates the number of NetScaler Buffer (NSB) chain allocations that failed during the last measurement period.	Number	
Client connect failed	Indicates the number of times the connection between the NetScaler and the auditserver tool (the NetScaler's custom logging tool) failed to establish during the last measurement period.	Number	
Multiprocessor buffer flush command count	Indicates the number of auditlog buffer flushes during the last measurement period.	Number	In a multiprocessor NetScaler appliance, both the main processor and the co-processor can generate auditlog messages and fill up the auditlog buffers. But only the primary processor can free up the buffers by sending auditlog messages to the auditlog server(s). The number of auditlog buffers is fixed. If the co-processor detects that all the auditlog buffers are full, then it issues a flush command to the main processor.

3.4.2 VPN Test

This test monitors VPN sessions and reports errors in the sessions.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Login page requests received	Indicates the number of VPN login requests received during the last meaurement period.	Number	
Login page delivery failures	Indicates the number of times the VPN login page failed to be displayed during the last measuement period.	Number	Ideally, the value should be 0.
Client configuration requests received	Indicates the number of client configuration requests received by the VPN server during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
DNS queries resolved:	Indicates the number of DNS queries resolved by the VPN server during the last measurement period.	Number	
WINS queries resolved	Indicates the number of WINS queries resolved by the VPN server during the last measurement period.	Number	
SSL VPN tunnels	Indicates the number of SSL VPN tunnels formed between the VPN server and the client during the last measurement period.	Number	
Backend HTTP server probes	Indicates the number of probes from the VPN server to the Back-end HTTP servers that have been accessed by the VPN client during the last measurement period.	Number	
Backend Non-HTTP server probes	Indicates the number of probes from the VPN server to the Back-end non-HTTP servers that have been accessed by the VPN client during the last measurement period.	Number	
Backend server probe successes	Indicates the number of probes that were successful on all the Back-end servers during the last measurement period.	Number	
File system requests received:	Indicates the number of file system requests received by the VPN server during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
Number of times MIP used and IIP disabled	Indicates the number of times the Mapped IP (MIP) was used in lieu of the Intranet IP (IIP) that was disabled during the last measurement period.	Number	
Number of times MIP used and IIP failed	Indicates the number of times the MIP was used in lieu of the IIP as the IIP assignment failed during the last measurement period.	Number	
Number of times MIP used and IIP spillover	Indicates the number of times the MIP was used on the IIP spillover during the last measurement period.	Number	When IP pooling is configured and the mapped IP is used as an intranet IP address, the mapped IP address is used when an intranet IP address cannot be assigned. This process is termed as a Spillover.
Both MIP and IIP disabled	Indicates the number of times both the MIP and the IIP was disabled during the last measurement period.	Number	
Number of times IIP failed and MIP disabled:	Indicates the number of times the IIP assignment failed with the MIP was disabled during the last measurement period.	Number	
SOCKS method requests received	Indicates the number of requests received through the SOCKS method during the last measurement period.	Number	SOCKet Secure (SOCKS) is an Internet protocol that routes network packets between a client and server through a proxy server.
SOCKS method requests sent	Indicates the number of requests sent through the SOCKS method during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
SOCKS method responses received	Indicates the number of responses received through the SOCKS during the last measurement period.	Number	
SOCKS method responses sent	Indicates the number of responses sent through the SOCKS during the last measurement period.	Number	
SOCKS connect requests received	Indicates the number of connect requests received through the SOCKS during the last measurement period.	Number	
SOCKS connect requests sent:	Indicates the number of connect requests sent through the SOCKS during the last measurement period.	Number	
SOCKS connect responses received	Indicates the connect responses received through the SOCKS during the last measurement period.	Number	
SOCKS connect responses sent	Indicates the connect responses sent through the SOCKS during the last measurement period.	Number	
SOCKS server errors	Indicates the number of server errors received through SOCKS protocol during the last measurement period.	Number	Ideally, this value should be 0.
SOCKS client errors	Indicates the number of client errors received through SOCKS protocol during the last	Number	

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		
STA connection successes	Indicates the number of successful STA (Secure Ticketing Authority) connections during the last measurement period.	Number	The Secure Ticket Authority (STA) is an XML Web service that exchanges MetaFrame server information for randomly generated tickets. It is used to control access for a Citrix Secure Gateway server.
STA connection failures	Indicates the number of STA connections that failed during the last measurement period.	Number	Ideally, this value should be 0.
CPS connection successes	Indicates the number of successful CPS (Citrix Provisioning Server) connections during the last measurement period.	Number	
CPS connection failures	Indicates the number of successful CPS connections during the last measurement period.	Number	Ideally, this value should be 0.
STA requests sent	Indicates the number of STA requests sent during the last measurement period.	Number	
STA responses received	Indicates the number of STA responses received during the last measurement period.	Number	
ICA license failures	Indicates the number of ICA license failures during the last measurement period.	Number	Ideally, this value should be 0.

3.4.3 VPN Virtual Servers Test

A vserver is a named NetScaler entity that external clients can use to access applications hosted on the servers. It is represented by an alphanumeric name, virtual IP address (VIP), port, and protocol. When a client attempts to access applications on a server, it sends a request to the VIP instead of the IP address of the physical server. When the NetScaler receives a request on the VIP, it terminates the connection at the vserver and uses its own connection with the server on behalf of the client. The port and protocol settings of the vserver determine the applications that the vserver represents.

vservers can be grouped into various categories. One such vserver is the virtual private network (VPN) virtual server. This server decrypts tunneled traffic and sends it to intranet applications. To understand the workload of each of these VPN virtual servers and isolate overloaded servers, use the **VPN Virtual Servers** test.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each VPN virtual server.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Parameter	Description
Show Up Server Oly	The default setting of this flag is **No**; this indicates that this test, by default, monitors all the VPN virtual servers configured on the NetScaler appliance. If you want the test to monitor only those VPN virtual servers that are up and running currently, then set this value to **Yes**.
Exclude Servers	Provide a comma-separated list of VPN virtual server names or name patterns that need to be excluded from monitoring. By default, this is set to *none*, indicating that all VPN virtual servers are by default monitored.
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
Server state	Indicates the current state of this virtual server.		If the virtual server is up, then the value of this measure is Up. If the virtual server is down, then the value of this measure is Down. The numeric values that correspond to these measure values have been listed in the table below: **Numeric Value** / **Measure Value**: 0 – Up; 1 – Down

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>2</td><td>Out of service</td></tr>
<tr><td>3</td><td>Transition out of service</td></tr>
<tr><td>4</td><td>Down when going out of service</td></tr>
<tr><td>-1</td><td>Unknown</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether a virtual server is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.
Use the detailed diagnosis of this measure to determine the primary IP address, primary port, and protocol type of each VPN virtual server being monitored.
</td></tr>
<tr><td>Request data received</td><td>Indicates the amount of request data received by this virtual server during the last measurement period.</td><td>MB</td><td rowspan="3">These are good measures of the request and response load on a virtual server. By comparing the value of each of these measures across virtual servers, you can instantly identify overloaded servers.</td></tr>
<tr><td>Response data received</td><td>Indicates the amount of response data received by this virtual server during the last measurement period.</td><td>MB</td></tr>
<tr><td>Requests received</td><td>Indicates the number of requests received by this virtual server during the last measurement period.</td><td>Number</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Responses received	Indicates the number of responses received by this virtual server during the last measurement period.	Number	
Currently AAA users logged in	Indicates the number of AAA users who are currently logged into this virtual server.	Number	A high value is indicative of high user load on the virtual server. By continuously tracking changes to this measure alongside the value of the Maximum users allowed to login measure, you can figure out when this upper threshold (i.e., Maximum users allowed to login measure) is likely to be reached/crossed.
Maximum users allowed to login	Indicates the number of concurrent users who are allowed to login to this virtual server.	Number	
Total users connected	Indicates the total number of users connected to this virtual server.	Number	This measure is a good indicator of total workload on the virtual server. Compare the value of this measure across the virtual servers to know which virtual server is overloaded.

3.4.4 Virtual Server Authentications Test

The AAA feature supports authentication, authorization, and auditing for all application traffic. To use AAA, you must configure authentication virtual servers to handle the authentication process. This process ensures that the access is granted only to an authorized user of the network.

If an authentication virtual server is rendered unavailable for a while or is unable to process authentication requests owing to an overload condition, unauthorized users may end up gaining access to critical data on the NetScaler. Sometimes, valid users may also be denied access. To avoid this, you can use this test to continuously monitor the state of your authentication virtual servers, track the flow of requests to and responses from each of these servers, and capture abnormalities before users notice anything amiss.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each authentication virtual server configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Up Server Oly	The default setting of this flag is **No**; this indicates that this test, by default, monitors all the VPN virtual servers configured on the NetScaler appliance. If you want the test to monitor only those VPN virtual servers that are up and running currently, then set this value to **Yes**.
Exclude Servers	Provide a comma-separated list of VPN virtual server names or name patterns that need to be excluded from monitoring. By default, this is set to *none*, indicating that all VPN virtual servers are by default monitored.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Server state	Indicates the current state of this authentication virtual server.		The values that this measure reports and their corresponding numeric equivalents have been listed in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Up</td></tr>
<tr><td>1</td><td>Down</td></tr>
<tr><td>2</td><td>Out of service</td></tr>
<tr><td>3</td><td>Transition out of service</td></tr>
<tr><td>4</td><td>Down when going out of service</td></tr>
<tr><td>-1</td><td>Unknown</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values to indicate the current state of a virtual server. However, in the graph of this measure, virtual server states will be represented using the corresponding numeric equivalents only.
The Detailed Diagnosis of this measure shows the Service Type, Primary Port and the Primary IP address of the virtual server.
</td></tr>
<tr><td>Data received</td><td>Indicates the amount of data received by this virtual server during the last measurement period.</td><td>MB</td><td></td></tr>
<tr><td>Data transmitted</td><td>Indicates the amount of data transmitted by this virtual server during the last measurement period.</td><td>MB</td><td></td></tr>
<tr><td>Requests</td><td>Indicates the number of authentication requests received by this virtual server during the last measurement period.</td><td>Number</td><td>A high value of this measure could indicate a potential overload.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Responses	Indicates the number of authentication responses sent out by this virtual server during the last measurement period.	Number	If the number of **Responses** is way too less than the number of **Requests**, it could indicate a processing bottleneck on the authentication virtual server. This could result in genuine users being denied access or gaining delayed access to resources.

3.4.5 Connections Test

When the NetScaler experiences an ICA connection overload, administrators may want to instantly determine which current user is contributing the most to this overload by establishing the maximum number of connections to the NetScaler appliance. The **Connections** test provides administrators with this information. This test auto-discovers the users who are currently connected to the appliance via ICA, and reports the count of ICA connections established per user. A quick comparison of the connection count across users will reveal, which user holds the maximum number of connections.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each user currently connected to the appliance via ICA.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text

Parameter	Description
	boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Secure Gateway DD	
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current connections	Indicates the number of ICA connections currently established by this user with the NetScaler appliance.	Number	In the event of an overload, compare the value of this measure across users to know which user has established the maximum number of connections and is contributing to the overload. You can use the detailed diagnosis of this measure to know from which client and Citrix XenApp server each user connected.
Active user sessions	Indicates the number of	Number	Compare the value of this measure

Measurement	Description	Measurement Unit	Interpretation
	ICA sessions that are currently active for this user on the NetScaler appliance.		across users to know which user has launched the maximum number of ICA sessions on the NetScaler appliance.

The detailed diagnosis of the *Current connections* measure reveals which client and Citrix XenApp server each user connected from.

Shows the list of ICA connections details

TIME	DOMAIN NAME	CLIENT IP ADDRESS	CLIENT PORT	CPS IP ADDRESS	CPS PORT
May 23, 2014 15:44:05					
	citrix	192.168.9.44	51429	192.168.8.126	2598
	citrix	192.168.9.44	63770	192.168.8.126	2598

Figure 3.7: The detailed diagnosis of the Current Connections measure of the NS ICA Connections Test

3.4.6 Top Sources Test

In environments where the NetScaler appliance is deployed, external clients (Source IPs) communicate via the NetScaler appliance to access applications hosted on the servers (Destination IPs). The NetScaler appliance is deployed in front of the web or database servers and routes requests to the servers from the external clients through it. When multiple clients are sending requests to the servers through the NetScaler device, it is impossible for the administrators to keep track of which client is sending the maximum amount of data to the server through its requests. In order to identify the topmost clients who are contributing to the maximum data transfer between the client and the server, adminsitrators can use the **Top Sources** test. This test dynamically lists the top clients and for each client reports the amount of data sent from the client to the server and amount of data received by the client from the server. By comparing the measures reported by this test, administrators can identify the topmost client in terms of data transfer.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the

Component type, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

For this test to run and report metrics, the NetScaler device should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged.

3.4.6.1 Creating a Syslog file in a remote Syslog server

To configure the Syslog server where this Syslog file should be created, do the following:

1. Connect to the NetScaler management console from your browser using the URL: <http://NetScaler host:port>.
2. Login to the NetScaler as an administrator.
3. Figure 3.8 will then appear.

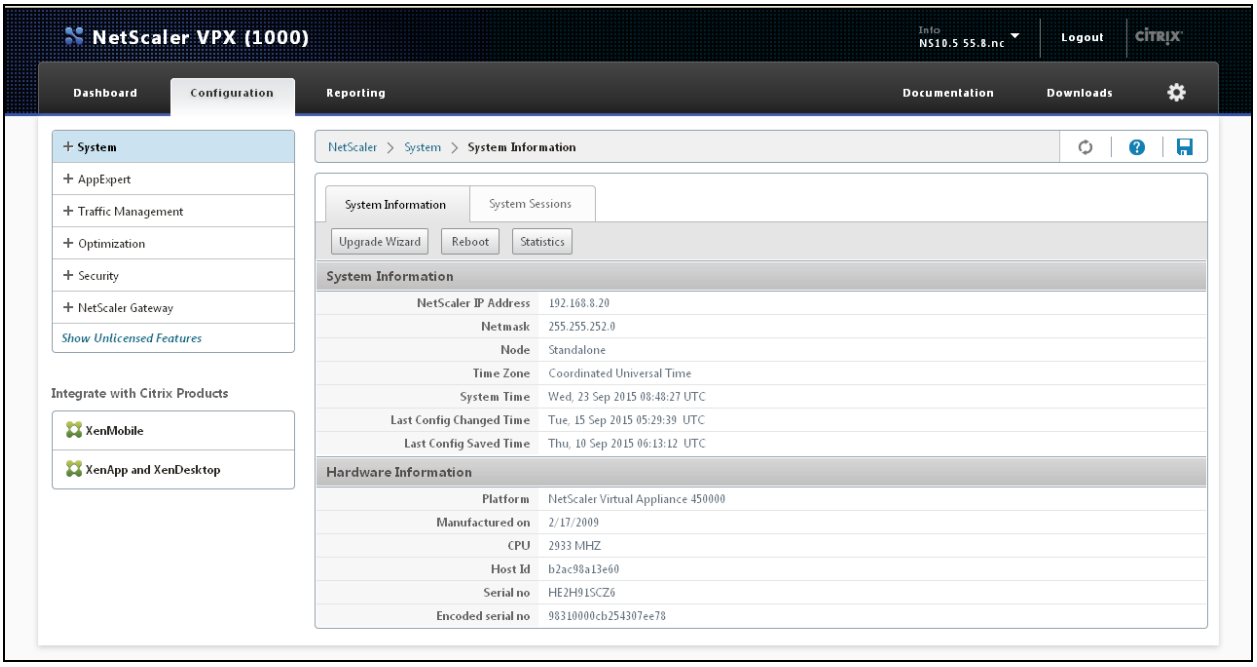

Figure 3.8: The NetScaler management console

4. Expand the **System** node in Figure 3.8 and further expanding the Auditing node will lead you to the **Syslog** option. Then, click the **Syslog** option.

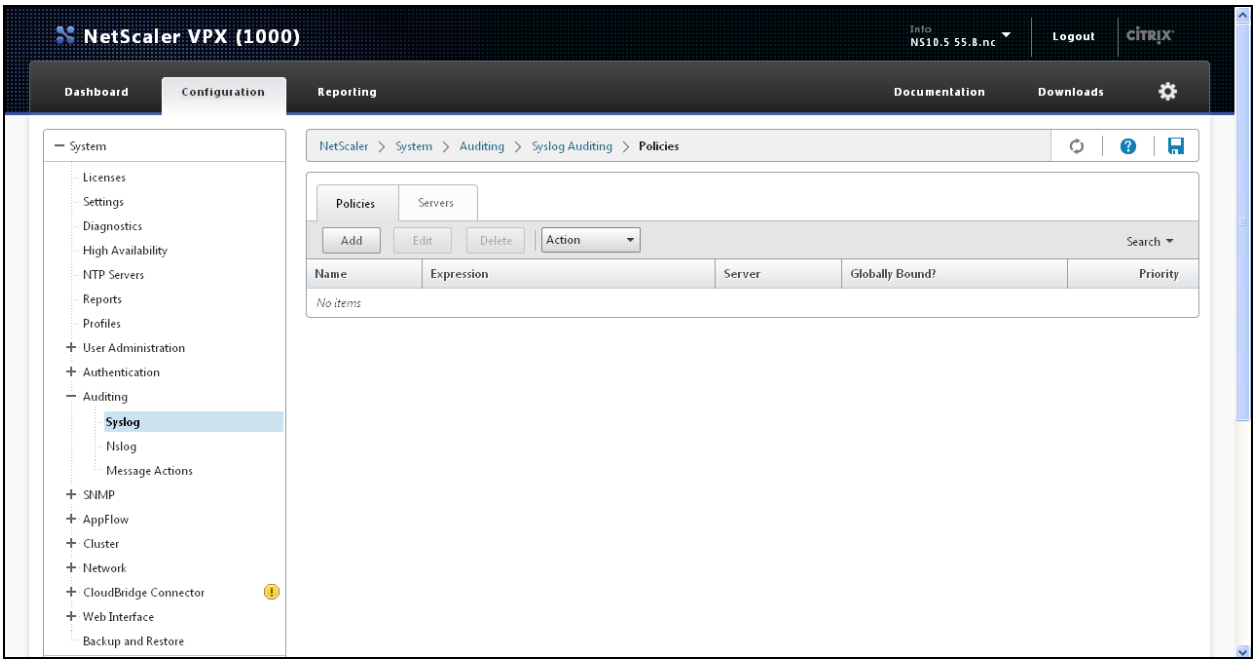

Figure 3.9: Figuring out the SysLog option

5. This will bring up a **Policies** tab and a **Servers** tab in the right panel, where you can configure Policies for a Syslog server, configure a remote Syslog server and enable Syslog file creation on the server. Selecting the Servers tab will lead you to Figure 3.10.

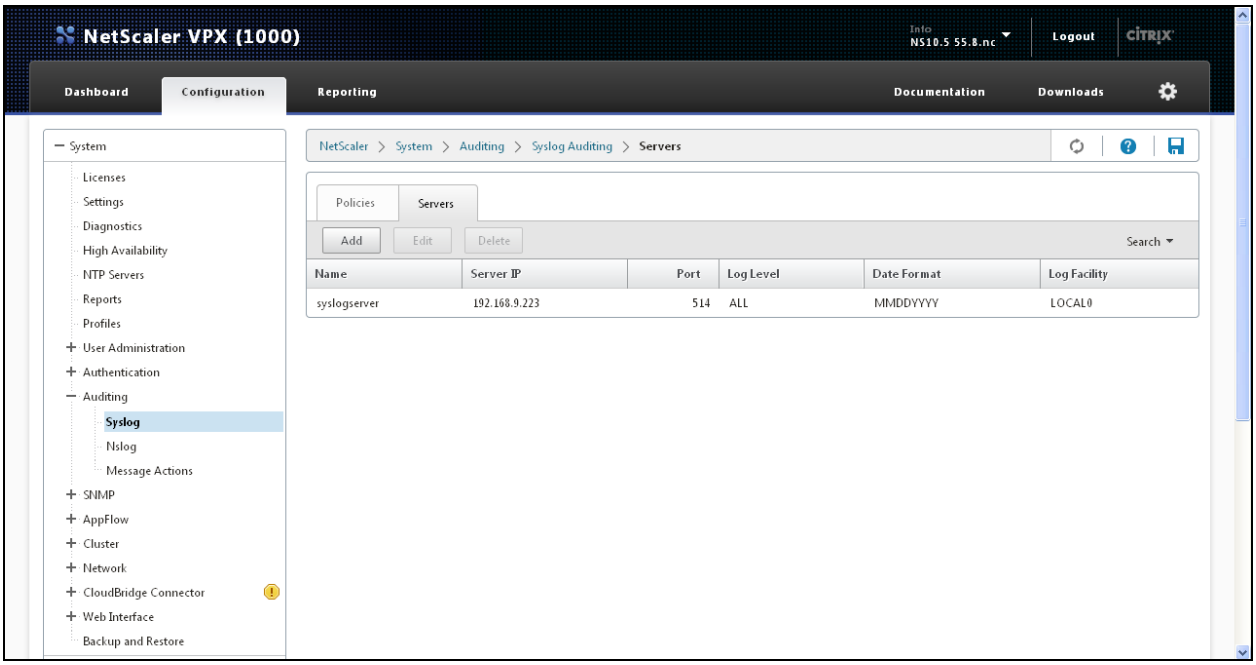

Figure 3.10: Configuring the Syslog server where the Syslog file is to be created

6. To configure a new Syslog server, click the **Add** button in Figure 3.10. Figure 3.11 will then appear.

Figure 3.11: Creating the new Syslog server

7. Enter the **Name** of the Syslog server.
8. Then enter the IP address of the Syslog server in the **IP Address** text box of Figure 3.10.
9. Enter the **Port** at which the Syslog server listens.
10. Then, indicate what details should be logged in the Syslog file. For the eG tests to work, set the **Log Levels** flag to **ALL.**
11. Set the **Time Zone** to **Local** and select the check boxes against **TCP Logging**, **ACL Logging**, **User Configurable Log Messages** and **AppFlow Logging**. Even though the syslog file is populated with the log messages, the metrics will be displayed in the eG Monitor interface only when these checkboxed are checked.
12. Click the **Create** button in Figure 3.10 to configure the Syslog server.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Source IP address configured on the target Citrix NetScaler

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST -** The host for which the test is to be configured.
3. **PORT** – The port at which the **HOST** listens. By default, this is **NULL**.
4. **LOG FILE PATH** – This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
5. **SEARCH STRING** – By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the **SEARCH STRING** text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
6. **SEARCH STRING INDEX** – Here, specify the cursor position after which the eG agent should search for the specified **SEARCH STRING** (or the position upto which the eG agent should ignore while searching for the specified **SEARCH STRING**) in the syslog file. For example, if the specified **SEARCH STRING** appears in the syslog file at the 17th position, then you may need to specify the **SEARCH STRING INDEX** as 16.
7. **REPORT TOP N FLOWS** – Specify the number of descriptors that need to be taken into consideration while monitoring the target NetScaler appliance. By default, this test will report the Top -10 source IP addresses.
8. **DD FREQUENCY** – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.
9. **DETAILED DIAGNOSIS** – To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

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

   - The eG manager license should allow the detailed diagnosis capability

- Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total data sent:	Indicates the total amount of data sent from this client (source IP) to the server (destination IP).	KB	Compare the value of this measure across the source IP addresses to know which source IP address is sending the maximum amount of data. The detailed diagnosis of this measure if enabled, lists the Source IP, Destination IP, Start time, End Time, Total Bytes sent and Total Bytes received.
Total data received:	Indicates the total amount of data received by this client (source IP) from the server (destination IP).	KB	Compare the value of this measure across the source IP addrsesses to know which source IP address is receiving the maximum amount of data. The detailed diagnosis of this measure if enabled, lists the Source IP, Destination IP, Start time, End Time, Total Bytes sent and Total Bytes received.

3.4.7 Top Destinations Test

In environments where the NetScaler appliance is deployed, external clients (Source IPs) communicate via the NetScaler appliance to access applications hosted on the servers (Destination IPs). The NetScaler appliance is deployed in front of the web or database servers and routes requests to the servers from the external clients through it. When multiple servers are responding to the requests received from the clients through the NetScaler device, it is impossible for the administrators to keep track of which server is sending the maximum amount of data to the clients

while responding to the requests. In order to identify the topmost servers that are contributing to the maximum data transfer between the server and the client, adminsitrators can use the **Top Destinations** test. This test dynamically lists the top servers and for each server reports the amount of data sent from the serverto the client and amount of data received by the server from the client. By comparing the measures reported by this test, administrators can identify the topmost server (Destination IP) in terms of data transfer.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Destination IP address configured on the target Citrix NetScaler.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the

Parameter	Description
	specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total data sent	Indicates the total amount of data sent from this destination IP to the source IP.	KB	Compare the value of this measure across the destination IP addresses to know which destination IP address is sending the maximum amount of data. The detailed diagnosis of this measure if enabled, lists the Destination IP, Source IP, start time, End time, total byte sent and the total byte received.
Total data received	Indicates the total amount of data received by this destination IP from the	KB	Compare the value of this measure across the destination IP addresses to know which destination IP address is

Measurement	Description	Measurement Unit	Interpretation
	source IP.		receiving the maximum amount of data. The detailed diagnosis of this measure if enabled, lists the Destination IP, Source IP, start time, End time, total byte sent and the total byte received.

3.4.8 Connection Delinks or Terminates Test

Whenever multiple connections are established to the NetScaler appliance, it is the onus of the administrators to keep track of the connections, figure out the connections that were delinked, connections that were terminated etc. When a TCP connection between the client side and server side delinks after a response is sent from the NetScaler device, that particular server connection can be reused for another client. This helps administrators to effectively utilize the connections to the NetScaler device. For an administrator to keep tab on the efficiency of the NetScaler device based on the connections that are established, the **Connection delinks or Terminates** test is the best bet to rely upon!

This test helps administrators to figure out the number of server and client TCP conenctions that are delinked, the connections that are terminated, the number of TCP connections for RNAT that are delinked etc.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Connection delinks	Indicates the number of times the server side and client side connections are delinked after the response is sent from the NetScaler device.	Number	By default, connections like HTTP are tracked by the NetScaler device. The detailed diagnosis of this measure if enabled, lists the Source IP, vServer, NAT IP, Destination IP, Delink time, Data sent and Data received.
Other connection delinks	Indicates the number of times the server side and client side TCP connections were delinked.	Number	These are connections like FTP, Telnet etc that are not tracked by the NetScaler device.
Connection terminates	Indicates the number of times the TCP connection was terminated.	Number	The detailed diagnosis of this measure if enabled lists the Source IP, Destination IP, the start time, the end time, the data sent and data received.
NAT connection delinks	Indicates the number of times the server side and client side TCP connection for RNAT was delinked.	Number	
NAT other connection delinks	Indicates the number of times the server side and client side TCP connection that are not tracked by the NetScaler device for RNAT was delinked.	Number	These are connections that are not tracked by the NetScaler device.

3.4.9 NetScaler License Information Test

To ensure continuous application/service delivery, the NetScaler device should ideally possess a valid license, which should be renewed in time. Invalid/out-dated lincenses may disable the key features enabled on the NetScaler device or lead the NetScaler to deliver applications with out-dated license information. Sometimes, the invalid licenses may completely halt the operations of the NetScaler. Consequently, the services/applications, delivered by the NetScaler device, become unavailable to users, thus, the user experience will get affected. To avoid this, administrators should

be able to proactively figure out the validity of licenses at the earliest. This is where the **NetScaler License Information** test helps administrators! This test continuously tracks the license usage of NetScaler device, and determines when each license is likely to expire.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each lincense of the target Citrix NetScaler being monitored..

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. However, to collect metrics for this test, the user is required to access the Citrix License file available in the NetScaler device. To access the Citrix License file, the user should be vested with *superuser* permission. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Expired License	By default, the Show Expired License flag is set to **No** indicating that this test will not monitor the licenses of NetScaler device that expired. Set this flag to **Yes**, if the expired licenses are also needed to be included for monitoring. As a result, this test will report -6 against the *Days to expire* measure to indicate the licenses that expired.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *6:1*. 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

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

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Days to expire	Indicates the number of days by which this license will expire.	Number	A very low value for this measure indicates that the license is nearing expiry. You may have to request for a license extension if you want to continue using the NetScaler. If the Show Expired License parameter is set to **Yes**, this test will display -6 against this measure to denote the expired licenses (if any). The detailed diagnosis of this measure lists the NetScaler IP to which each license is mapped to and the feature of the licensed appliance.

3.5 The NetScaler Events Layer

The tests mapped to this layer scan the NetScaler appliance for security holes. Administrators can capture failed authentication sessions on the NetScaler, measure the efficiency of the Application Firewall profiles configured on the NetScaler, and isolate policy labels that are frequently invoked by the appliance, with the help of tests mapped to this layer.

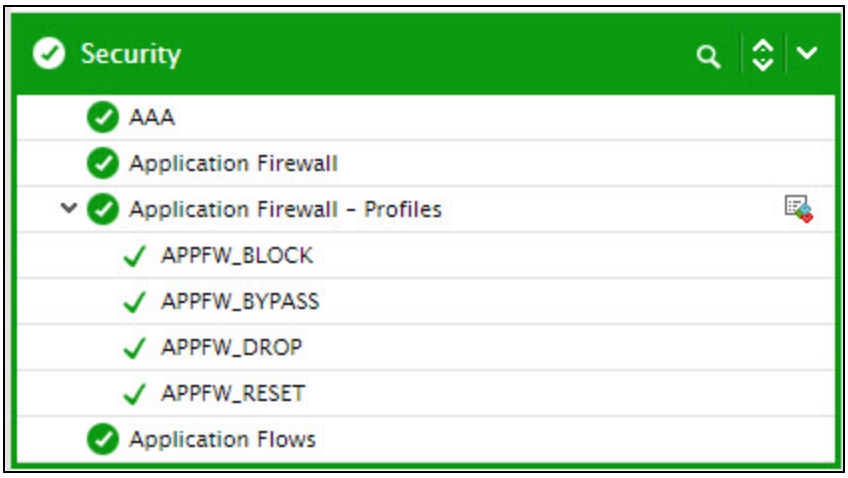

Figure 3.12: The tests mapped to the NetScaler Events layer

3.5.1 Device Events Test

In most environments where the NetScaler appliance is deployed, administrators may want to figure out the performance and efficiency of the NetScaler device. If the status of the NetScaler device is up, then they can figure out that the device is operating without a glitch! If the NetScaler device experiences any issues, then the environment may also get affected and administrators may find it diffuclt to load balance the servers sitting in the backend. For an environment to operate without any technical issues, it is the onus of the administrators to continuously monitor the NetScaler device. The **Device Events** test helps administrators in this regard!

This test scans the syslog file for the events related to the target NetScaler device and reports the number of times the status of the device was up/down/out of service. In addition, this test also helps administrators to identify the number of times the NetScaler system was started/stopped, the number of times the Route6 was up/down and the number of times the ns.conf file was read for configuration information by the NetScaler device. Using this test, administrators can identify if the performance of the NetScaler device is maintained consistently or if any performance degradation is noticed.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section **3.4.6.1**.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the

Component type, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server,

Parameter	Description
	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
Device down events	Indicates the number of times the NetScaler device was down.	Number	Ideally, the value of this measure should be zero.
Device out of service events	Indicates the number of times the status of the NetScaler device was Out of Service.	Number	
Device up events	Indicates the number of times the NetScaler device was up.	Number	
NetScaler system stopped events	Indicates the number of times the NetScaler system was stopped.	Number	
NetScaler system started events:	Indicates the number of times the NetScaler system was started.	Number	
CPU started events	Indicates the number of times the CPU of the NetScaler device was started.	Number	
Route6 down events	Indicates the number of times Route6 was down.	Number	
Route6 up events	Indicates the number of times Route6 was up.	Number	

Measurement	Description	Measurement Unit	Interpretation
SNMP module started an alarm:	Indicates the number of times an SNMP module started an alert.	Number	An alert will be generally raised when the value of a monitored attribute crosses the threshold value.
SNMP module stopped an alarm	Indicates number of times an SNMP module stopped an alarm.	Number	
NetScaler started reading configuration Events	Indicates the number of times the NetScaler device has started reading the configuration from *ns.conf* file (during boot up).	Number	A high value for this measure is a cause of concern as this may indicate that the NetScaler device is either corrupted or is often restarting.
NetScaler completed reading configuration Events	Indicates the number of times the NetScaler device has completed reading the configuration from *ns.conf* file (during boot up).	Number	

3.5.2 DHCP Events Test

DHCP provides an automated way to distribute and update IP addresses and other configuration information on a network. When you deploy Dynamic Host Configuration Protocol (DHCP) servers on your network, you can automatically provide client computers and other TCP/IP based network devices with valid IP addresses. You can also provide the additional configuration parameters these clients and devices need, called DHCP options, that allow them to connect to other network resources, such as DNS servers, WINS servers, and routers.

DHCP is a client-server technology that allows DHCP servers to assign, or lease, IP addresses to computers and other devices that are enabled as DHCP clients. With DHCP, you can do the following:

- Lease IP addresses for a specific amount of time to DHCP clients, and then automatically renew the IP addresses when the client requests a renewal.
- Update DHCP client parameters automatically by changing a server or scope option at the DHCP server rather than performing this action individually on all DHCP clients.

- Reserve IP addresses for specific computers or other devices so that they always have the same IP address and also receive the most up-to-date DHCP options.
- Exclude IP addresses or address ranges from distribution by the DHCP server so that these IP addresses and ranges can be used to statically configure servers, routers, and other devices that require static IP addresses.
- Provide DHCP services to many subnets, if all routers between the DHCP server and the subnet for which you want to provide service are configured to forward DHCP messages.
- Configure the DHCP server to perform DNS name registration services for DHCP clients.
- Provide multicast address assignment for IP-based DHCP clients.

A DHCP client initiates a conversation with a DHCP server when it is seeking a new lease, renewing a lease, rebinding, or restarting. The DHCP conversation consists of a series of DHCP messages passed between the DHCP client and DHCP servers. In an environment where the Citrix NetScaler VPX/MDX is deployed, the DHCP client- server communication happens via the NetScaler appliance. Whenever communication between a DHCP client and a DHCP server suffers a setback, say for example, a DHCP client cannot renew the acquired lease as the DHCP server that provided the original lease is offline, then, the NetScaler appliance is required to initiate a communication between the DHCP client and another DHCP server. If the DHCP client has to wait for a longer period to acquire a lease, then the DHCP client jmay move out of the network by releasing its lease. If too many DHCP clients face difficulty in acquiring a lease or renewing the lease, then severe bottlenecks will be detected in the communication between the DHCP client and DHCP server. To proactively detect such communication failure, administrators can use the **DHCP Events** test!

This test scans the syslog file for the communication between a DHCP client and DHCP server via the NetScaler and reports the number of times the DHCP server has sent an invalid setting in lieu of a DHCP request; the number of times a lease was acquired and released by a DHCP client and the number of times a Policy Based Routing policy is still dependent on a leased IP even after the release of the lease. Using this test, administrators can analyze the communication between the DHCP server and DHCP client, identify if there are any performance bottlenecks between the DHCP client and server and rectify the same.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1 .

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the

Component type, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server,

Parameter	Description
	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	Decription	Measurement Unit	Interpretation
DHCP server acquires a lease	Indicates the number of times the DHCP client acquired a lease.	Number	A DHCP-enabled client obtains a lease for an IP address from a DHCP server. Before the lease expires, the DHCP server must renew the lease for the client or the client must obtain a new lease. Leases are retained in the DHCP server database approximately one day after expiration. This grace period protects a client's lease in case the client and server are in different time zones, their internal clocks are not synchronized, or the client is off the network when the lease expires.
DHCP server releases a lease	Indicates the number of times the DHCP client released a lease.	Number	The DHCP client acquires a lease for an IP address from the DHCP server and configures its TCP/IP properties by using the DHCP option information in the reply received from the DHCP server, and completes its initialization of TCP/IP. The IP address obtained as lease from the DHCP server remains allocated to the client until the client manually releases the address, or until the lease time expires after which the DHCP server cancels the lease. Mostly, the DHCP client releases a

Measurement	Decription	Measurement Unit	Interpretation
			lease when moving to a different network. A high value for this measure therefore may indicate that the DHCP clients are moving away from the network or the DHCP server is currently offline.
DHCP lease is released and a PBR is dependent on the lease IP	Indicates the number of times the DHCP client released a lease while a Policy Based Routing (PPBR) is still dependent on the leased IP.	Number	PBR is a concept that closely relates to Access Control List (ACL) on a NetScaler appliance. PBR can be leveraged to take routing decision (next hop router) based on certain criteria such as Source IP, Source Port, Destination IP, Destination Port, Protocol, Interface, VLAN and Source MAC. Using PBR, a NetScaler appliance can either ALLOW or DENY access to network packets.
DHCP server sends an invalid setting	Indicates the number of times the DHCP server sent an invalid setting.	Number	The DHCP client acquires a lease from the DHCP server and configures its TCP/IP properties by using the DHCP option information in the reply received from the DHCP server, and completes its initialization of TCP/IP. In rare cases, a DHCP server might return a negative acknowledgment to the client. This can happen if a client requests an invalid or duplicate address. If a client receives a negative acknowledgment (DHCPNack), the client must begin the entire lease process again. This measure is a good indicator of a conflict in communication between the DHCP client and DHCP server.

3.5.3 HA Events Test

A high availability (HA) deployment of two Citrix NetScaler appliances can provide uninterrupted operation in any transaction. With one appliance configured as the primary node and the other as the

secondary node, the primary node accepts connections and manages servers while the secondary node monitors the primary. If, for any reason, the primary node is unable to accept connections, the secondary node takes over.

The secondary node monitors the primary by sending periodic messages (often called heartbeat messages or health checks) to determine whether the primary node is accepting connections. If a health check fails, the secondary node retries the connection for a specified period, after which it determines that the primary node is not functioning normally. The secondary node then takes over for the primary (a process called failover). When the secondary takes over from the primary, the configuration of both the nodes should be the same. If there exists a non-sync between the configuration of the devices, then the performance of the devices will be affected due to various external reasons like network connectivity, authentication failure etc. To avoid such non-synchronization, administrators have to frequently monitor the success/failure of the command propagation feature which helps in the synchronization process. The HA Events test helps administrators in this regard!

By carefully analyzing the syslog file, this test reports the number of times the NetScaler system in a HA setup has stopped and the number of times the command propagation failed/was successful. In addition, this test reports the number of times the NetScaler device has switched over from primary to secondary in a HA setup. Using this test, administrators may be able to figure out the effectiveness of the High availability setup of the NetScaler device.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
NetScaler system stopped events	Indicates the number of times the NetScaler system in a HA setup was stopped.	Number	
HA propagation failed	Indicates the number of times the HA Command Propagation failed.	Number	Command propagation is a feature of the NetScaler appliance that ensures that the commands run on the primary NetScaler appliance of the high availability setup are automatically run on the secondary NetScaler appliance. When you run a command on the primary appliance, this feature ensures that the command runs on the secondary appliance before it runs on the primary appliance. Ideally, the value of this measure should be zero. A HA Propagation may occur due to the following reasons: a. Network connectivity issues between the primary and secondary NetScaler appliances; b. Authentication failure between the primary and secondary appliances; c. Resources, such as Secure Socket Layer (SSL) certificates and initialization script customization are missing on the secondary appliance. Administrators therefore are required to do the following in order to maintain the least possible value for this measure:

Measurement	Description	Measurement Unit	Interpretation
			a. Check the network connectivity between the primary and secondary NetScaler appliances; b. Verify the Remote Procedure Call (RPC) node settings on both the appliances. c. Run the command directly on the secondary appliance and verify the error message. The error might have occurred because a resource required for the command exists on the primary appliance but not on the secondary appliance. Ensure that the required resource exists on the secondary appliance as well. If command execution fails on the secondary or times out when executing on the secondary, it may cause a non-sync between the configuration of the primary and the secondary.
HA propagation successful	Indicates the number of times the HA Command Propagation was successful.	Number	A high value is desired for this measure. A high success rate indicates that the configuration of the primary and secondary are in sync.
HA state changed	Indicates the number of times the HA state has changed for the NetScaler device i.e, the NetScaler device has changed from primary to secondary and vice versa.	Number	Frequent change in the high availability state of a NetScaler device indicates serious load balancing and network issues which may sometimes lead to non – synchronization between the primary and secondary devices.
Cluster state changed	Indicates the number of times the cluster state has changed.	Number	

3.5.4 Interface Events Test

Network interfaces in the NetScaler appliance are numbered in <slot>/<port> notation. By reading the statisitics polled in the syslog file, this test reports various numerical statistics relating to the states of the network interfaces such as started, stopped, hung etc. In addition, this test reports the number of times the network interface was reset and the number of times the network interface was bound/unbound to a channel.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.

Parameter	Description
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Network interface stopped events	Indicates the number of times the network interface was stopped.	Number	
Network interface started events	Indicates the number of times the network interface was started.	Number	
Network interface in hung state events	Indicates the number of times the network interface was in a hung state.	Number	

Measurement	Description	Measurement Unit	Interpretation
Network interface reset events	Indicates the number of times the Network interface was reset.	Number	A network interface is reset if its settings are to be renegotiated.
Interface bound or unbound from channel events	Indicates the number of times the network interface was bound/unbound to a channel.	Number	Link aggregation combines data coming from multiple ports into a single high-speed link. Configuring link aggregation increases the capacity and availability of the communication channel between the NetScaler appliance and other connected devices. An aggregated link is also referred to as a “channel”. When a network interface is bound to a channel, the channel parameters have precedence over the network interface parameters. (That is, the network interface parameters are ignored.) A network interface can be bound only to one channel.
Interface throughput is less than the min required	Indicates the number of times the throughput of the network interface was less than the minimum required throughput.	Number	
NIC throughput is equal or greater than the min required	Indicates the number of times the throughput of the network interface is equal or greater than the minimum required throughput.	Number	
Interface is powered on event	Indicates the number of times the network interface was powered on.	Number	
Interface is powered off event	Indicates the number of times the network interface was powered off.	Number	

3.5.5 Memory Events Test

For a NetScaler appliance to perform smoothly, the memory utilization of the NetScaler appliance should be optimal. Though high memory utilization may at times indicate that the appliance is running out of resources, it may also indicate other severe issues relating to memory. For example, when the same memory block is freed from a memory pool more than once, the the data in that memory block will be corrupted. If issues such as these are not detected at the earliest and rectified, data corruption may occur in most of the memory blocks leading to the performance degradation of the NetScaler appliance. To avoid such performance degradation due to memory related issues, administrators can use the **Memory Events** test.

This test proactively monitors the memory of the NetScaler appliance and reports the number of times bad memory was freed, the number of times duplicate memory was freed and the number of times memory was freed from a wrong memory pool. This way, administrators can figure out if there are impending issues with the memory and resolve it before serious data corruption leading to the downgrade of the NetScaler appliance occurs.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.

Parameter	Description
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Bad memory freed	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
events	times bad memory was freed.		
Duplicate memory free events	Indicates the number of times duplicate memory was freed i.e., the number of times the same memory block was freed twice.	Number	When a memory is freed, the memory is returned to the memory pool from which it was allocated. This memory can be used by some other resources. If the same memory is freed again and in the meantime is allocated to another resource, then it indicates that the same memory is allocated to multiple uses and once the second memory free command is received, data in the memory block will be corrupted. A high value for this measure is therefore a cause of concern.
Memory freed from a wrong pool	Indicates the number of times memory was freed from a wrong pool.	Number	A high value for this measure is a cause of concern as this may indicate severe data corruption.

3.5.6 Monitor Events Test

The NetScaler appliance has a set of default monitors that are automatically bound to the respective service as soon as you create the service. For example, the tcp-default monitor is bound to all TCP services and the ping-default monitor is bound to all non-TCP services. The NetScaler appliance allows multiple monitors to be bound to a service. These monitors, by default, monitor the service they are bound to, but can be configured to monitor any destination IP address, port or both. A practical example of this is a Web server which makes a call to a backend database server where, if the database server fails, you also want the Web server to be marked as Down. To be precise, the health of one service is tied to the health of another potentially related service. A service monitoring threshold is set for each service which when combined with the weight of the monitor that is bound to the service, helps you to control how many monitors must fail before you consider the service as Down. If the monitor bound to the service satisfies the threshold, then it can be inferred that the service is Up. If any monitor fails to satisfy the threshold, then the service is considered as down. Mostly services fail when there are network connectivity issues between the backend servers on the port hosting the application, or connectivity issues between the NetScaler appliance and the backend servers, or if an SSL certkey is not bound to the virtual server etc. If too many services fail, then the performance of the NetScaler appliance will begin to degrade rapidly. To keep check on the

performance degradation of the NetScaler appliance, administrators should constantly monitor the number of services that are down. The **Monitor Events** test helps administrators in this regard

This test monitors the monitors that are bound to a service in the NetScaler appliance and reports the number of times the monitor bound to the service had hit the threshold limit and the number of times the monitor bound to the service was up/down.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while

Parameter	Description
	searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Monitor bound to the service has hit threshold limit event	Indicates the number of times the monitor bound to the service had hit the threshold limit.	Number	This measure is a good check on the health of the service.
Monitor bound to the service is down event	Indicates the number of times the monitor bound to the service was down.	Number	
Monitor bound to the service is up event	Indicates the number of times the monitor bound to the service was up.	Number	

3.5.7 PITBOSS Activities Test

Pitboss, a watchdog deamon controls the processes on a NetScaler appliance. If the pitboss detects a failing process, it will try to restart it, and if the pitboss fails to restart a process in 5 attempts, on the sixth failure, the NetScaler will undergo a full reboot. If the failure of the pitboss forces the process as well as the NetScaler to reboot frequently, then with each reboot, the performance of the NetScaler appliance would be on the downhill! To detect such failures at the earliest and identify the processes that are failing, administrators can use the PITBOSS Activities test. Using this test, administrators can figure out the number of times a pitboss watch was added to a process and deleted to a process. In addition, this test throws light on the number of times the process has reached the maximum number of restarts thus allowing the pitboss to restart the NetScaler appliance.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the

Parameter	Description
	target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Pitboss watch added	Indicates the number of times pitboss watch was added on a process with the process ID.	Number	A high value for this measure is a cause of concern as this may indicate that a large number of processes running on the NetScaler appliance are failing.

Measurement	Description	Measurement Unit	Interpretation
Pitboss watch deleted	Indicates the number of times pitboss watch was deleted from a process with the process ID.	Number	
Pitboss system restarts	Indicates the number of times the process with pid had reached the maximum number of restarts thus leading to system reboot.	Number	Ideally, the value of this measure should be zero.
Pitboss process restarts	Indicates the number of times the process with pid had reached the maximum number of restarts thus leading to process reboot.	Number	Ideally, the value of this measure should be zero.

3.5.8 Routing Test

By default, the NetScaler has the ability to participate in Layer 3 routing i.e., learning and advertising routes using routing protocols such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP). BGP is the de facto standard for routing IP traffic across the global Internet. It is a routing protocol designed to share IP network address groups, or prefixes, between multiple different organizations. Each of these organizations (usually an ISP or a large private network) classify their addressing space as an **Autonomous System**–a network, or group of networks, that has an addressing and routing structure unique from another organizations. BGP is a very complex and powerful routing protocol, which some characteristics that make it quite unique compared to other protocols like RIP or OSPF. All aspects of its configuration are done manually–there is no “auto-discovery” of neighboring routers. Connections to other BGP neighbors, once configured, use a TCP connection to exchange network prefixes. Because of the possibility of mistakenly sending or receiving route information that could bring down the routing tables for neighboring routers, or even the global Internet, there are many methods of screening and applying routing policies to updates sent or received from BGP peers.

A NetScaler can run BGP as a routing protocol to learn routes from other BGP routers, as well as advertise routes that the NetScaler knows about (networks downstream, vservers, and so on). Configuring NetScaler for BGP routing involves enabling dynamic routing, adding the BGP process to the routing engine, and configuring the BGP process with the essential BGP settings: peer router addresses and Autonomous System numbers, whether to redistribute routes from the kernel and/or

defined static routes, and whether to learn routes from connected BGP peers. When a route is advertised, the BGP peers connected to the network, learns those routes. Frequent advertising of routes indicates that the peers are always kept aware of the routes in the network. If there are any errors that occur when a router advertises a route, then the availability of the router may alternate frwuently between up and down. Due to the instability of the router, the network topology is distorted which forces the routes to be withdrawn. If the instability continues, then the routes that are advertised will decrease and the routes that are withdrawn may subsequently go up resulting in high network latency. The Routing test helps administrators to keep a check on the number of routes that are advertised and the routes that are withdrawn over a period of time.

Using this test, administrators can figure out the routes that were advertised and the routes that were relearnt. In addition, this test will help administrators to identify the number of times the routes were withdrawn and the number of times the HA state of the routes was changed.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that

Parameter	Description
	are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Route advertised	Indicates the number of routes advertised.	Number	A route advertisement is when a router, using a routing protocol such as RIP or EIGRP, sends information to another router indicating that a specific network is reachable, and what the

Measurement	Description	Measurement Unit	Interpretation
			next "hop" or IP address is to use to get to the final destination. A high value is desired for this measure. The detailed diagnosis of this measure if enabled lists the routes that were advertised.
Route withdrawn	Indicates the number of routes that were withdrawn.	Number	A route flap occurs when a router alternately advertises a destination network via one route then another, or when there's an interface error on the router that alternates the availability of the router as up or down. When fluctuations are noticed repeatedly, the routing topology is distorted which in turn forces the router to withdraw the routes. Now, determining the next possible route will take longer than usual for the network thus leading to network latency or downtime. Therefore, the value of this measure should be kept at the least possible value. The detailed diagnosis of this measure if enabled, lists the routes that were withdrawn.
Route relearnt	Indicates the number of routes that were relearnt.	Number	The detailed diagnosis of this measure if enabled, lists the routes that were relearnt.
HA state changed	Indicates the number of times the HA state to the route was changed.	Number	The detailed diagnosis of this measure if enabled, lists the route for which the HA state changed.

3.6 Security Layer

The tests mapped to this layer scan the NetScaler appliance for security holes. Administrators can capture failed authentication sessions on the NetScaler, measure the efficiency of the Application

Firewall profiles configured on the NetScaler, and isolate policy labels that are frequently invoked by the appliance, with the help of tests mapped to this layer.

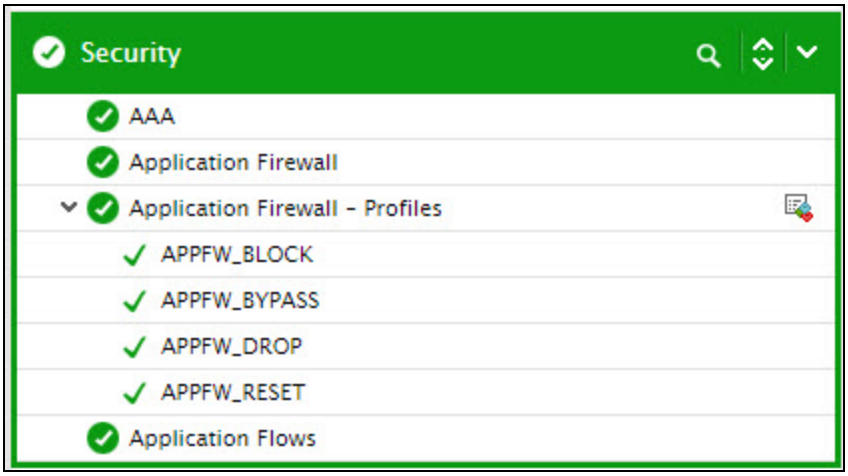

Figure 3.13: The tests mapped to the Security layer

3.6.1 AAA Stats Test

Data security is one of the important features that you must consider while making the data accessible to public over the Internet. Authenticating the user requesting for an access to the network resources is one of the methods to secure data for unauthorized access.

On a NetScaler appliance, you can use various authentication techniques to secure the data. One of the techniques is the Authentication, Authorization, and Accounting (AAA) technique which can be used when the Secure Socket Layer (SSL) Virtual Private Network (VPN) is deployed on your network. The AAA authentication technique includes three steps to secure the network. The first process, Authentication, ensures that the access is granted only to an authorized user of the network. The second process, Authorization, ensures that depending on the profile of the user, the user is authorized to perform only a set of specific tasks on the network. And finally the third process, Accounting, measures the resources the user has used during a session.

This test enables administrators to measure the effectiveness of the AAA authentication technique. This test monitors the AAA sessions on the NetScaler and reports the count and percentage of authentications that were successful and those that failed on the NetScaler. This way, the test turns the spotlight on unauthorized access attempts that were detected and prevented by the AAA technique.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Authentication successes	Indicates the number of user authentications that were successful during the last measurement period.	Number	
Authentication failures	Indicates the user authentications that failed during the last measurement period.	Number	A high value is indicative of too many authentication failures. You may want to investigate the reason for this phenomenon.
Percent of authentication successes	Indicates the percentage of user authentications that is currently successful.	Percent	A high value is desired for this measure.
HTTP authorization successes	Indicates the number of HTTP connections from the user that were	Number	

Measurement	Description	Measurement Unit	Interpretation
	authorized successfully during the last measurement period.		
HTTP authorization failures	Indicates the number of HTTP connections from the user that failed authorization during the last measurement period.	Number	A high value is indicative of too many authentication failures. You may want to investigate the reason for this phenomenon.
Percent of HTTP authorization successes	Indicates the percentage of current HTTP connections from the user that is authorized successfully.	Number	A high value is desired for this measure.
Non HTTP authorization successes	Indicates the number of connections other than the HTTP connections that were authorized successfully during the last measurement period.	Number	
Non HTTP authorization failures	Indicates the number of connections other than the HTTP connections that failed authorization during the last measurement period.	Number	A high value is indicative of too many authentication failures. You may want to investigate the reason for this phenomenon.
AAA sessions	Indicates the number of AAA sessions during the last measurement period.	Number	
Timed out AAA sessions	Indicates the number of AAA sessions that timed out during the last measurement period.	Number	NetScaler maintains a session timeout after which users must authenticate again to regain access to the intranet. This timeout is configurable. If the value of this measure is very high - i.e., timeouts appear to be occurring too often - you may want to consider changing this timeout value.

3.6.2 Application Firewall Test

NetScaler protects against a wide variety of threats with integrated security capabilities (like Application firewall) that protect applications resources, augmenting existing network-layer security protections. The NetScaler Application Firewall secures web applications, prevents inadvertent or intentional disclosure of confidential information and aids in compliance with information security regulations such as PCI-DSS.

This test tracks the network traffic flowing through the Application Firewall, and reports the different types of security check violations that have been detected by the Application Firewall. The statistics reported by this test thus serve as a good measure of the efficiency of the Application Firewall.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the amount of data received for requests by this Application firewall during the last measurement period.	MB	
Data transmitted	Indicates the amount of data transferred for responses received by this Application firewall during the last measurement period.	MB	
Requests	Indicates the number of HTTP/HTTPS requests sent to the web servers through this Application firewall during the last measurement period.	Number	
Responses	Indicates the number of HTTP/HTTPS responses sent by the web servers through this Application firewall during the last measurement period.	Number	
Aborts	Indicates the number of incomplete HTTP/HTTPS requests aborted by the client before this Application Firewall completes processing during the last measurement period.	Number	A high value for this measure could warrant an investigation.
Redirects	Indicates the number of HTTP/HTTPS requests redirected by this Application Firewall to a	Number	

Measurement	Description	Measurement Unit	Interpretation
	different web page or web server during the last measurement period.		
Long term average response time	Indicates the average backend response time of the Application firewall since reboot.	Secs	Ideally, this value should be low.
Recent average response time	Indicates the average backend response time of this Application firewall over the last 7 seconds.	Secs	Ideally, this value should be low.
Start URL	Indicates the number of Start URL security check violations detected by this Application firewall during the last measurement period.	Number	A URL that a user accesses must be explicitly listed as a Start URL. If it is not, the Citrix Application Firewall software blocks the request. By default, the Citrix Application Firewall blocks the direct access to the URLs of the Web applications. You must define the URLs you want to allow the users to access directly. Such URLs include the Web application home page, error page of the Web application; any page that you expect that the user might bookmark, and is allowed to bookmark, and any Web page that can be included in another Web site.
Deny URL	Indicates the number of Deny URL security check violations detected by this Application firewall during the last measurement period.	Number	The Deny URL check examines and blocks connections to URLs that are commonly accessed by hackers and malicious code. This check contains a list of URLs that are common targets of hackers or malicious code and that rarely if ever appear in legitimate requests. You can also add URLs or URL patterns to the list. The Deny URL check prevents attacks against

Measurement	Description	Measurement Unit	Interpretation
			various security weaknesses known to exist in web server software or on many web sites. The Deny URL check takes priority over the Start URL check, and thus denies malicious connection attempts even when a Start URL relaxation would normally allow a request to proceed.
Referer header	Indicates the number of Referrer Header security check violations detected by this Application firewall during the last measurement period.	Number	The Start URL check contains a parameter called Referer Header, which if configured, tells the Application Firewall to verify that the **Referer header** in a request that contains Web form data comes from your protected Web server rather than from another Web site. This verifies that your Web site, rather than an outside attacker, is the source of that Web form. This protects against cross-site request forgeries (CSRF) without requiring form tagging, which is more CPU-intensive than header checks.
Buffer overflow	Indicates the number of buffer overflows detected by this Application firewall during the last measurement period.	Number	The buffer overflow might result due to any of the following reasons: a. The problem arises when there is a slow Web logging client and relatively smaller chunks of data are sent from the circular buffer. With slow clients, there can be a condition when the writing of the new transactions to the buffer is much faster than the data sent from the buffer to the client. The static buffer might get circled and

Measurement	Description	Measurement Unit	Interpretation
			overwritten with new transactions before the old data is sent to the client. This results in buffer overflow, which results in loss of logging data. b. The buffer overflow might also happen due to network congestion between the Web logging client and the NetScaler appliance. To overcome the buffer overflow, you can increase the value of the buffer size. This can delay the buffer overflow condition. However, it starts again if the processing speed of the client does not match to the rate at which the data is written. The slow processing speed of the client inhibits the speed at which the data is sent to the Web logging client. You can expect some instances of buffer overflow under very high traffic conditions. If the buffer value increases continuously, then consider increasing the processing speed and the RAM of the client. Additionally, check if network congestion is preventing the client from reading data smoothly.
Cookie consistency	Indicates the number of Cookie Consistency security check violations detected by this Application firewall during the last measurement period.	Number	The Cookie Consistency check examines cookies returned by users, to verify that they match the cookies that your web site set for that user. If a modified cookie is found, it is stripped from the request before the request is forwarded to the web server. You can also configure the Cookie Consistency

Measurement	Description	Measurement Unit	Interpretation
			check to transform all of the server cookies that it processes, by encrypting the cookies, proxying the cookies, or adding flags to the cookies. This check applies to requests and responses. An attacker would normally modify a cookie to gain access to sensitive private information by posing as a previously authenticated user, or to cause a buffer overflow. The Buffer Overflow check protects against attempts to cause a buffer overflow by using a very long cookie. The Cookie Consistency check focuses on the first scenario.
CSRF form tag	Indicates the number of Cross Site Request Forgery (CSRF) security check violations detected by this Application firewall during the last measurement period.	Number	The CSRF Form Tagging check tags each web form sent by a protected web site to users with a unique and unpredictable FormID, and then examines the web forms returned by users to ensure that the supplied FormID is correct. This check protects against Cross Site Request Forgery (CSRF) attacks. This check applies only to HTML requests that contain a web form, with or without data. It does not apply to XML requests. The CSRF Form Tagging check prevents attackers from using their own web forms to send high volume form responses with data to your protected web sites. This check requires relatively little CPU processing capacity compared to certain other security checks that analyze web forms in depth. It is therefore able to handle high volume

Measurement	Description	Measurement Unit	Interpretation
			attacks without seriously degrading the performance of the protected web site or the application firewall itself.
HTML cross-site scripting	Indicates the number of html cross-site scripting attacks detected by this Application firewall during the last measurement period.	Number	A cross-site scripting attack (XSS), sends a web application an unvalidated script that activates when it is read by the browser or application to steal user identities, hijack user sessions, poison cookies, redirect users to malicious web sites, access restricted sites and even launch false advertisements. Application Firewall has dynamic context sensitive XSS attack protections that looks for anything that looks like an HTML tag and checks against allowed HTML attributes and tags to detect XSS attacks. Custom XSS patterns can be stored to modify this default list of tags and attributes. Both HTML and XML payloads are inspected. Field format protection and form field consistency is included.
HTML SQL injection	Indicates the number of HTML sql injection security check violation detected by this Application firewall during the last measurement period.	Number	The HTML SQL Injection check provides special defenses against injection of unauthorized SQL code that might break security. It examines both the headers and the POST bodies of requests for injected SQL code. If the application firewall detects unauthorized SQL code in a user request, it either transforms the request, to render the SQL code inactive, or blocks the request. Many web applications have web forms that use SQL to communicate with relational database servers. Often, the scripts that pass web form

Measurement	Description	Measurement Unit	Interpretation
			information to the database do not validate the information provided by the user before sending it to the database. Malicious code or a hacker can use the insecure web form to send SQL commands to the web server.
Field format	Indicates the number of field format security check violation detected by the Application firewall during the last measurement period.	Number	The Field Formats check verifies the data that users send to your web sites in a web form. It examines both the length and type of data to ensure that it is appropriate for the form field in which it appears. If the application firewall detects inappropriate web form data in a user request, it blocks the request. This check applies to HTML requests only. It does not apply to XML requests. By preventing an attacker from sending inappropriate web form data to your web site, the Field Formats check prevents certain types of attacks on your web site and database servers. For example, if a particular field expects the user to enter a phone number, the Field Formats check examines the user's response to ensure that the data matches the format for a phone number. If a particular field expects a first name, the Field Formats check ensures that the data in that field is of a type and length appropriate for a first name. It does the same thing for each form field that you configure it to protect.
Field consistency	Indicates the number of Form Field Consistency security check violations detected by this	Number	The Form Field Consistency check examines the web forms returned by users of your web site, and verifies that the web form was not modified

Measurement	Description	Measurement Unit	Interpretation
	Application firewall during the last measurement period.		inappropriately by the client. This check applies only to HTML requests that contain a web form, with or without data. It does not apply to XML requests. The Form Field Consistency check prevents clients from making unauthorized changes to the structure of the web forms on your web site when they are filling out a web form and submitting data by using that form. It also ensures that the data a user submits meets the HTML restrictions for length and type, and that data in hidden fields is not modified. This prevents an attacker from tampering with a web form and using the modified form to gain unauthorized access to web site, redirect the output of a contact form that uses an insecure script and thereby send unsolicited bulk email, or exploit a vulnerability in your web server software to gain control of the web server or the underlying operating system. Web forms are a weak link on many web sites and attract a wide range of attacks.
Credit card	Indicates the number of credit card security check violations detected by this Application firewall during the last measurement period.	Number	The Credit Card check provides special handling for credit card numbers. A Webapplication does not usually send a credit card number in a response to a user request, even when the user supplies a credit card number in the request. The Application Firewall examines Web server responses, including headers, for credit card numbers. If it finds a credit card

Measurement	Description	Measurement Unit	Interpretation
			number in the response, and the administrator has not configured it to allow credit card numbers to be sent, it responds in one of two ways: a. It blocks the response. b. It replaces all but the final group of digits in the credit card with x's. For example, a credit card number of *9876-5432-1234-5678* would be rendered, *xxxx-xxxx-xxxx-5678*. The Credit Card check prevents attackers from exploiting a security flaw in your Web server software or on your Web site to obtain credit card numbers of your customers. If your Web sites do not have access to credit card information, you do not need to configure this check. If your Web sites do have access to credit card information, such as via a shopping cart application, or your Web sites have access to back-end database servers that contain customer credit card numbers, you should configure protection for each type of credit card that you accept.
Safe object	Indicates the number of safe object security check violations detected by this Application firewall during the last measurement period.	Number	The Safe Object check provides user-configurable protection for sensitive business information, such as customer numbers, order numbers, and country- or region-specific telephone numbers or postal codes. A user-defined regular expression or custom plug-in tells the Application

Measurement	Description	Measurement Unit	Interpretation
			Firewall the format of this information, and defines the rules to be used to protect it. If the Application Firewall detects a string in a user request that matches a safe object definition, depending on how you configured that particular Safe Object rule, it either blocks the response, masks the protected information, or removes the protected information from the response before sending it to the user.
Signature violations	Indicates the number of signature security check violations detected by the Application firewall during the last measurement period.	Number	The Application Firewall Signatures function provides specific, configurable rules that protect your Web sites against known attacks. When properly configured, the signatures may provide all the protection that a simple Web site needs. They also provide a good level of immediate protection for more complex Web sites, allowing you to implement that protection without delay while you configure additional protections as needed.
XML format	Indicates the number of XML format security check violations detected by this Application firewall during the last measurement period.	Number	The XML Format check examines the XML format of incoming requests, and blocks those requests that are not well formed, or that do not meet the criteria in the XML specification for properly-formed XML documents. Some of those criteria are: a. An XML document must contain only properly-encoded Unicode characters that match the Unicode specification. b. No special XML syntax

Measurement	Description	Measurement Unit	Interpretation
			characters—such as “<”, “>” and "&"—can be included in the document except when used in XML markup. c. All begin, end, and empty-element tags must be correctly nested, with none missing or overlapping d. XML element tags are case-sensitive; all beginning and end tags must match exactly. e. A single root element must contain all the other elements in the XML document. A document that does not meet the XML well-formedness criteria does not meet the definition of an XML document. Strictly speaking, it is not XML. However, not all XML applications and web services enforce the well-formedness standard, and not all handle poorly-formed or invalid XML correctly. Inappropriate handling of a poorly-formed XML document can cause security breaches. The purpose of the XML Format check is to prevent a malicious user from using a poorly-formed XML request to breach security on your XML application or web service.
XML denial of service	Indicates the number of XML Denial-of-Service security check violations detected by this Application firewall during	Number	The XML Denial of Service (XML DoS or XDoS) check examines incoming XML requests to determine whether they match the characteristics of a

Measurement	Description	Measurement Unit	Interpretation
	the last measurement period.		denial-of-service (DoS) attack, and blocks those requests that do. The purpose of the XML DoS check is to prevent an attacker from using XML requests to launch a denial-of-service attack on your server or application.
XML message violations	Indicates the number of XML message validation security check violations detected by this Application firewall during the last measurement period.	Number	The XML Message Validation check examines requests that contain XML messages to ensure that they are valid. If a request contains an invalid XML message, the Application Firewall blocks the request. The purpose of the XML Validation check is to prevent an attacker from using specially-constructed invalid XML messages to breach security on your application.
Web services interoperability	Indicates the number of Web Services Interoperability (WS-I) security check violations detected by this Application firewall during the last measurement period.	Number	The Web Services Interoperability (WS-I) check examines both requests and responses for adherence to the WS-I standard, and blocks those requests and responses that do not adhere to this standard. The purpose of the WS-I check is to block requests that might not interact with other XML appropriately. An attacker can use inconsistencies in interoperability to launch an attack on your XML application.
XML SQL injection	Indicates the number of XML SQL Injection security check violations detected by this Application firewall during the last measurement period.	Number	The XML SQL Injection check examines both the headers and the bodies of user requests for possible XML SQL Injection attacks. If it finds injected SQL, it blocks the request. To prevent misusing the scripts on your protected web services to breach security on your web services, the

Measurement	Description	Measurement Unit	Interpretation
			XML SQL Injection check blocks scripts that violate the same origin rule, which states that scripts should not access or modify content on any server but the server on which they are located. Any script that violates the same origin rule is called a cross-site script, and the practice of using scripts to access or modify content on another server is called XML SQL Injection. The reason XML SQL Injection is a security issue is that a web server that allows XML SQL Injection can be attacked with a script that is not on that web server, but on a different web server, such as one owned and controlled by the attacker. Unfortunately, many companies have a large installed base of Javascript-enhanced web content that violates the same origin rule. If you enable the XML SQL Injection check on such a site, you have to generate the appropriate exceptions so that the check does not block legitimate activity. In addition, to prevent blocking of legitimate requests, this check ignores cookies that were set by the server, even if they contain elements that the Cookie Consistency check would otherwise block.
XML cross-site scripting	Indicates the number of XML Cross-site Scripting security check violations detected by this Application firewall during the last measurement period.	Number	The XML Cross-Site Scripting check examines both the headers and the bodies of user requests for possible cross-site scripting attacks. If it finds a possible cross-site scripting attack, it blocks the request.

Measurement	Description	Measurement Unit	Interpretation
			To prevent misuse of the scripts on your protected web services to breach security on your web services, the XML Cross-Site Scripting check blocks scripts that violate the same origin rule, which states that scripts should not access or modify content on any server but the server on which they are located. Any script that violates the same origin rule is called a cross-site script, and the practice of using scripts to access or modify content on another server is called cross-site scripting. The reason cross-site scripting is a security issue is that a web server that allows cross-site scripting can be attacked with a script that is not on that web server, but on a different web server, such as one owned and controlled by the attacker. Unfortunately, many companies have a large installed base of Javascript-enhanced web content that violates the same origin rule. If you enable the XML Cross-Site Scripting check on such a site, you have to generate the appropriate exceptions so that the check does not block legitimate activity. In addition, to prevent blocking of legitimate requests, this check ignores cookies that were set by the server, even if they contain elements that the Cookie Consistency check would otherwise block.
XML attachment	Indicates the number of XML Attachment security check violations detected by this Application firewall	Number	The XML Attachment check examines incoming requests for malicious attachments, and it blocks those requests that contain attachments that

Measurement	Description	Measurement Unit	Interpretation
	during the last measurement period.		might breach applications security. The purpose of the XML Attachment check is to prevent an attacker from using an XML attachment to breach security on your server.
SOAP fault violations	Indicates the number of XML Soap Fault Filtering security check violations detected by this Application firewall during the last measurement period.	Number	The XML SOAP Fault Filtering check examines responses from your protected web services and filters out XML SOAP faults. This prevents leaking of sensitive information to attackers.
XML generic violations	Indicates the number of requests returning XML generic error from the backend server through this Application firewall during the last measurement period.	Number	
Total violations	Indicates the total number of security check violations detected by this Application firewall during the last measurement period.	Number	This is a good measure of how secure your environment is.
HTTP client errors (4xx)	Indicates the number of requests returning HTTP 4xx error from the backend server during the last measurement period.	Number	The 4xx codes are intended for cases in which the client seems to have erred. Response status codes beginning with the digit "5" indicate cases in which the server is aware that it has encountered an error or is otherwise incapable of performing the request. A high value for either of these measures is a cause for concern, and would require scrutiny.
HTTP server errors (5xx)	Indicates the number of requests returning HTTP 5xx error from the backend server during the last measurement period.	Number	

3.6.3 Application Firewall – Profiles Test

An Application Firewall profile is a collection of settings that tell the Application Firewall which security checks to use when filtering a particular request or response, and how to handle a request or response that fails a security check.

The built-in profiles provide an easy way to process types of content that do not require complex filtering. The four built-in profiles are:

- **APPFW_BYPASS**: Allows requests to proceed without any filtering. You should use this profile only for requests and responses that do not require any Application Firewall protection.
- **APPFW_RESET**: Resets the connection, requiring the user to re-establish his or her session by visiting a designated start URL.
- **APPFW_DROP**: Drops the connection without response.
- **APPFW_BLOCK**: Redirects the connection to the designated error page.

For more complex types of content, you can create the following types of profile:

- **HTML profile**: Protects standard HTML-based web content.
- **XML profile**: Protects XML-based applications.
- **Web 2.0 profile**: Protects Web 2.0 content containing both XML and HTML content

This test auto-discovers the profiles (both built-in and user-configured) supported by the Application Firewall, and reports the different types of security check violations that have been detected by each profile. With the help of this test, administrators can isolate the most effective profiles and those that may require some fine-tuning.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Application Firewall Profile configure on the NetScaler appliance being monitored

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed

Parameter	Description
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the amount of data received for requests governed by this Application firewall profile during the last measurement period.	MB	
Data transmitted	Indicates the amount of data transferred for responses received by this Application firewall profile during the last measurement period.	MB	
Requests	Indicates the number of HTTP/HTTPS requests sent to the web servers through this Application firewall profile during the last measurement period.	Number	
Responses	Indicates the number of HTTP/HTTPS responses	Number	

Measurement	Description	Measurement Unit	Interpretation
	sent by the web servers through this Application firewall profile during the last measurement period.		
Aborts	Indicates the number of incomplete HTTP/HTTPS requests aborted by the client before this Application firewall profile completes processing during the last measurement period.	Number	A high value for this measure could warrant an investigation.
Redirects	Indicates the number of HTTP/HTTPS requests redirected by this Application firewall profile to a different web page or web server during the last measurement period.	Number	
Long term average response time	Indicates the average backend response time of this Application firewall profile since reboot.	Secs	Ideally, this value should be low.
Recent average response time	Indicates the average backend response time of this Application firewall profile over the last 7 seconds.	Secs	Ideally, this value should be low.
Start URL	Indicates the number of Start URL security check violations detected by this Application firewall profile during the last measurement period.	Number	A URL that a user accesses must be explicitly listed as a Start URL. If it is not, the Citrix Application Firewall software blocks the request. By default, the Citrix Application Firewall blocks the direct access to the URLs of the Web applications. You must define the URLs you want to allow the users to access directly.

Measurement	Description	Measurement Unit	Interpretation
			Such URLs include the Web application home page, error page of the Web application; any page that you expect that the user might bookmark, and is allowed to bookmark, and any Web page that can be included in another Web site.
Deny URL	Indicates the number of Deny URL security check violations detected by this Application firewall profile during the last measurement period.	Number	The Deny URL check examines and blocks connections to URLs that are commonly accessed by hackers and malicious code. This check contains a list of URLs that are common targets of hackers or malicious code and that rarely if ever appear in legitimate requests. You can also add URLs or URL patterns to the list. The Deny URL check prevents attacks against various security weaknesses known to exist in web server software or on many web sites. The Deny URL check takes priority over the Start URL check, and thus denies malicious connection attempts even when a Start URL relaxation would normally allow a request to proceed.
Referer header	Indicates the number of Referrer Header security check violations detected by this Application firewall profile during the last measurement period.	Number	The Start URL check contains a parameter called Referer Header, which if configured, tells the Application Firewall to verify that the Referer header in a request that contains Web form data comes from your protected Web server rather than from another Web site. This verifies that your Web site, rather than an outside attacker, is the source of that Web form. This protects against cross-site request forgeries (CSRF)

Measurement	Description	Measurement Unit	Interpretation
			without requiring form tagging, which is more CPU-intensive than header checks.
Buffer overflow	Indicates the number of buffer overflows detected by this Application firewall profile during the last measurement period.	Number	The buffer overflow might result due to any of the following reasons: a. The problem arises when there is a slow Web logging client and relatively smaller chunks of data are sent from the circular buffer. With slow clients, there can be a condition when the writing of the new transactions to the buffer is much faster than the data sent from the buffer to the client. The static buffer might get circled and overwritten with new transactions before the old data is sent to the client. This results in buffer overflow, which results in loss of logging data. b. The buffer overflow might also happen due to network congestion between the Web logging client and the NetScaler appliance. To overcome the buffer overflow, you can increase the value of the buffer size. This can delay the buffer overflow condition. However, it starts again if the processing speed of the client does not match to the rate at which the data is written. The slow processing speed of the client inhibits the speed at which the data is sent to

Measurement	Description	Measurement Unit	Interpretation
			the Web logging client. You can expect some instances of buffer overflow under very high traffic conditions. If the buffer value increases continuously, then consider increasing the processing speed and the RAM of the client. Additionally, check if network congestion is preventing the client from reading data smoothly.
Cookie consistency	Indicates the number of Cookie Consistency security check violations detected by this Application firewall profile during the last measurement period.	Number	The Cookie Consistency check examines cookies returned by users, to verify that they match the cookies that your web site set for that user. If a modified cookie is found, it is stripped from the request before the request is forwarded to the web server. You can also configure the Cookie Consistency check to transform all of the server cookies that it processes, by encrypting the cookies, proxying the cookies, or adding flags to the cookies. This check applies to requests and responses. An attacker would normally modify a cookie to gain access to sensitive private information by posing as a previously authenticated user, or to cause a buffer overflow. The Buffer Overflow check protects against attempts to cause a buffer overflow by using a very long cookie. The Cookie Consistency check focuses on the first scenario.
CSRF form tag	Indicates the number of Cross Site Request Forgery (CSRF) security check violations detected	Number	The CSRF Form Tagging check tags each web form sent by a protected web site to users with a unique and

Measurement	Description	Measurement Unit	Interpretation
	by this Application firewall profile during the last measurement period.		unpredictable FormID, and then examines the web forms returned by users to ensure that the supplied FormID is correct. This check protects against Cross Site Request Forgery (CSRF) attacks. This check applies only to HTML requests that contain a web form, with or without data. It does not apply to XML requests. The CSRF Form Tagging check prevents attackers from using their own web forms to send high volume form responses with data to your protected web sites. This check requires relatively little CPU processing capacity compared to certain other security checks that analyze web forms in depth. It is therefore able to handle high volume attacks without seriously degrading the performance of the protected web site or the application firewall itself.
HTML cross-site scripting	Indicates the number of html cross-site scripting attacks detected by this Application firewall profile during the last measurement period.	Number	A cross-site scripting attack (XSS), sends a web application an unvalidated script that activates when it is read by the browser or application to steal user identities, hijack user sessions, poison cookies, redirect users to malicious web sites, access restricted sites and even launch false advertisements. Application Firewall has dynamic context sensitive XSS attack protections that looks for anything that looks like an HTML tag and checks against allowed HTML attributes and tags to detect XSS attacks. Custom XSS patterns can be stored to modify this default list of tags

Measurement	Description	Measurement Unit	Interpretation
			and attributes. Both HTML and XML payloads are inspected. Field format protection and form field consistency is included.
HTML SQL injection	Indicates the number of HTML sql injection security check violation detected by this Application firewall profile during the last measurement period.	Number	The HTML SQL Injection check provides special defenses against injection of unauthorized SQL code that might break security. It examines both the headers and the POST bodies of requests for injected SQL code. If the application firewall detects unauthorized SQL code in a user request, it either transforms the request, to render the SQL code inactive, or blocks the request. Many web applications have web forms that use SQL to communicate with relational database servers. Often, the scripts that pass web form information to the database do not validate the information provided by the user before sending it to the database. Malicious code or a hacker can use the insecure web form to send SQL commands to the web server.
Field format	Indicates the number of field format security check violation detected by this Application firewall profile during the last measurement period.	Number	The Field Formats check verifies the data that users send to your web sites in a web form. It examines both the length and type of data to ensure that it is appropriate for the form field in which it appears. If the application firewall detects inappropriate web form data in a user request, it blocks the request. This check applies to HTML requests only. It does not apply to XML requests. By preventing an attacker from

Measurement	Description	Measurement Unit	Interpretation
			sending inappropriate web form data to your web site, the Field Formats check prevents certain types of attacks on your web site and database servers. For example, if a particular field expects the user to enter a phone number, the Field Formats check examines the user's response to ensure that the data matches the format for a phone number. If a particular field expects a first name, the Field Formats check ensures that the data in that field is of a type and length appropriate for a first name. It does the same thing for each form field that you configure it to protect.
Field consistency	Indicates the number of Form Field Consistency security check violations detected by this Application firewall profile during the last measurement period.	Number	The Form Field Consistency check examines the web forms returned by users of your web site, and verifies that the web form was not modified inappropriately by the client. This check applies only to HTML requests that contain a web form, with or without data. It does not apply to XML requests. The Form Field Consistency check prevents clients from making unauthorized changes to the structure of the web forms on your web site when they are filling out a web form and submitting data by using that form. It also ensures that the data a user submits meets the HTML restrictions for length and type, and that data in hidden fields is not modified. This prevents an attacker from tampering with a web form and using the modified form to gain unauthorized access to

Measurement	Description	Measurement Unit	Interpretation
			web site, redirect the output of a contact form that uses an insecure script and thereby send unsolicited bulk email, or exploit a vulnerability in your web server software to gain control of the web server or the underlying operating system. Web forms are a weak link on many web sites and attract a wide range of attacks.
Credit card	Indicates the number of credit card security check violations detected by this Application firewall profile during the last measurement period.	Number	The Credit Card check provides special handling for credit card numbers. A Webapplication does not usually send a credit card number in a response to a user request, even when the user supplies a credit card number in the request. The Application Firewall examines Web server responses, including headers, for credit card numbers. If it finds a credit card number in the response, and the administrator has not configured it to allow credit card numbers to be sent, it responds in one of two ways: a. It blocks the response. b. It replaces all but the final group of digits in the credit card with x's. For example, a credit card number of 9876-5432-1234-5678 would be rendered, xxxx-xxxx-xxxx-5678. The Credit Card check prevents attackers from exploiting a security flaw in your Web server software or on your Web site to obtain credit card numbers of your customers. If your Web sites do not have access to credit

Measurement	Description	Measurement Unit	Interpretation
			card information, you do not need to configure this check. If your Web sites do have access to credit card information, such as via a shopping cart application, or your Web sites have access to back-end database servers that contain customer credit card numbers, you should configure protection for each type of credit card that you accept.
Safe object	Indicates the number of safe object security check violations detected by this Application firewall profile during the last measurement period.	Number	The Safe Object check provides user-configurable protection for sensitive business information, such as customer numbers, order numbers, and country- or region-specific telephone numbers or postal codes. A user-defined regular expression or custom plug-in tells the Application Firewall the format of this information, and defines the rules to be used to protect it. If the Application Firewall detects a string in a user request that matches a safe object definition, depending on how you configured that particular Safe Object rule, it either blocks the response, masks the protected information, or removes the protected information from the response before sending it to the user.
Signature violations	Indicates the number of signature security check violations detected by this Application firewall profile during the last measurement period.	Number	The Application Firewall Signatures function provides specific, configurable rules that protect your Web sites against known attacks. When properly configured, the signatures may provide all the protection that a simple Web site needs. They also provide a good level of immediate protection for more complex Web sites, allowing you to

Measurement	Description	Measurement Unit	Interpretation
			implement that protection without delay while you configure additional protections as needed.
XML format	Indicates the number of XML format security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML Format check examines the XML format of incoming requests, and blocks those requests that are not well formed, or that do not meet the criteria in the XML specification for properly-formed XML documents. Some of those criteria are: a. An XML document must contain only properly-encoded Unicode characters that match the Unicode specification. b. No special XML syntax characters—such as “<“, “>” and "&"—can be included in the document except when used in XML markup. c. All begin, end, and empty-element tags must be correctly nested, with none missing or overlapping d. XML element tags are case-sensitive; all beginning and end tags must match exactly. e. A single root element must contain all the other elements in the XML document. A document that does not meet the XML well-formedness criteria does not meet the definition of an XML

Measurement	Description	Measurement Unit	Interpretation
			document. Strictly speaking, it is not XML. However, not all XML applications and web services enforce the well-formedness standard, and not all handle poorly-formed or invalid XML correctly. Inappropriate handling of a poorly-formed XML document can cause security breaches. The purpose of the XML Format check is to prevent a malicious user from using a poorly-formed XML request to breach security on your XML application or web service.
XML denial of service	Indicates the number of XML Denial-of-Service security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML Denial of Service (XML DoS or XDoS) check examines incoming XML requests to determine whether they match the characteristics of a denial-of-service (DoS) attack, and blocks those requests that do. The purpose of the XML DoS check is to prevent an attacker from using XML requests to launch a denial-of-service attack on your server or application.
XML message violations	Indicates the number of XML message validation security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML Message Validation check examines requests that contain XML messages to ensure that they are valid. If a request contains an invalid XML message, the Application Firewall blocks the request. The purpose of the XML Validation check is to prevent an attacker from using specially-constructed invalid XML messages to breach security on your application.
Web services interoperability:	Indicates the number of XML message validation security check violations detected by this	Number	The Web Services Interoperability (WS-I) check examines both requests and responses for adherence to the WS-I standard, and blocks those

Measurement	Description	Measurement Unit	Interpretation
	Application firewall profile during the last measurement period.		requests and responses that do not adhere to this standard. The purpose of the WS-I check is to block requests that might not interact with other XML appropriately. An attacker can use inconsistencies in interoperability to launch an attack on your XML application.
XML SQL injection	Indicates the number of XML SQL Injection security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML SQL Injection check examines both the headers and the bodies of user requests for possible XML SQL Injection attacks. If it finds injected SQL, it blocks the request. To prevent misusing the scripts on your protected web services to breach security on your web services, the XML SQL Injection check blocks scripts that violate the same origin rule, which states that scripts should not access or modify content on any server but the server on which they are located. Any script that violates the same origin rule is called a cross-site script, and the practice of using scripts to access or modify content on another server is called XML SQL Injection. The reason XML SQL Injection is a security issue is that a web server that allows XML SQL Injection can be attacked with a script that is not on that web server, but on a different web server, such as one owned and controlled by the attacker. Unfortunately, many companies have a large installed base of Javascript-enhanced web content that violates the same origin rule. If you enable the XML SQL Injection check on such a

Measurement	Description	Measurement Unit	Interpretation
			site, you have to generate the appropriate exceptions so that the check does not block legitimate activity. In addition, to prevent blocking of legitimate requests, this check ignores cookies that were set by the server, even if they contain elements that the Cookie Consistency check would otherwise block.
XML cross-site scripting	Indicates the number of XML Cross-site Scripting security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML Cross-Site Scripting check examines both the headers and the bodies of user requests for possible cross-site scripting attacks. If it finds a possible cross-site scripting attack, it blocks the request. To prevent misuse of the scripts on your protected web services to breach security on your web services, the XML Cross-Site Scripting check blocks scripts that violate the same origin rule, which states that scripts should not access or modify content on any server but the server on which they are located. Any script that violates the same origin rule is called a cross-site script, and the practice of using scripts to access or modify content on another server is called cross-site scripting. The reason cross-site scripting is a security issue is that a web server that allows cross-site scripting can be attacked with a script that is not on that web server, but on a different web server, such as one owned and controlled by the attacker. Unfortunately, many companies have a large installed base of Javascript-enhanced web content that violates

Measurement	Description	Measurement Unit	Interpretation
			the same origin rule. If you enable the XML Cross-Site Scripting check on such a site, you have to generate the appropriate exceptions so that the check does not block legitimate activity. In addition, to prevent blocking of legitimate requests, this check ignores cookies that were set by the server, even if they contain elements that the Cookie Consistency check would otherwise block.
XML attachment	Indicates the number of XML Attachment security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML Attachment check examines incoming requests for malicious attachments, and it blocks those requests that contain attachments that might breach applications security. The purpose of the XML Attachment check is to prevent an attacker from using an XML attachment to breach security on your server.
SOAP fault violations	Indicates the number of XML Soap Fault Filtering security check violations detected by this Application firewall profile during the last measurement period.	Number	The XML SOAP Fault Filtering check examines responses from your protected web services and filters out XML SOAP faults. This prevents leaking of sensitive information to attackers.
XML generic violations	Indicates the number of requests returning XML generic error from the backend server through this Application firewall profile during the last measurement period.	Number	
Total violations	Indicates the total number of security check violations detected by this Application firewall profile	Number	This measure is a good indicator of how secure your environment is.

Measurement	Description	Measurement Unit	Interpretation
	during the last measurement period.		
HTTP client errors (4xx)	Indicates the number of requests returning HTTP 4xx error from the backend server during the last measurement period.	Number	The 4xx codes are intended for cases in which the client seems to have erred. Response status codes beginning with the digit "5" indicate cases in which the server is aware that it has encountered an error or is otherwise incapable of performing the request. A high value for either of these measures is a cause for concern, and would require scrutiny.
HTTP server errors (5xx)	Indicates the number of requests returning HTTP 5xx error from the backend server during the last measurement period.	Number	

3.6.4 Application Firewall Violations Test

NetScaler protects against a wide variety of threats with integrated security capabilities (like Application firewall) that protect applications resources, augmenting existing network-layer security protections. The NetScaler Application Firewall secures web applications, prevents inadvertent or intentional disclosure of confidential information and aids in compliance with information security regulations such as PCI-DSS.

This test reports the different types of security check violations that have been detected by the Application Firewall. The statistics reported by this test thus serve as a good measure of the efficiency of the Application Firewall.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability

Parameter	Description
	• 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
Buffer overflow	Indicates the number of buffer overflows detected by the Application firewall.	Number	The buffer overflow might result due to any of the following reasons: a. The problem arises when there is a slow Web logging client and relatively smaller chunks of data are sent from the circular buffer. With slow clients, there can be a condition when the writing of the new transactions to the buffer is much faster than the data sent from the buffer to the client. The static buffer might get circled and overwritten with new transactions before the old data is sent to the client. This results in buffer overflow, which results in loss of logging data. b. The buffer overflow might also happen due to network congestion between the Web logging client and the NetScaler appliance. To overcome the buffer overflow, you can increase the value of the buffer size. This can delay the buffer overflow condition. However, it starts again if the processing speed of the

Measurement	Description	Measurement Unit	Interpretation
			client does not match to the rate at which the data is written. The slow processing speed of the client inhibits the speed at which the data is sent to the Web logging client. You can expect some instances of buffer overflow under very high traffic conditions. If the buffer value increases continuously, then consider increasing the processing speed and the RAM of the client. Additionally, check if network congestion is preventing the client from reading data smoothly.
Field consistency	Indicates the number of Form Field Consistency security check violations detected by the Application firewall.	Number	The Form Field Consistency check examines the web forms returned by users of your web site, and verifies that the web form was not modified inappropriately by the client. This check applies only to HTML requests that contain a web form, with or without data. It does not apply to XML requests. The Form Field Consistency check prevents clients from making unauthorized changes to the structure of the web forms on your web site when they are filling out a web form and submitting data by using that form. It also ensures that the data a user submits meets the HTML restrictions for length and type, and that data in hidden fields is not modified. This prevents an attacker from tampering with a web form and using the modified form to gain unauthorized access to web site, redirect the output of a contact form that uses an insecure script and thereby send unsolicited

Measurement	Description	Measurement Unit	Interpretation
			bulk email, or exploit a vulnerability in your web server software to gain control of the web server or the underlying operating system. Web forms are a weak link on many web sites and attract a wide range of attacks.
Start URL	Indicates the number of Start URL security check violations detected by the Application firewall.	Number	A URL that a user accesses must be explicitly listed as a Start URL. If it is not, the Citrix Application Firewall software blocks the request. By default, the Citrix Application Firewall blocks the direct access to the URLs of the Web applications. You must define the URLs you want to allow the users to access directly. Such URLs include the Web application home page, error page of the Web application; any page that you expect that the user might bookmark, and is allowed to bookmark, and any Web page that can be included in another Web site.
Deny URL	Indicates the number of Deny URL security check violations detected by the Application firewall.	Number	The Deny URL check examines and blocks connections to URLs that are commonly accessed by hackers and malicious code. This check contains a list of URLs that are common targets of hackers or malicious code and that rarely if ever appear in legitimate requests. You can also add URLs or URL patterns to the list. The Deny URL check prevents attacks against various security weaknesses known to exist in web server software or on many web sites. The Deny URL check takes priority over the Start URL check, and thus

Measurement	Description	Measurement Unit	Interpretation
			denies malicious connection attempts even when a Start URL relaxation would normally allow a request to proceed.
XML SQL injection	Indicates the number of XML SQL Injection security check violations detected by the Application firewall.	Number	The XML SQL Injection check examines both the headers and the bodies of user requests for possible XML SQL Injection attacks. If it finds injected SQL, it blocks the request. To prevent misusing the scripts on your protected web services to breach security on your web services, the XML SQL Injection check blocks scripts that violate the same origin rule, which states that scripts should not access or modify content on any server but the server on which they are located. Any script that violates the same origin rule is called a cross-site script, and the practice of using scripts to access or modify content on another server is called XML SQL Injection. The reason XML SQL Injection is a security issue is that a web server that allows XML SQL Injection can be attacked with a script that is not on that web server, but on a different web server, such as one owned and controlled by the attacker. Unfortunately, many companies have a large installed base of Javascript-enhanced web content that violates the same origin rule. If you enable the XML SQL Injection check on such a site, you have to generate the appropriate exceptions so that the check does not block legitimate activity. In addition, to prevent

Measurement	Description	Measurement Unit	Interpretation
			blocking of legitimate requests, this check ignores cookies that were set by the server, even if they contain elements that the Cookie Consistency check would otherwise block.
XML cross-site script	Indicates the number of XML Cross-site Scripting security check violations detected by the Application firewall.	Number	The XML Cross-Site Scripting check examines both the headers and the bodies of user requests for possible cross-site scripting attacks. If it finds a possible cross-site scripting attack, it blocks the request. To prevent misuse of the scripts on your protected web services to breach security on your web services, the XML Cross-Site Scripting check blocks scripts that violate the same origin rule, which states that scripts should not access or modify content on any server but the server on which they are located. Any script that violates the same origin rule is called a cross-site script, and the practice of using scripts to access or modify content on another server is called cross-site scripting. The reason cross-site scripting is a security issue is that a web server that allows cross-site scripting can be attacked with a script that is not on that web server, but on a different web server, such as one owned and controlled by the attacker. Unfortunately, many companies have a large installed base of Javascript-enhanced web content that violates the same origin rule. If you enable the XML Cross-Site Scripting check on such a site, you have to generate the appropriate exceptions so that the

Measurement	Description	Measurement Unit	Interpretation
			check does not block legitimate activity. In addition, to prevent blocking of legitimate requests, this check ignores cookies that were set by the server, even if they contain elements that the Cookie Consistency check would otherwise block.
SQL injection	Indicates the number of HTML sql injection security check violation detected by the Application firewall.	Number	The HTML SQL Injection check provides special defenses against injection of unauthorized SQL code that might break security. It examines both the headers and the POST bodies of requests for injected SQL code. If the application firewall detects unauthorized SQL code in a user request, it either transforms the request, to render the SQL code inactive, or blocks the request. Many web applications have web forms that use SQL to communicate with relational database servers. Often, the scripts that pass web form information to the database do not validate the information provided by the user before sending it to the database. Malicious code or a hacker can use the insecure web form to send SQL commands to the web server.
Cookie consistency	Indicates the number of Cookie Consistency security check violations detected by the Application firewall.	Number	The Cookie Consistency check examines cookies returned by users, to verify that they match the cookies that your web site set for that user. If a modified cookie is found, it is stripped from the request before the request is forwarded to the web server. You can also configure the Cookie Consistency check to transform all of the server cookies that it processes, by

Measurement	Description	Measurement Unit	Interpretation
			encrypting the cookies, proxying the cookies, or adding flags to the cookies. This check applies to requests and responses. An attacker would normally modify a cookie to gain access to sensitive private information by posing as a previously authenticated user, or to cause a buffer overflow. The Buffer Overflow check protects against attempts to cause a buffer overflow by using a very long cookie. The Cookie Consistency check focuses on the first scenario.
CSRF form tag	Indicates the number of Cross Site Request Forgery (CSRF) security check violations detected by this Application firewall.	Number	The CSRF Form Tagging check tags each web form sent by a protected web site to users with a unique and unpredictable FormID, and then examines the web forms returned by users to ensure that the supplied FormID is correct. This check protects against Cross Site Request Forgery (CSRF) attacks. This check applies only to HTML requests that contain a web form, with or without data. It does not apply to XML requests. The CSRF Form Tagging check prevents attackers from using their own web forms to send high volume form responses with data to your protected web sites. This check requires relatively little CPU processing capacity compared to certain other security checks that analyze web forms in depth. It is therefore able to handle high volume attacks without seriously degrading the performance of the protected web

Measurement	Description	Measurement Unit	Interpretation
			site or the application firewall itself.
Field format	Indicates the number of field format security check violation detected by the Application firewall.	Number	The Field Formats check verifies the data that users send to your web sites in a web form. It examines both the length and type of data to ensure that it is appropriate for the form field in which it appears. If the application firewall detects inappropriate web form data in a user request, it blocks the request. This check applies to HTML requests only. It does not apply to XML requests. By preventing an attacker from sending inappropriate web form data to your web site, the Field Formats check prevents certain types of attacks on your web site and database servers. For example, if a particular field expects the user to enter a phone number, the Field Formats check examines the user's response to ensure that the data matches the format for a phone number. If a particular field expects a first name, the Field Formats check ensures that the data in that field is of a type and length appropriate for a first name. It does the same thing for each form field that you configure it to protect.
Credit card	Indicates the number of credit card security check violations detected by the Application firewall.	Number	The Credit Card check provides special handling for credit card numbers. A Webapplication does not usually send a credit card number in a response to a user request, even when the user supplies a credit card number in the request. The Application Firewall examines Web server responses, including headers, for credit card

Measurement	Description	Measurement Unit	Interpretation
			numbers. If it finds a credit card number in the response, and the administrator has not configured it to allow credit card numbers to be sent, it responds in one of two ways: a. It blocks the response. b. It replaces all but the final group of digits in the credit card with x's. For example, a credit card number of 9876-5432-1234-5678 would be rendered, xxxx-xxxx-xxxx-5678. The Credit Card check prevents attackers from exploiting a security flaw in your Web server software or on your Web site to obtain credit card numbers of your customers. If your Web sites do not have access to credit card information, you do not need to configure this check. If your Web sites do have access to credit card information, such as via a shopping cart application, or your Web sites have access to back-end database servers that contain customer credit card numbers, you should configure protection for each type of credit card that you accept.
XML format	Indicates the number of XML format security check violations detected by this Application firewall.	Number	The XML Format check examines the XML format of incoming requests, and blocks those requests that are not well formed, or that do not meet the criteria in the XML specification for properly-formed XML documents. Some of those criteria are: a. An XML document must contain

Measurement	Description	Measurement Unit	Interpretation
			only properly-encoded Unicode characters that match the Unicode specification. b. No special XML syntax characters—such as “<”, “>” and "&"—can be included in the document except when used in XML markup. c. All begin, end, and empty-element tags must be correctly nested, with none missing or overlapping d. XML element tags are case-sensitive; all beginning and end tags must match exactly. e. A single root element must contain all the other elements in the XML document. A document that does not meet the XML well-formedness criteria does not meet the definition of an XML document. Strictly speaking, it is not XML. However, not all XML applications and web services enforce the well-formedness standard, and not all handle poorly-formed or invalid XML correctly. Inappropriate handling of a poorly-formed XML document can cause security breaches. The purpose of the XML Format check is to prevent a malicious user from using a poorly-formed XML request to breach security on your XML application or web service.

Measurement	Description	Measurement Unit	Interpretation
Cross-site script	Indicates the number of html cross-site scripting attacks detected by this Application firewall.	Number	A cross-site scripting attack (XSS), sends a web application an unvalidated script that activates when it is read by the browser or application to steal user identities, hijack user sessions, poison cookies, redirect users to malicious web sites, access restricted sites and even launch false advertisements. Application Firewall has dynamic context sensitive XSS attack protections that looks for anything that looks like an HTML tag and checks against allowed HTML attributes and tags to detect XSS attacks. Custom XSS patterns can be stored to modify this default list of tags and attributes. Both HTML and XML payloads are inspected. Field format protection and form field consistency is included.
Referer header	Indicates the number of Referrer Header security check violations detected by the Application firewall.	Number	The Start URL check contains a parameter called **Referer Header**, which if configured, tells the Application Firewall to verify that the Referer header in a request that contains Web form data comes from your protected Web server rather than from another Web site. This verifies that your Web site, rather than an outside attacker, is the source of that Web form. This protects against cross-site request forgeries (CSRF) without requiring form tagging, which is more CPU-intensive than header checks.
Safe object	Indicates the number of safe object security check violations detected by the Application firewall.	Number	The Safe Object check provides user-configurable protection for sensitive business information, such as customer numbers, order numbers,

Measurement	Description	Measurement Unit	Interpretation
			and country- or region-specific telephone numbers or postal codes. A user-defined regular expression or custom plug-in tells the Application Firewall the format of this information, and defines the rules to be used to protect it. If the Application Firewall detects a string in a user request that matches a safe object definition, depending on how you configured that particular Safe Object rule, it either blocks the response, masks the protected information, or removes the protected information from the response before sending it to the user.
XML DOS	Indicates the number of XML Denial-of-Service security check violations detected by the Application firewall.	Number	The XML Denial of Service (XML DoS or XDoS) check examines incoming XML requests to determine whether they match the characteristics of a denial-of-service (DoS) attack, and blocks those requests that do. The purpose of the XML DoS check is to prevent an attacker from using XML requests to launch a denial-of-service attack on your server or application.
WSI	Indicates the number of Web Services Interoperability (WS-I) security check violations detected by the Application firewall.	Number	The Web Services Interoperability (WS-I) check examines both requests and responses for adherence to the WS-I standard, and blocks those requests and responses that do not adhere to this standard. The purpose of the WS-I check is to block requests that might not interact with other XML appropriately. An attacker can use inconsistencies in interoperability to launch an attack on your XML application.
XML attachments	Indicates the number of	Number	The XML Attachment check examines

Measurement	Description	Measurement Unit	Interpretation
	XML Attachment security check violations detected by the Application firewall.		incoming requests for malicious attachments, and it blocks those requests that contain attachments that might breach applications security. The purpose of the XML Attachment check is to prevent an attacker from using an XML attachment to breach security on your server.
SOAP faults	Indicates the number of XML Soap Fault Filtering security check violations detected by the Application firewall.	Number	The XML SOAP Fault Filtering check examines responses from your protected web services and filters out XML SOAP faults. This prevents leaking of sensitive information to attackers.
Policy hits	Indicates the number of policy hit violations detected by the Application firewall.	Number	
XML Generic	Indicates the number of requests returning XML generic error from the backend server through the Application firewall.	Number	

3.6.5 Application Flows Test

The Citrix NetScaler appliance is a central point of control for all application traffic in the data center. It collects flow and user-session level information valuable for application performance monitoring, analytics, and business intelligence applications. The AppFlow freature, when enabled on the NetScaler, transmits this information by using the Internet Protocol Flow Information eXport (IPFIX) format, which is an open Internet Engineering Task Force (IETF) standard defined in RFC 5101. Using UDP as the transport protocol, AppFlow transmits the collected data, called flow records, to one or more IPv4 collectors. The collectors aggregate the flow records and generate real-time or historical reports.

To enable administrators to promptly detect bottlenecks (if any) in the transmission of flow records to configured collectors, the **Application Flows** test runs periodic checks on data transmissions from the NetScaler appliance.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Appflow transmitted	Indicates the number of Appflow flows transmitted by the NetScaler during the last measurement period.	Number	
Appflow not transmitted	Indicates the number of Appflows that were not transmitted during the last measurement period.	Number	A high value is indicative of a transmission bottleneck. This could be owing to improper collector configuration or a bad network link

Measurement	Description	Measurement Unit	Interpretation
			between the appliance and the collector.
Appflow data transmitted	Indicates the amount of Appflow octets transmitted by the NetScaler during the last measurement period.	MB	
Appflow data not transmitted	Indicates the amount of Appflow octets that were not transmitted during the last measurement period.	MB	A high value is indicative of a transmission bottleneck. This could be owing to improper collector configuration or a bad network link between the appliance and the collector.
Appflow data ignored	Indicates the amount of Appflow octets that were ignored by the NetScaler during the last measurement period.	MB	
Appflow packets ignored	Indicates the number of Appflow packets that were ignored by the NetScaler during the last measurement period.	Number	
Appflow packets not transmitted	Indicates the number of Appflow packets that were not transmitted by the NetScaler during the last measurement period.	Number	A high value is indicative of a transmission bottleneck. This could be owing to improper collector configuration or a bad network link between the appliance and the collector.
Appflow messages transmitted	Indicates the number of Appflow messages transmitted by the NetScaler during the last measurement period.	Number	

3.6.6 Authorization Policies Test

For many NetScaler features, policies control how a feature evaluates data, which ultimately determines what the feature does with the data. A policy uses a logical expression, also called a rule, to evaluate requests, responses, or other data, and applies one or more actions determined by the outcome of the evaluation. Alternatively, a policy can apply a profile, which defines a complex action.

A policy label is a user-defined bind point to which policies are bound. When a policy label is invoked, all the policies bound to it are evaluated in the order of the priority you configured. A policy label can include one or more policies, each of which can be assigned its own result. A match on one policy in the policy label can result in proceeding to the next policy, invoking a different policy label or appropriate resource, or an immediate end to policy evaluation and return of control to the policy that invoked the policy label. You can create policy labels only for advanced policies.

This test auto-discovers the policy labels that have been configured on the NetScaler, and reports the number of a times each policy label was invoked.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each policy label configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Policy label hits	Indicates the number of times this policy label of the NetScaler was invoked.	Number	Compare the value of this measure across labels to identify which label was invoked often.

3.6.7 SSL Logs Test

A Citrix NetScaler appliance communicates via a secure communication channel with other servers and clients. The NetScaler appliance uses SSL for a safe and secure transaction. If the SSL communication channels suffer a set back with an expired SSL certificate or a number of SSL handshake failures, then the NetsScaler appliance may be prone to malicious attacks. In order to secure the NetScaler appliance, administrators should constantly keep a check on the SSL certficates, handskahes and the Certificate Revocation lists. The **SSL Logs** test exactly helps administrators in this regard. Using this test, administrators can figure out the success and failure count of the SSL handshakes and also be proactively warned of an impending SSL certificate expiry. In addition, this test reports the number of times the Certifcate Revocation List (CRL) was updated successfully and the number of times the CRL update failed. This way, administrators can be proactively alerted to potential security threats (if any) and secure the NetScaler appliance from malicious attacks.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability

Parameter	Description
	• 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
Successful SSL handshakes	Indicates the number of SSL handshakes that were successful on the NetScaler appliance.	Number	
Failed SSL handshakes	Indicates the number of SSL handshakes that failed on the NetScaler appliance.	Number	Ideally, the value of this measure should be zero. A high value for this measure is a cause of concern as this may affect the communication between the server and client.
Imminent SSL certificate expiry	Indicates the number of SSL certificates that are about to expire.	Number	The detailed diagnosis of this measure if enabled, lists the SSL Certificate key pairs that are about to expire and the number of days for expiry.
Successful SSL CRL updates	Indicates the number of times the SSL Certificate Revocation List was updated successfully.	Number	From time to time, Certificate Authorities (CAs) issue certificate revocation lists (CRLs). CRLs contain information about certificates that can no longer be trusted. A certificate can be revoked if the private key is compromised or if that certificate expired and a new one is in use. A high value for this measure indicates that the CRLs are updated continuously which implies that the NetScaler device is highly secure.
Failed SSL CRL updates	Indicates the number of times the SSL Certificate Revocation List failed to update.	Number	Ideally, the value of this meaure should be zero. A high value for this measure indicates a serious threat to the security of the NetScaler device.

Measurement	Description	Measurement Unit	Interpretation
SSL VPN license limit reached	Indicates the number of times the SSL VPN license limit was reached.	Number	

3.6.8 SSL VPN Errors Test

The SSL VPN provides remote users access to authorized resources on a private intranet network, over a secure connection. The SSL VPN feature uses certain security policies that are enforced by the policy engine on the NetScaler appliance. If too many resource accesses (be it HTTP or non-HTTP) through SSL VPN are denied by the packet engine due to violation of the security policies, it indicates that the NetScaler appliance is highly prone to vulnerability which would eventually result in a poor performance show of the NetScaler. In order to closely monitor the performance of the NetScaler appliance, administrators should constantly keep a vigil on the errors that occur when resources are accessed through SSL VPN. The **SSL VPN Errors** test helps administrators in this regard. Using this test, administrators may be proactively alerted to the number of HTTP/non HTTP resource accesses denied by the policy engine and the number of times the Client Computer Security Check plug in for a SSL VPN failed to enforce a security policy.

For this test to run and report metrics, the NetScaler appliance should be configured to create a Syslog file in a remote Syslog server, where the details of all interactions with the NetScaler appliance will be logged. To know how to configure a remote Syslog server for the use of the NetScaler appliance, refer to Section 3.4.6.1.

This test is disabled by default. To enable the test, follow the Agents -> Tests -> Enable/Disable menu sequence in the eG administrative interface, pick *Citrix NetScaler VPX/MPX* as the **Component type**, select *Performance* as the **Test type**, choose this test from the list of **DISABLED TESTS** list, and click on the **<** button.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed
Host	The IP address of the host for which the test is being configured.
Port	The port at which the host listens. By default, this is *NULL*.
Log File Path	This test reports metrics by parsing a Syslog file. Specify the full path to the Syslog file here.
Search String	By default, the Syslog file may contain information relating to a number of servers that are inter linked with the target NetScaler appliance. In order to obtain the metrics of the target NetScaler appliance alone, specify the hostname or the IP address of the target NetScaler appliance for which the logs are to be read from the syslog file, in the Search String text box. Using this search string the information in the Syslog file may be parsed and metrics may be collected.
Search String Index	Here, specify the cursor position after which the eG agent should search for the specified Search String (or the position upto which the eG agent should ignore while searching for the specified Search String) in the syslog file. For example, if the specified Search String appears in the syslog file at the 17th position, then you may need to specify the Search String as 16.
DD Frequency	Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against DD Frequency.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Non HTTP resource access denied	Indicates the number of non-HTTP resource accesses that were denied by the policy engine.	Number	The Policy Engine (PE) provides a common framework for creating policy expressions that can be utilized by any of the features of the Citrix NetScaler Application Switch. The Policy Engine refers to the architecture in the Citrix NetScaler Application Switch for versions up to 8.x. The features that use policies are: • Load Balancing • Content Switching • Content Filtering • AppCompress • Cache Redirection • SSL VPN • Priority Queuing • DoS Protection • Sure Connect A Policy consists of an expression and an action. Expressions are “shared” among features on the switch. Actions are “feature-specific”. So we can create an expression to determine certain file types that are being processed by the NetScaler and as an action you can compress or optimize those files. The packet engine is created to perform TCP/IP processing, optimization tasks and acceleration of packages, next to

Measurement	Description	Measurement Unit	Interpretation
			this it enforces security policies too. This is a continuous process of grabbing packets, handling them accordingly and putting the packets in place again, the packet engine is designed to run an entire instance of NetScaler's packet engine on each processor core (nCore technology) and runs as a kernel component on the NetScaler. The Packet Processing Engine is responsible for all load balancing acceleration, server offload and security tasks. The detailed diagnosis of this measure if enabled lists the User, NAT IP, vServer, Source, Destination, the data sent, the data received and the policy that denied access to the non-HTTP resource.
HTTP resource access denied	Indicates the number of HTTP resource accesses that were denied by the policy engine.	Number	The detailed diagnosis of this measure if enabled lists the User, vServer, the data sent, the name of the remote host, the denied URL, and the policy that denied access to the HTTP resource.
Client security check for a SSL VPN fails	Indicates the number times the client computer security check for a SSL VPN failed.	Number	The SSL VPN administrator can configure the Client Computer Security Check plug-in to enforce a security policy on the client computer. A security policy is typically meant to ensure that security applications are installed and running. Security applications typically include personal firewalls, anti-virus packages, and customized applications or services. The plug-in performs a security check to ensure that the security policy is adhered to. These security checks can be performed once on login to the SSL VPN and also at

Measurement	Description	Measurement Unit	Interpretation
			periodic intervals during an active SSL VPN session as specified by the administrator.If a security check fails at any of these points, the plug-in will not be able to access the SSL VPN, even if successfully authenticated. If you are currently logged in and a security check fails, you will be disconnected from the SSL VPN. Frequent failures are a cause of concern and administrators should rectify such errors as soon as possible.
Client security expression evaluates to false	Indicates the number of times the client security check for a SSL VPN evaluated to false.	Number	Ideally, the value of this measure should be zero.

3.7 The Optimization Layer

By monitoring the effectiveness of the Integrated Cache and the Compression algorithm at work on the NetScaler appliance, the tests mapped to this layer provide useful pointers to how the cache usage and compression algorithms can be tweaked to optimize NetScaler performance.

Figure 3.14: The test mapped to the Optimization layer

3.7.1 Integrated Cache Test

The integrated cache provides in-memory storage on the Citrix NetScaler appliance and serves Web content to users without requiring a round trip to an origin server. The cache monitors HTTP and SQL requests that flow through the Citrix NetScaler appliance and compares the requests with stored policies. Depending on the outcome, the integrated cache feature either searches the cache for the response or forwards the request to the origin server.

For best performance, majority of requests for web content should be serviced by the cache, without disturbing the origin server. If the cache is not sized right however, this might not be possible! To

quickly detect irregularities in cache usage and sizing, administrators should know how much memory is used up by objects in cache, how much cache memory is free, and whether/not the cache is able to service requests efficiently with the memory available to it. This is exactly what the **Integrated Cache** test reveals!

This test monitors the requests to the cache, checks how well the cache processes the requests, and reveals whether cache misses are more than cache hits or vice versa. In the event that many requests are not serviced by the cache (i.e., if cache misses are more), administrators can use the memory usage statistics reported by this test to figure out if the cache is sized commensurate to its load.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Hits	Indicates the number of	Number	Ideally, the value of this measure should be high.

Measurement	Description	Measurement Unit	Interpretation
	requests serviced by the cache since the last measurement period.		
Misses	Indicates the number of requests serviced by origin server during the last measurement period.	Number	Ideally, the value of this measure should be 0 or much lower than the value of the Hits measure. A very high value is a cause for concern as it indicates that the cache is poorly utilized.
Requests	Indicates the total number of requests for web content that is passing through the NetScaler appliance.	Number	This is the sum of the values of the Hits and Misses measures.
Hit ratio	Indicates the percentage of requests that have been serviced by the cache.	Percent	Ideally, the value of this measure should be 80% and above. A value lesser than 80% indicates that the cache is unable to service many requests, maybe because the objects requested are not in cache. This could be due to poor cache size.
Origin bandwidth saved	Indicates the percentage of origin bandwidth saved, expressed as the ratio of the number of bytes served from the integrated cache divided by all bytes served.	Percent	The integrated cache saves bandwidth by servicing requests that would otherwise have fetched data from the origin server and consumed considerable bandwidth resources in the process. This is why, the value of this measure should be high, ideally. A very low value indicates excessive bandwidth consumption and abysmal bandwidth saving owing to the inability of the cache to service majority of the requests.
Non-304 hits	Indicates the total number of full (non-304) responses served by the cache during the last measurement	Number	

Measurement	Description	Measurement Unit	Interpretation
	period.		
304 hits	Indicates the total number of 304 responses served by the cache during the last measurement period.	Number	A 304 status code indicates that a response has not been modified since the last time it was served.
304 hit ratio	Indicates the percentage of 304 responses served by the cache during the last measurement period.	Percent	Ideally, the value of this measure should be close to 100%.
Data served by NetScaler	Indicates the total number of HTTP response bytes served by NetScaler from both the origin and the cache during the last measurement period.	MB	
Data served by cache	Indicates the total number of HTTP response bytes served by the cache since the last measurement period.	MB	
Data hit ratio	Indicates what percentage of the total data served by NetScaler has been served y the cache.	Percent	A value close to 100% is desired for this measure. If compression is On in the NetScaler, this ratio may not reflect the bytes served by the compression module. If the compression is Off, this ratio is the same the ratio of the Origin bandwidth saved.
Compressed data from cache	Indicates the number of compressed bytes served from the cache during the last measurement period.	MB	
Storable misses	Indicates the number of cache misses in the last measurement period, for	Number	A high value is desired for this measure, as it may reduce cache misses going forward.

Measurement	Description	Measurement Unit	Interpretation
	which the response fetched from the origin server was stored in the cache before serving to the client.		
Non-storable misses	Indicates the number of cache misses in the last measurement period, for which the response fetched from the origin server was not stored in the cache before serving to the client.	Number	Ideally, the value of this measure should be 0. This is because, non-storable misses are inappropriate for caching. By default, any response that contains the following status codes is a non-storable cache miss: 201, 202, 204, 205, 206 status codes All 4xx codes, except 403, 404 and 410 5xx status codes
Revalidations	Indicates the number of responses that an intervening cache revalidated with the integrated cache before serving to the client.	Number	Max-Age setting in a Cache-Control header determines, in number of seconds, when an intervening cache must revalidate the content with the integrated cache before serving it to the user. The value of this measure signifies the number of times the Max-Age setting was violated, causing revalidation to be attempted.
Successful revalidations	Indicates the number of successful revalidations in the last measurement period.	Percent	Ideally, the value of this measure should be the same as the value of the Revalidations measure.
Conversions to conditional requests	Indicates the number of user-agent requests for a cached Poll Every Time (PET) response that were sent to the origin server as conditional requests.	Number	You can configure the NetScaler appliance to always consult the origin server before serving a stored response. This is known as Poll Every Time (PET). When the NetScaler appliance consults the origin server and the PET response has not expired, a full response from

Measurement	Description	Measurement Unit	Interpretation
			the origin server does not overwrite cached content. This property is useful when serving client-specific content. After a PET response expires, the NetScaler appliance refreshes it when the first full response arrives from the origin server. A client issues a conditional request to ensure that the response that it has is the most recent copy. A user-agent request for a cached PET response is always converted to a conditional request and sent to the origin server. A conditional request has validators in If-Modified-Since or If-None-Match headers. The If-Modified-Since header contains the time from the Last-Modified header. An If-None-Match header contains the response's ETag header value. If the client's copy of the response is fresh, the origin server replies with 304 Not Modified. If the copy is stale, a conditional response generates a 200 OK that contains the entire response.
Storable miss ratio	Indicates the percentage of all cache misses, for which responses were fetched from the origin, stored in the cache, and then served to the client.	Percent	Higher the value of this measure, lesser will be the count of cache misses going forward. A value close to 100% will be ideal.
Successful revalidation ratio	Indicates the percentage of times stored content was successfully revalidated by a 304 (Object Not Modifed) response rather than by a	Percent	The value 100% will be ideal for this measure.

Measurement	Description	Measurement Unit	Interpretation
	full response.		
Expire at last byte	Indicates the number of times content expired immediately fter receiving the last body byte due to the Expire at Last Byte setting of the content group.	Number	
Flashcache misses	Indicates the number of requests to a content group with flash cache enabled that were not serviced by the cache.	Number	The phenomenon of Flash crowds occurs when a large number of clients access the same content. The result is a sudden surge in traffic toward the server. The Flash Cache feature enables the NetScaler appliance to improve performance in such situations by sending only one request to the server. All other requests are queued on the appliance and the single response is served to all of the requests. Ideally, the value of this measure should be 0.
Flashcache hits	Indicates the number of requests to a content group with flash cache enabled that were serviced by the cache.	Number	Ideally, the value of this measure should be high.
Parameterized invalidation requests	Indicates the number of requests in the last measurement period that match a policy with an invalidation (INVAL) action and a content group that uses an invalidation selector or parameters.	Number	Policies enable the integrated cache to determine whether to try to serve a response from the cache or the origin. Actions determine what the NetScaler appliance does when traffic matches a policy. Policies that you associate with the INVAL action for instance, immediately expire cached responses and refresh them from the origin server.

Measurement	Description	Measurement Unit	Interpretation
			An invalidation selector/parameter in a content group is a filter that locates particular objects in a content group that need to be expired. The value of this measure represents the number of requests that match an invalidation policy configured with an invalidation selector or parameter. A non-zero value for this measure indicates that one/more responses in the content group match the invalidation selector/parameter specification in the invalidation policy and have hence been expired.
Full invalidation requests	Indicates the number of requests in the last measurement period that match an invalidation policy where the invalGroups parameter is configured and expires one or more content groups.	Number	The invalGroups parameter in an INVAL policy configuration indicates the content groups to be invalidated/expired. A non-zero value for this measure therefore indicates that content groups configured for invalidation have been found and have been expired.
Invalidation requests	Indicates the requests that in the last measurement period match an invalidation policy and result in expiration of specific cached responses or entire content groups.	Number	A non-zero value for this measure indicates that content groups and/or one/more responses in content groups have been found to be match with configured invalidation policies and the matching groups/responses have been expired.
Parameterized requests	Indicates the number of cache requests in the last measurement period that were processed using a policy with a parameterized content group – i.e., a content group configured	Number	

Measurement	Description	Measurement Unit	Interpretation
	with hit and/or invalidation parameters/selectors.		
Parameterized non-304 hits	Indicates the number of cache requests in the last measurement period that were processed using a policy with a parameterized content group, where full cached response was found, and the response was not a 304 (object not updated) response.	Number	
Total parameterized hits	Indicates the number of cache requests in the last measurement period that were processed using a policy with a parameterized content group, where the cached object was found.	Number	
Parameterized 304 hit ratio	Indicates the percentage of 304 (object not updated) responses that were found using a parameterized policy, relative to all cache hits.	Percent	
Poll every time requests	Indicates the number of cache requests in the last measurement period that triggered a search for a content group that has Poll Every Time (PET) enabled.	Number	You can configure the NetScaler appliance to always consult the origin server before serving a stored response. This is known as Poll Every Time (PET). When the NetScaler appliance consults the origin server and the PET response has not expired, a full response from the origin server does not overwrite cached content. This property is useful when serving client-specific content.

Measurement	Description	Measurement Unit	Interpretation
			After a PET response expires, the NetScaler appliance refreshes it when the first full response arrives from the origin server.
Poll every time hits	Indicates the number of times a cache hit was found using the Poll Every Time method.	Number	
Poll every time hit ratio	Indicates the percentage of cache hits in content groups that have Poll Every Time enabled, relative to all searches of content groups with Poll Every Time enabled.	Number	A high value is desired for this measure.
Maximum memory	Indicates the largest amount of memory the NetScaler can dedicate to caching.	MB	Typically, upto 50% of memory available to the NetScaler appliance can be allocated for caching. If the global memory limit of the cache is set to 0, it indicates that the integrated cache feature is disabled.
Maximum memory active	Indicates the maximum amount of memory (active value) that will be set after the memory is actually allocated to the cache.	MB	Once the global memory limit is set, you can reset it to any other positive value later. However, this change will not alter the existing memory allocation to the Integrated cache immediately. For instance, you can change the global memory limit from 4000 MB to 6000 MB. But, if you query the memory limit soon after, you will find that while the Maximum memory measure reports the new memory limit of 6000 MB, the actual memory usage limit in effect currently continues to be 4000 MB and is by Maximum memory active measure. To ensure that the change to the global memory limit is effected, save

Measurement	Description	Measurement Unit	Interpretation
			the new configuration and restart the appliance. Once this is done, then before the Maximum memory and Maximum memory active measures will report the same value.
Utilized memory	Indicates the amount of memory the cache is currently using.	MB	If the value of this measure is very close to the Maximum memory active measure, it indicates that the cache is running out of memory and will not be able to hold any more requested objects. This will drastically reduce cache hits and increase cache misses. Also, accesses to the origin server will increase, thereby adversely impacting bandwidth usage. Under such circumstances, you may want to consider resizing the cache by allocating more memory to it.
Memory allocation failures	Indicates the total number of times in the last measurement period, the cache failed to allocate memory to store responses.	Number	Ideally, the value of this measure should be 0. A high value for this measure could indicate that the cache does not have adequate memory to allocate for responses. In such a situation, consider resizing the cache by allocating more memory to it.
Largest response so far	Indicates the the size (in bytes) of the largest response sent to client from the cache or the origin server.	MB	
Cached objects	Indicates the number of responses currently in the cache.	Number	One of the reasons for a memory drain on the cache is the presence of too many objects in the cache. Whenever Memory allocation failures are high, check the value of this measure to figure out if too many

Measurement	Description	Measurement Unit	Interpretation
			objects have been cached.
Marker objects	Indicates the number of marker objects in the marker cell.	Number	Marker objects are created when a response exceeds the maximum or minimum size for entries in its content group or has not yet received the minimum number of hits required for items in its content group. Though a marker object is cacheable, it will not be cached until it meets the configured maximum or minimum size specification or the minimum number of hits specification. If too many cache misses occur because of a marker object, you may want to change the Minimum Response Size, Maximum Response Size, and/or the MinHits for the content group.
Hits being served	Indicates the number of requests currently being served by the cache.	Number	
Misses being handled	Indicates the number of requests for which responses are fetched from the origin server but are currently being served by the cache.	Number	

3.7.2 Compression Test

The Citrix NetScaler appliance compression feature compresses the size of HTTP responses sent from servers to compression-aware browsers and thereby improves the performance of Web sites by reducing the download time of Web content. Bandwidth is also saved in the process. Another indirect benefit of HTTP compression is that the data passed between the Web server and the browser is encrypted by virtue of the compression algorithm, adding more security to the data.

If you enable the **Compression** feature, the Citrix NetScaler intercepts requests from clients and determines whether the client can accept compressed content. After receiving the HTTP response from the server, the NetScaler examines the content to determine whether it is compressible. If the content is compressible, the NetScaler compresses it, modifies the response header to indicate the type of compression performed, and forwards the compressed content to the client.

The real test of the effectiveness of the compression feature lies in how much data is compressed and how much bandwidth is saved in the process. Besides the compression algorithm employed, when compression occurs and what type of data is compressed also influence how compression is performed and how beneficial it really is. Moreover, errors encountered when decompressing data also serve as factors impacting user-perception on compression and its usefulness. Using the **Compression** test, administrators can accurately assess the effectiveness of the compression feature, keep a close watch on all factors influencing compression, and understand how these parameters can be fine-tuned to improve compression ratios and increase bandwidth savings.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
HTTP compression requests	Indicates the number of HTTP compression requests the NetScaler received in the last measurement period for which the response was successfully compressed.	Number	After you enable compression and configure services, if you send requests to the NetScaler with the following header information: Accept-Encoding: gzip, deflate, and NetScaler compresses the corresponding response, this counter is incremented.
HTTP compressible data received	Indicates the amount of HTTP data the NetScaler received in the last measurement period, which can be compressed.	MB	This measure reports the content length of the response that the NetScaler receives from server.
HTTP compressed data transmitted	Indicates the amount of HTTP data the NetScaler sent to the client in the last measurement period after compressing the response from the server.	MB	
HTTP compressible packets received	Indicates the number of HTTP packets received by NetScaler in the last measurement period, that can be compressed.	Number	
HTTP compressed packets transmitted	Indicates the number of HTTP packets the NetScaler sent to the client in the last measurement period after compressing the response from the server.	Number	
HTTP bandwidth saving	Indicates the percentage of bandwidth resources saved due to compression of HTTP traffic.	Percent	A high value is desired for this measure, as it indicates that HTTP compression has managed to save considerable bandwidth resources. This serves as adequate proof of the

Measurement	Description	Measurement Unit	Interpretation
			effectiveness of the HTTP compression algorithm. On the other hand, a very low value of this measure is indicative of ineffective compression of HTTP traffic.
HTTP compression ratio	Indicates the ratio of the compressible HTTP data received from the server to the compressed HTTP data sent to the client.	Percent	The higher the compression ratio, the greater the bandwidth savings! This is why, a value close to 100% for this measure is indicative of a very effective HTTP compression mechanism, which is bound to result in significant savings in bandwidth usage. On the other hand, a very low value for this measure indicates that only a small amount of the compressible HTTP data has been successfully compressed, thus resulting in poor bandwidth savings.
Total HTTP compression ratio	Indicates the ratio of total HTTP data received to total HTTP data transmitted.	Percent	The value of this measure indicates what percentage of HTTP data has been compressed. A high value is hence desired for this measure.
TCP compressible data received	Indicates the amount of TCP data the NetScaler received in the last measurement period, which can be compressed.	MB	This measure reports the content length of the TCP response that the NetScaler receives from server.
TCP compressed data transmitted	Indicates the amount of TCP data the NetScaler sent to the client in the last measurement period after compressing the response from the server.	MB	
TCP compressible packets received	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
	TCP packets received by NetScaler in the last measurement period that can be compressed.		
TCP compressed packets transmitted	Indicates the number of TCP packets the NetScaler sent to the client in the last measurement period after compressing the response from the server.	Number	
TCP bandwidth saving	Indicates the percentage of bandwidth resources saved due to compression of TCP traffic.	Percent	A high value is desired for this measure, as it indicates that TCP compression has managed to save considerable bandwidth resources. This serves as adequate proof of the effectiveness of the TCP compression algorithm. On the other hand, a very low value of this measure is indicative of ineffective compression of TCP traffic.
TCP compression ratio	Indicates the ratio of the compressible TCP data received from the server to the compressed TCP data sent to the client.	Percent	The higher the compression ratio, the greater the bandwidth savings! This is why, a value close to 100% for this measure is indicative of a very effective TCP compression mechanism, which is bound to result in significant savings in bandwidth usage. On the other hand, a very low value for this measure indicates that only a small amount of the compressible TCP data has been successfully compressed, thus resulting in poor bandwidth savings.
Quantum	Indicates the number of	Number	The NetScaler appliance buffers the

Measurement	Description	Measurement Unit	Interpretation
compression	times NetScaler compresses a quantum of data.		data received from the server until it reaches a configured quantum size and then compresses the buffered data and transmits it to the client. The minimum value of the quantumSize configuration is 1 KB, the maximum is 63488 KB and the default value is 57344 KB. Since compression is a CPU-intensive operation, setting a quantumSize ensures that compression does not take place every time a response packet comes in, and instead occurs only when the response packets are of the quantumSize configured. A high value for this measure could indicate a very low quantumSize configuration, which is causing compression to occur frequently. This in turn, can put a strain on the CPU resources, if quantumSize is not increased soon.
Push flag compression	Indicates the number of times NetScaler compressed data in the last measurement period, upon receiving a TCP PUSH flag from the server.	Number	The PUSH flag ensures that data is compressed immediately without waiting for the buffered data size to reach the quantumSize. If the compression occurs frequently, despite setting a high quantumSize, you may want to check the PUSH flag status. If this flag is enabled, it is a clear indicator that the quantumSize has been ignored. To reduce the frequency of compression and to conserve CPU resources, you may want to keep the PUSH flag disabled.

Measurement	Description	Measurement Unit	Interpretation
End of input compression	Indicates the number of times NetScaler compressed data in the last measurement period, upon receiving End of Input (FIN packet).	Number	When the NetScaler receives End Of Input (FIN packet), it compresses the buffered data immediately without waiting for the buffered data size to reach the quantum size. If the compression occurs frequently, despite setting a high quantumSize and disabling the PUSH flag, it could be because the the End of input packet has been received.
Timer compression	Indicates the number of times in the last measurement period, NetScaler compressed data because the data accumulation timer expired.	Number	The timer expires if the server response is very slow and consequently, the NetScaler does not receive response for a certain amount of time. Under such a condition, the NetScaler compresses the buffered data immediately without waiting for the buffered data size to reach the quantum size. If frequent timer expiry is causing compression to occur frequently, you may want to reset the timer to wait for a longer time before compressing buffered data.
Compressed data received	Indicates the total amount of compressed data received by NetScaler in the last measurement period.	MB	
Decompressed data transmitted	Indicates the total amount of decompressed data transmitted by NetScaler during the last measurement period.	MB	
Compressed packets received	Indicates the number of compressed packets received by NetScaler	Number	

Measurement	Description	Measurement Unit	Interpretation
	during the last measurement period.		
Decompressed packets transmitted	Indicates the number of decompressed packets transmitted by NetScaler during the last measurement period.	Number	
Decompression ratio	Indicates the ratio of decompressed data transmitted to compressed data received.	Percent	A high value is indicative of an efficient decompressing mechanism.
Wrong data	Indicates the number of data errors encountered when decompressing during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Less data	Indicates the number of times in the last measurement period, NetScaler received less data than declared by the protocol.	Number	
More data	Indicates the number of times in the last measurement period, NetScaler received more data than declared by the protocol.	Number	
Memory failures	Indicates the number of times memory failures occurred when decompressing, during the last measurement period.	Number	Ideally, the value of this measure should be 0.
Unknown	Indicates the number of times unknown errors occurred while decompressing, during the last measurement period.	Number	Ideally, the value of this measure should be 0.

3.8 The Load Balancing Layer

The tests mapped to this layer monitor how the NetScaler performs load-balancing and reveals irregularities in the process.

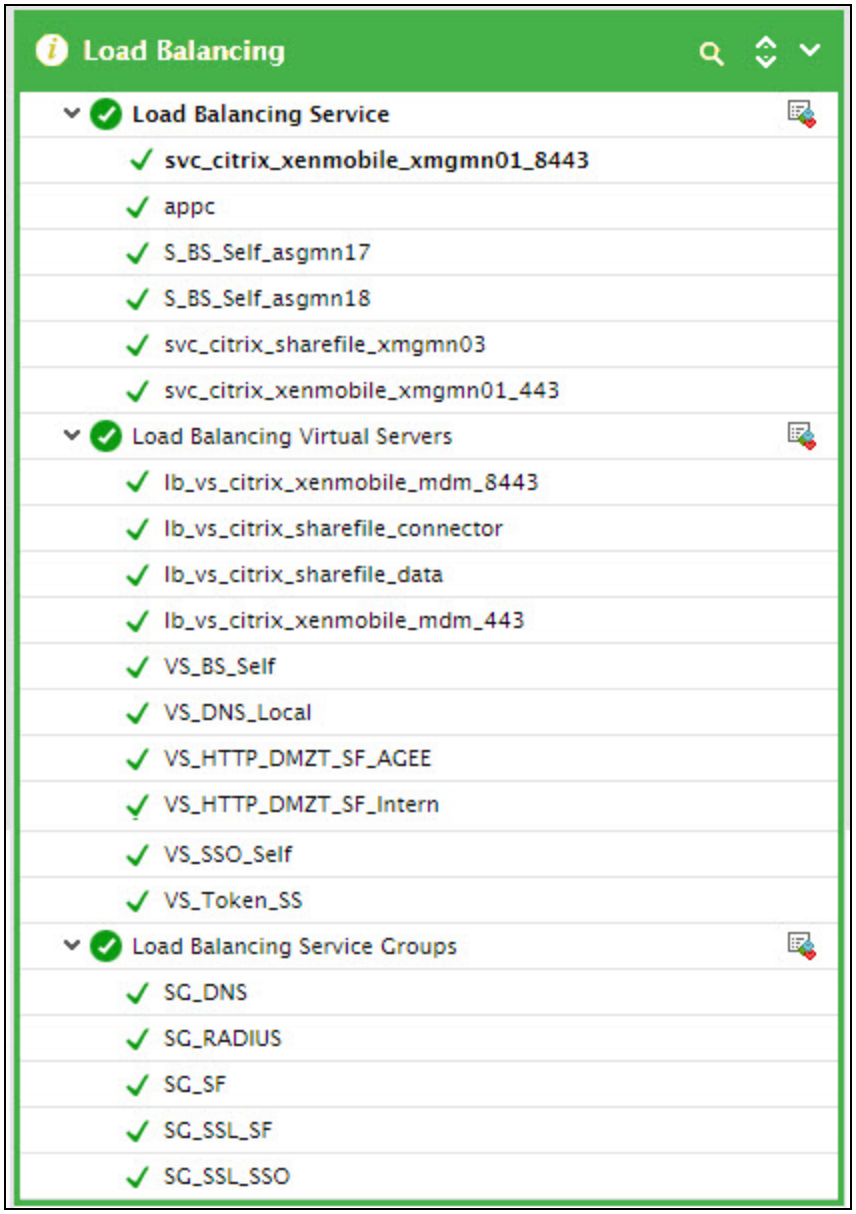

Figure 3.15: The tests mapped to the Load Balancing layer

3.8.1 Load Balancing Service Test

A service is a NetScaler entity that represents applications on a server. While services are normally combined with virtual servers, in the absence of a virtual server, a service can still manage

application-specific traffic.

This test reports the current state of the virtual server for each service and also reports the load on each service along with critical metrics that provide detailed information on the connections handled by each service and the data transfer happening on each service.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each service in each service group configured on the NetScaler appliance being monitored

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **NETSCALER USERNAME** and **NETSCALER PASSWORD** - To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the **NETSCALER USERNAME** and **NETSCALER PASSWORD** text boxes.
4. **CONFIRM PASSWORD** - Confirm the **NETSCALER PASSWORD** by retyping it here.
5. **SSL** - The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the **SSL** flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the **SSL** flag to **No**.
6. **SHOW UP SERVER ONLY** – The default setting of this flag is **No**; this indicates that this test, by default, monitors all the services configured in a NetScaler device. If you want the test to monitor only those services that are up and running currently, then set this value to **Yes**.
7. 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
Server state:	Indicates the current state of the virtual server bound to this service in the NetScaler device.		The values that this measure can report and their numeric equivalents are listed in the table below: Numeric Value / Measure Value 0 – Up 1 – Down 2 – Out of service 3 – Transition out of service 4 – Down when going out of service -1 – Unknown **Note:** By default, this measure reports the above-mentioned Measure Values to indicate the state of the virtual server. However, in the graph of this measure the virtual server state will be represented using the corresponding numeric equivalents only.

Measurement	Description	Measurement Unit	Interpretation
			The detailed diagnosis of this measure shows the Service Type, Primary Port and the Primary IP address of the virtual server.
Server connections:	Indicates the number of current connections to the actual servers behind the virtual server bound to this service.	Number	These measures serve as good indicators of the current workload of a service.
Client connections:	Indicates the number of current client connections to this service.	Number	
Active connections:	Indicates the number of active transactions (including those in the surge queue) handled by this service.	Number	
Load on service:	Indicates the load on this service.	Number	NetScaler uses a special type of monitor known as a Load Monitor to calculate the load on each service in the network. Load calculation enables the NetScaler Load Balancing engine to make load balancing decisions and distribute the traffic appropriately between the services. This measure reports the service load as calculated by the Load Monitor. By comparing the value of this measure across services, you can quickly determine whether/not the load is balanced across services, isolate the overloaded

Measurement	Description	Measurement Unit	Interpretation
			services, and initiate measures to fix the load-balancing irregularities.
Requests in idle queue/reuse pool:	Indicates the number of requests in the idle queue or the reuse pool.	Number	Ideally, the value of this measure should be low. A very high value indicates that too many idle connections have been placed in the idle queue/reuse pool and have not been reused yet. Such requests drain resources unnecessarily. To avoid this, you can limit the number of connections that can be added to the connection reuse pool of the appliance. You can add an HTTP profile and attach it at a virtual server (VServer) or a service level. You can use this profile to limit the number of connections that can be added to the connection reuse pool.
Avg TTFB between NetScaler and server:	Indicates the average response Time to First Byte (TTFB) to the NetScaler device from the server.	Secs	Time to First Byte is a metric used to measure response time, calculated based on the time it takes to send a GET Request and receive a GET Response back from a server. A high value is a clear indication of network congestion or connection failure.
Maximum open connections allowed:	Indicates the maximum number of open connections that are allowed on this service.	Number	
Surge queue:	Indicates the number of requests in the surge queue of this service.	Number	The NetScaler device can be used to limit the number of simultaneous

Measurement	Description	Measurement Unit	Interpretation
			requests that are passed on to a server. When a request is completed, additional requests are forwarded to the server. If a request arrives and the server is handling the maximum configured number of requests, the NetScaler device places the new request in a surge queue, where the request waits for its turn to be sent to the server for processing. The surge queue allows a server to run at peak capacity without the risk of having it spiral out of control because of a surge of incoming requests. If the surge queue is consistently greater than 0, this indicates that the server is not able to keep up with the workload and additional server capacity is required. On the other hand, a periodic surge is not a cause for concern. When a surge in client requests overloads a server, server response becomes slow, and the server is unable to respond to new requests. The Surge Protection feature ensures that connections to the server occur at a rate that the server can handle. The response rate depends on how surge protection is configured. The NetScaler appliance also tracks the number of connections to the

Measurement	Description	Measurement Unit	Interpretation
			server, and uses that information to adjust the rate at which it opens new server connections.
Server connections in established state:	Indicates the number of connections to the virtual server bound to this service that are currently in ESTABLISHED state.	Number	
Request data received:	Indicates the amount of data received as requests on this service or virtual server during the last measurement period.	MB	
Response data received:	Indicates the amount of data received as responses on this service or virtual server during the last measurement period.	MB	
Requests received:	Indicates the number of requests received on this service or virtual server during the last measurement period.	Number	
Responses received:	Indicates the number of responses received on this service or virtual server during the last measurement period.	Number	
Packets received:	Indicates the number of packets received on this service or virtual server during the last	Number	

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		
Packets sent:	Indicates the number of packets sent from this service or virtual server during the last measurement period.	Number	
Throughput:	Indicates the rate at which the data is received from and sent by this service or virtual server during the last measurement period.	Mbps	Comparing the throughput across services will enable you to instantly isolate the service that is consuming bandwidth excessively.
Service hits:	Indicates the number of times this service has been provided during the last measurement period.	Number	

3.8.2 Load Balancing Virtual Servers Test

The load balancing feature is a core feature of the NetScaler appliance. The load balancing feature distributes user requests for Web site pages and other protected applications across multiple servers that all host (or mirror) the same content.

A load balancing setup includes a load-balancing virtual server and multiple load-balanced application servers. The virtual server receives incoming client requests, uses the load balancing algorithm to select an application server, and forwards the requests to the selected application server. The following conceptual drawing illustrates a typical load balancing deployment.

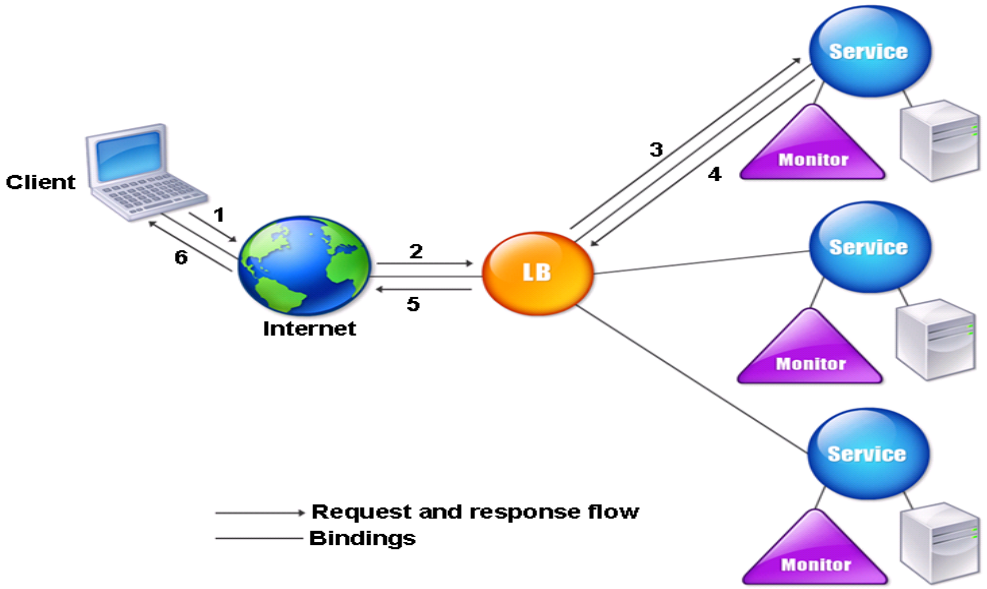

Figure 3.16: Load balancing architecture

The entities that you configure in a typical NetScaler load balancing setup are:

- **Load balancing virtual server**. The IP address, port, and protocol combination to which a client sends connection requests for a particular load-balanced website or application. If the application is accessible from the Internet, the virtual server IP address (VIP) is a public IP address. If the application is accessible only from the local area network (LAN) or wide area network (WAN), the VIP is usually a private (ICANN non-routable) IP address.
- **Service**. The IP address, port, and protocol combination used to route requests to a specific load-balanced application server. A service can be a logical representation of the application server itself, or of an application running on a server that hosts multiple applications. Each service is bound to a specific virtual server.
- **Server object.** An entity that identifies a physical server and provides the server's IP address. If you want to use the server's IP address as the name of the server object, you can enter the server's IP address when you create a service, and the server object is then created automatically. Alternatively, you can create the server object first and assign it an FQDN or other name, and then specify that name instead of the IP address when you create the service.
- **Monitor**. An entity on the NetScaler appliance that tracks a service and ensures that it is operating correctly. The monitor periodically probes (or performs a health check on) each service to which you assign it. If the service does not respond within the time specified by the time-out, and a specified number of health checks fail, that service is marked DOWN. The NetScaler

appliance then skips that service when performing load balancing, until the issues that caused the service to quit responding are fixed.

The load balancing virtual server can use any of a number of algorithms (or methods) to determine how to distribute load among the load-balanced servers that it manages. The default load balancing method is the least connection method, in which the NetScaler appliance forwards each incoming client connection to whichever load-balanced application server currently has the fewest active user connections.

Since load balancing manages user requests to heavily used applications, it prevents poor performance and outages. Irregularities in load-balancing can hence cause significant delays in request processing, thus adversely impacting the user experience with a load-balanced application. To avoid this, you need to configure the periodic execution of the **Load Balancing Virtual Servers** test. For each virtual server configured on the NetScaler appliance, this test does the following:

- Verifies and promptly reports the non-availability / abnormal state of the virtual servers;
- Continuously monitors the load on the load-balancing virtual servers and reveals how well each server processes client requests;
- Detects inconsistencies in load-balancing early on;
- Warns administrators of these deviations proactively;
- Helps initiate changes in the load-balancing algorithm (if required);

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each load balancing virtual server configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.

Parameter	Description
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Up Server Only	The default setting of this flag is **No**; this indicates that this test, by default, monitors all the services configured in a NetScaler device. If you want the test to monitor only those services that are up and running currently, then set this value to **Yes**.
Exclude Servers	Provide a comma-separated list of virtual server names or name patterns that need to be excluded from monitoring. By default, this is set to *none*, indicating that all virtual servers are by default monitored.
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
Server state	Indicates the current state of this virtual server.		The values reported by this measure and their numeric equivalents are as shown in the table:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Up</td></tr>
<tr><td>1</td><td>Down</td></tr>
<tr><td>2</td><td>Out of service</td></tr>
<tr><td>3</td><td>Transition out of service</td></tr>
<tr><td>4</td><td>Down when going out of service</td></tr>
<tr><td>-1</td><td>Unknown</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating the virtual server state. However, in the graph of this, the states will be represented using the corresponding numeric equivalents only.
The detailed diagnosis of this measure shows the Service Type, Primary Port and the Primary IP address of the virtual server.
</td></tr>
<tr><td>Virtual server health status</td><td>Indicates the current health of this virtual server.</td><td>Percent</td><td>A high value is desired for this measure, as low values are indicative of unhealthy state.
Use the detailed diagnosis of this measure to know which service is bounding with the virtual server.</td></tr>
<tr><td>Client connections</td><td>Indicates the current number of client connections to this virtual server.</td><td>Number</td><td>This is a good indicator of the load on the virtual server. Compare the values of this measure across virtual servers to determine which virtual server is heavily loaded currently.</td></tr>
<tr><td>Server connections</td><td>Indicates the number of</td><td>Number</td><td></td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	current connections to the load-balanced application servers behind this virtual server.		
Client connections in established state	Indicates the number of client connections that were in the *ESTABLISHED* state during the last measurement period.	Number	
Server connections in established state	Indicates the number of server connections that were in the *ESTABLISHED* state during the last measurement period.	Number	
Data received	Indicates the amount of request data received on this service or virtual server during the last measurement period.	MB	These measures serve as effective indicators of data/packet load on a virtual server.
Data transmitted	Indicates the amount of response data transmitted by this service or virtual server during the last measurement period.	MB	
Packets received:	Indicates the number of packets received on this service or virtual server during the last measurement period.	Number	
Packets sent	Indicates the number of packets sent by this service or virtual server during the last measurement period.	Number	

Measurement	Description	Measurement Unit	Interpretation
Requests received:	Indicates the number of requests received on this service or virtual server during the last measurement period.	Number	If the number of Requests received is a lot higher than the Responses received for a virtual server, it indicates that many requests are still pending processing on the virtual server; one of the reasons for latencies in request processing is an inefficient load-balancing algorithm.
Responses received	Indicates the amount of responses received by this service or virtual server during the last measurement period.	Number	
Spill over threshold	Indicates the spillover threshold that was set on the virtual server during the last measurement period.	Number	Spillover is a means to divert new connections to a vserver to a backup vserver when the number of connections to the vserver exceeds the threshold value. Spillover can either be connection-based or dynamic.
Virtual server experienced spill over	Indicates the number of times the spill over threshold was exceeded by this virtual server during the last measurement period.	Number	A connection overload can cause a spillover. A high value of this measure indicates that the virtual server was often overloaded with connections. This can in turn be caused by inefficient load-balancing by the virtual server.
Deferred requests	Indicates the number of deferred requests received on this virtual server during the last measurement period.	Number	
Labeled connections	Indicates the number of client connections to this virtual server that were labeled during the last measurement period.	Number	The NetScaler Web 2.0 push feature enables the server to label a client connection and subsequently identify and send data over that labeled connection. With NetScaler Web 2.0 push enabled, the client first

Measurement	Description	Measurement Unit	Interpretation
			establishes a TCP/IP connection and connects to the NetScaler appliance. The appliance uses the configured load balancing method or content switching policy to select a Web server (referred to as a notification server) to which the request is to be forwarded. The appliance then initiates the labeling protocol with the Web server. This protocol enables the Web server to label the connection and defer the response. The protocol also enables the server to process other requests without invoking push processing. Upon receipt of the deferred response from the Web server, the appliance starts waiting for updates from the Web server. When updates become available, the Web server uses the message push protocol to push the updates/messages to a push virtual server. A push virtual server is a load balancing virtual server with service type PUSH or SSL_PUSH. The appliance then processes updates/messages and uses the label to 'push' the updates to the client. This way, the NetScaler Web 2.0 push feature helps reduce the frequent polling of the Web server for updates, and thus minimizes the load on the server. From a load-balancing perspective therefore, a large number of Labelled connections and Push labels are desired.
Push labels	Indicates the number of labels for this push virtual server during the last measurement period.	Number	
Virtual server hits	Indicates the number of virtual server hits during	Number	

Measurement	Description	Measurement Unit	Interpretation
	the last measurement period.		

3.8.3 Load Balancing Service Groups Test

A service group is a representation of one/more services. For each service in a service group, this test reports the state of the virtual server bound to the service, tracks the packet/data traffic flowing into and out of each service, and thus reveals how uniformly load is distributed across all the services in a group.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each service in each service group configured on the NetScaler appliance being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
Confirm Password	Confirm the NetScaler Password by retyping it here.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Up Server Only	The default setting of this flag is **No**; this indicates that this test, by default, monitors all the services configured in a NetScaler device. If you want the test to monitor only

Parameter	Description
	those services that are up and running currently, then set this value to **Yes**.
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
Server state	Indicates the current state of the virtual server bound to this service in the NetScaler device.		The values that this measure can report and their numeric equivalents are listed in the table below: **Measure Value** / **Numeric Value** Up – 0 Down – 1 Out of service – 2 Transition out of service – 3 Down when going out of service – 4 **Note:** By default, this measure reports the above-mentioned **Measure Value**s to indicate the state of the virtual server. However, in the graph of this measure

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>the virtual server state will be represented using the corresponding numeric equivalents only.

The detailed diagnosis of this measure shows the Service Type, Primary Port, the Primary IP address, Weight (The higher the weight, the higher the percentage of requests sent to each service), Server ID and Hash ID of the virtual server.</td></tr>
<tr><td>Effective state</td><td>Indicates the effective state of virtual server that are bound to this service.</td><td></td><td>Effective state of a service is the state derived from the Load Balancing virtual server bound to that service in each group.

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

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>100</td></tr><tr><td>Partial-Up</td><td>75</td></tr><tr><td>Out of Service</td><td>50</td></tr><tr><td>Partial-Down</td><td>25</td></tr><tr><td>Down</td><td>10</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values to indicate the effective state of the virtual server. However, in the graph of this measure the effective state will be represented using the corresponding numeric equivalents only.</td></tr>
<tr><td>Is server enabled</td><td>Indicates whether the</td><td></td><td>The state of a service is the state</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>state?</td><td>virtual server bound to this service is enabled or not.</td><td></td><td>derived from the monitors bound to that service in each service group.

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

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

By default, this measure reports the above-mentioned Measure Values to indicate the state of the virtual server bound to each service in each service group. However, in the graph of this measure the service group state will be represented using the corresponding numeric equivalents only - i.e., 0 or 1.</td></tr>
<tr><td>Available servers</td><td>Indicates the number of virtual servers that are bound to this service.</td><td>Number</td><td>This measure will report a value only for the Summary descriptor.</td></tr>
<tr><td>Servers in UP state</td><td>Indicates the number of virtual servers that are in the UP state.</td><td>Number</td><td>This measure will report a value only for the Summary descriptor.</td></tr>
<tr><td>Percentage of UP state servers</td><td>Indicates the percentage of virtual servers that are in the UP state.</td><td>Percent</td><td>This measure will report a value only for the Summary descriptor.</td></tr>
<tr><td>Server connections</td><td>Indicates the number of current connections to the actual servers behind the virtual server bound to this service.</td><td>Number</td><td rowspan="2">These measures serve as good indicators of the current workload of a service.</td></tr>
<tr><td>Client connections</td><td>Indicates the number of current client connections</td><td>Number</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	to this service.		
Requests in idle queue/reuse pool	Indicates the number of requests in the idle queue or the reuse pool.	Number	Ideally, the value of this measure should be low. A very high value indicates that too many idle connections have been placed in the idle queue/reuse pool and have not been reused yet. Such requests drain resources unnecessarily. To avoid this, you can limit the number of connections that can be added to the connection reuse pool of the appliance. You can add an HTTP profile and attach it at a virtual server (VServer) or a service level. You can use this profile to limit the number of connections that can be added to the connection reuse pool.
Avg TTFB between NetScaler and server	Indicates the average response Time to First Byte (TTFB) to the NetScaler device from the server.	Seconds	Time to First Byte is a metric used to measure response time, calculated based on the time it takes to send a GET Request and receive a GET Response back from a server. A high value is a clear indication of network congestion or connection failure.
Maximum open connections allowed	Indicates the maximum number of open connections that are allowed on this service.	Number	
Surge queue	Indicates the number of requests in the surge queue of this service.	Number	The NetScaler device can be used to limit the number of simultaneous requests that are passed on to a server. When a request is completed, additional requests are forwarded to the server. If a request arrives and the server is handling the maximum configured number of requests, the NetScaler device places the new request in a surge queue, where the

Measurement	Description	Measurement Unit	Interpretation
			request waits for its turn to be sent to the server for processing. The surge queue allows a server to run at peak capacity without the risk of having it spiral out of control because of a surge of incoming requests.
Server connections in established state	Indicates the number of connections to the virtual server bound to this service that are currently in *ESTABLISHED* state.	Number	
Request data received	Indicates the amount of request data received on this service group or virtual server during the last measurement period.	MB	
Response data received	Indicates the amount of response data received on this service group or virtual server during the last measurement period.	Number	
Request packets	Indicates the number of packets received on this service group or virtual server during the last measurement period.	Number	
Response packets	Indicates the number of packets sent from this service group or virtual server during the last measurement period.	Number	

3.8.4 GSLB Domains Test

GSLB (Global Server Load Balancing) is a Domain Name Server (DNS)-based solution that load balances services between geographically distributed locations. GSLB relies on DNS for directing client requests. GSLB enables the NetScaler appliance to make intelligent network traffic direction

decisions based on the configured method. GSLB responds to DNS requests for a domain name with an IP address of a member service. Which service IP is returned is dependent on the load-balancing algorithm used - for example least connection, simple round robin or more commonly used, proximity to the client (or the clients local DNS to be precise).

To understand how efficiently a NetScaler appliance configured with GSLB performs load-balancing, you first need to understand the DNS request load on the appliance and the most-requested domain names. The **GSLB Domains** test provides this information. This test tracks the DNS requests received by the NetScaler appliance, automatically discovers the domain names that are requested, and reports the number of requests received for each domain name. Besides enabling you to instantly capture a sudden or a steady increase in load, this test also helps you point to those domains that are requested the maximum. This test can also provide provide early pointers to what could be a potential request processing bottleneck on the NetScaler appliance.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each domain that is requested for.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
DNS queries received	Indicates the number of DNS requests received for this domain name during the last measurement period.	Number	A high value for this measure could indicate an increase in the load on the NetScaler appliance. Under such circumstances, you can quickly compare the value of this measure across domains to identify the domain that has received the maximum requests. You may want to observe the traffic to this domain for a while to know whether there was a consistent increase in the load or was it just a sudden occurrence. A consistent rise in the number of requests could indicate a potential processing bottleneck.

3.8.5 GSLB Sites Test

A typical GSLB deployment contains the entities described in the following figure:

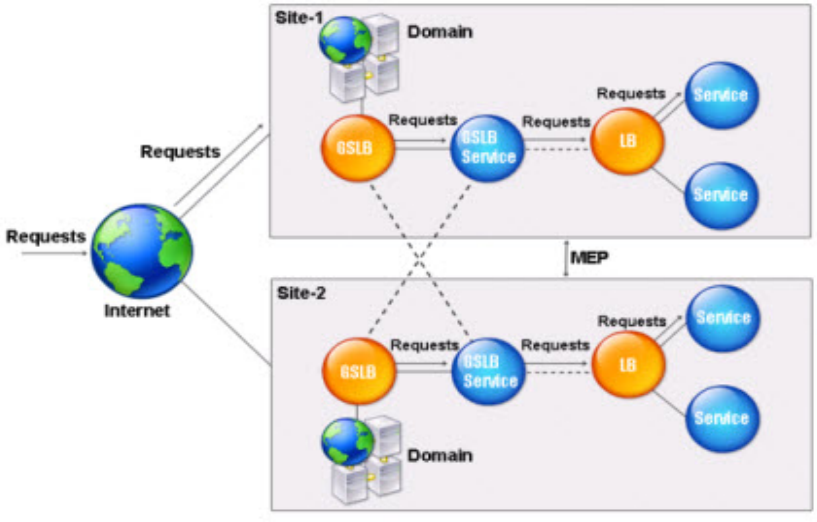

Figure 3.17: How GSLB works?

To configure GSLB, you must configure a GSLB site. As shown in the figure, a GSLB site is the logical collection of GSLB vserver, GSLB service, LB vserver, service, domain, and ADNS service. It is the central entity in a GSLB deployment, and is represented by a name and an IP address. To create a GSLB site, you must configure load balancing on the system. You must create GSLB vservers and GSLB services for each site. You must bind GSLB services to GSLB vservers.

Once multiple such GSLB sites are configured, you need to define load-balancing policies and methods, and metric exchange policies per site. GSLB methods are algorithms that control how the system load-balances client requests across distributed data centers. GSLB policies direct the traffic to a pre-defined target site. Multiple sites exchange metrics with each other using the Metric Exchange Protocol (MEP). The system uses this protocol to exchange load, network, and persistence information between GSLB sites. Once all the above are configured, GSLB enables the uniform distribution of traffic across these sites, manages disaster recovery, and ensures that applications are consistently accessible.

By continuously observing the DNS requests and responses received by each site configured on a NetScaler appliance, you can easily measure the efficiency of GSLB and the effectiveness of its policies and methods. The **GSLB Sites** test enables this analysis. This test reports the number of requests and responses received by each GSLB site that is configured on a target NetScaler appliance, and in the process, points to overloaded sites (if any). This way, the test reveals irregularities in load-balancing and helps determine whether it is because of improper metric exchange between sites.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each GSLB site.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface

Parameter	Description
	Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Detailed Diagnosis	To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option. The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled: • The eG manager license should allow the detailed diagnosis capability • Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current client connections	Indicates the number of current client connections to the virtual servers represented by all GSLB services associated with this GSLB site.	Number	These are good measures of the connection load on a site. By comparing the value of each of these measures across sites, you can instantly identify overloaded sites and promptly detect irregularities in load-balancing. You can use the detailed diagnosis of the Current client connections measure to figure out the site type, the IP address of the site, and the public IP address of the site.
Current server connections	Indicates the number of current connections to the real servers behind the virtual servers represented by all GSLB services associated with this GSLB site.	Number	
Request data received	Indicates the amount of request data received by	MB	These are good measures of the request and response load on a site. By

Measurement	Description	Measurement Unit	Interpretation
	the virtual servers represented by all GSLB services associated with this GSLB site during the last measurement period.		comparing the value of each of these measures across sites, you can instantly identify overloaded sites and promptly detect irregularities in load-balancing.
Response data received	Indicates the amount of response data received by the virtual servers represented by all GSLB services associated with this GSLB site during the last measurement period.	MB	In the event that such irregularities come to light, you may want to consider fine-tuning the GSLB policies and/or GSLB methods and/or metric exchange policies supported by the system to ensure that the appliance takes intelligent load-balancing decisions.
Requests received	Indicates the number of requests received by the virtual servers represented by all GSLB services associated with this GSLB site during the last measurement period.	Number	
Responses received:	Indicates the number of requests received by the virtual servers represented by all GSLB services associated with this GSLB site during the last measurement period.	Number	
Network metric exchange connection status	Indicates the current status of the network Metric Exchange connection at this GSLB site.		MEP (Metric exchange protocol) is a proprietary protocol used by the NetScaler appliances to exchange site metrics, network metrics, and persistence information to other sites participating in GSLB. Network metric exchange refers to the LDNS RTT information exchange, which is used in the dynamic proximity load balancing algorithm. Network proximity is a measure of how far a user is located

Measurement	Description	Measurement Unit	Interpretation
			from a data resource in terms of network distance or time. The GSLB feature monitors the real-time status of the network and directs the request of the client to the best site. GSLB uses Round Trip Time (RTT) metric to measure network proximity. The RTT between the Local DNS (LDNS) of the client and each of the GSLB sites is measured. The appliance then uses this metric to make the load balancing decision. The LDNS RTT information is exchanged every five seconds. This is a push exchange model. Every five seconds, each site sends its data to other participating sites. A connection for exchanging these network metrics can be initiated by either site that is involved in the exchange. If the connection at a GSLB site is down, then, this measure will report the value Down. Network metrics will no longer be exchanged between sites if the connection is down. In the absence of LDNS RTT information, intelligent load-balancing decisions can not be taken. To avoid this, ensure that the correct port 3011 or 3009 (if secure) must be open on the firewall (if any) between the appliances. Likewise, make sure that the public IP address of the site is allowed on the blocking firewall (if any). On the other hand, if the connection for exchanging network metrics is alive, then the value of this measure will be Up. The numeric values that

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>correspond to these measure values have been listed in the table below:
<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>0</td></tr><tr><td>Down</td><td>1</td></tr></table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether the exchange connection is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Site metric exchange connection status</td><td>Indicates the current status of the site Metric Exchange connection at this GSLB site.</td><td></td><td>Site metric exchange is a polling exchange model. The GSLB site metric exchange interval is one second. Site information consists of the current number of connections and current packet rate for a load balancing virtual server. For example if a site has a configuration for the services of another site, then after every one second, the first site requests the other site for the status of the GSLB services. The other site responds with the state and the other load details.
A connection for exchanging site metrics is always initiated by the site with the lower IP address. By default, the site uses a subnet IP address (SNIP) or a mapped IP address (MIP) to establish a connection to the IP address of a different site. However, you can configure a specific SNIP,</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>MIP, the NetScaler IP address (NSIP), or a virtual IP address (VIP) as the source IP address for metrics exchange.

If the connection for exchanging site metrics is not available, then this measure will report the value Down. On the other hand, if the connection is alive, then this measure will report the value Up. The numeric values that correspond to these measure values have been listed in the table below:

<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>0</td></tr><tr><td>Down</td><td>1</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether the exchange connection is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Metric exchange policy status</td><td>Indicates the current status of the metric exchange policy at this site.</td><td></td><td>MEP is a proprietary protocol used by the NetScaler appliances to exchange site metrics, network metrics, and persistence information to other sites participating in GSLB. To use the load balancing method, as defined on the load balancing virtual servers as a means to resolution, it is necessary to enable MEP on all GSLB sites. If MEP is enabled on a site, then the value of this measure will be Up. If MEP is</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>disabled, then the value of this measure will be Down. In this case, the entire load balancing method as defined by the virtual server is disabled and falls back to a Round Robin load balancing method.

The numeric values that correspond to these measure values have been listed in the table below:

<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>0</td></tr><tr><td>Down</td><td>1</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether the MEP status is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Is network metric exchange enabled?</td><td>Indicates whether network metric exchange is currently enabled or disabled at this GSLB site.</td><td></td><td>Network metric exchange refers to the LDNS RTT information exchange, which is used in the dynamic proximity load balancing algorithm. Network proximity is a measure of how far a user is located from a data resource in terms of network distance or time. The GSLB feature monitors the real-time status of the network and directs the request of the client to the best site. GSLB uses Round Trip Time (RTT) metric to measure network proximity. The RTT between the Local DNS (LDNS) of the client and each of the</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>GSLB sites is measured. The appliance then uses this metric to make the load balancing decision.

If the exchange of network metrics is disabled for a site, then this measure will report the value Down. In this case, the appliance will not be able to compute RTT between the Local DNS of the client and the GSLB site, and will hence be unable to take accurate load-balancing decisions. On the other hand, if network metric exchange is enabled on a site, then this measure will report the value Up.

The numeric values that correspond to these measure values have been listed in the table below:

<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>0</td></tr><tr><td>Down</td><td>1</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether the network metric exchange status is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Is metric exchange enabled?</td><td>Indicates whether site metric exchange is enabled or disabled at this GSLB site.</td><td></td><td>Site metric exchange is a polling exchange model. The GSLB site metric exchange interval is one second. Site information consists of the current number of connections and current</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>packet rate for a load balancing virtual server. For example if a site has a configuration for the services of another site, then after every one second, the first site requests the other site for the status of the GSLB services. The other site responds with the state and the other load details.

This measure reports the value Up if site metric exchange is currently enabled for a site. On the other hand, the measure reports the value Down if site metric exchange for a site is disabled.

The numeric values that correspond to these measure values have been listed in the table below:

<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>0</td></tr><tr><td>Down</td><td>1</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether metric exchange is enabled / disabled. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
<tr><td>Is persistence entry exchange enabled?</td><td>Indicates whether persistence entries exchange is enabled or disabled at this GSLB site.</td><td></td><td>Persistence ensures that a series of client requests for a particular domain name is sent to the same data center instead of being load balanced. If persistence is configured for a</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			particular domain, it takes precedence over the configured GSLB method. The NetScaler appliance supports persistence based on the source IP address or on HTTP cookies. With source-IP persistence, when a DNS request is received at a data center, the NetScaler appliance first looks for an entry in the persistence table and, if an entry for the local DNS server exists and the server mentioned in the entry is configured, the IP address of that server is sent as the DNS response. For the first request from a particular client, the NetScaler appliance selects the best GSLB site for the request and sends its IP address to the client. Since persistence is configured for the source IP address of the client, all subsequent requests by that client or another local DNS server in the same IP subnet are sent the IP address of the GSLB site that was selected for the first request. For source-IP address based persistence, the same set of persistence identifiers must be configured on the GSLB virtual servers in all data centers. A persistence identifier is a number used by the data centers to identify a particular GSLB virtual server. A cookie transmits the persistence identifier, enabling the NetScaler appliance to identify the domain so that it can forward all appropriate requests to the same domain. When persistence is enabled, the persistence information is also

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>exchanged as part of metrics exchange. In this case, the value of this measure will be Up. On the other hand, if persistence is not enabled, then persistence information will not be exchanged as part of metrics exchange. In this case, the value of this measure will be Down.

The numeric values that correspond to these measure values have been listed in the table below:

<table><tr><th>Measure value</th><th>Numeric Value</th></tr><tr><td>Up</td><td>0</td></tr><tr><td>Down</td><td>1</td></tr></table>
Note:

By default, this measure reports the above-mentioned Measure Values while indicating whether persistence exchange is enabled / disabled. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.</td></tr>
</table>

3.8.6 GSLB Services Test

A GSLB service is usually a representation of a load balancing or content switching virtual server, although it can represent any type of virtual server. The GSLB service identifies the virtual server's IP address, port number, and service type. GSLB services are bound to GSLB virtual servers on the NetScaler appliances managing the GSLB sites.

A GSLB service bound to a GSLB virtual server in the same data center is local to the GSLB virtual server. A GSLB service bound to a GSLB virtual server in a different data center is remote from that GSLB virtual server. A GSLB virtual server has one or more GSLB services bound to it, and load

balances traffic among those services. It evaluates the configured GSLB methods (algorithms) to select the appropriate service to which to send a client request. Because the GSLB services can represent either local or remote servers, selecting the optimal GSLB service for a request has the effect of selecting the data center that should serve the client request.

Using the **GSLB Services** test, you can monitor the request, response, and data load on the individual services configured on the NetScaler appliance, evaluate the efficiency with which the appliance distributes this load to the services, and promptly detect irregularities (if any) in load-balancing.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each GSLB service.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Up Server Only	The default setting of this flag is **No**; this indicates that this test, by default, monitors all the GSLB services configured on the NetScaler appliance. If you want the test to monitor only those GSLB services that are up and running currently, then set this value to **Yes**.
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

Parameter	Description
	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
Server state	Indicates the current state of the virtual server with which this service is bound.		If the virtual server is up, then the value of this measure is Up. If the virtual server is down, then the value of this measure is Down. The numeric values that correspond to these measure values have been listed in the table below: **Numeric Value** / **Measure Value** 0 / Up 1 / Down 2 / Out of service 3 / Transition out of service 4 / Down when going out of service -1 / Unknown

Measurement	Description	Measurement Unit	Interpretation
			Note: By default, this measure reports the above-mentioned **Measure Value**s while indicating whether a virtual server is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only. The detailed diagnosis of this measure reports the service type, the primary IP address of the service, and its primary port.
Current load on the service	Indicates the load on the service that is calculated from the bound load based monitor.	Percent	In a load-balanced setup, the value of this measure should be low or should be more or less the same for all services. A very high value for this measure could indicate that the service is overloaded. You may then want to consider fine-tuning your GSLB policies and metric exchange policies to ensure that the load-balancing irregularities are removed.
Current client connections	Indicates the number of current client connections to this service	Number	These are good measures of the connection load on a service. By comparing the value of each of these measures across services, you can instantly identify overloaded services and promptly detect irregularities in load-balancing.
Current server connections	Indicates the number of connections to the actual servers behind the virtual server that is bound to this service.	Number	
Client connections in established state	Indicates the number of client connections to this service that are currently in an *ESTABLISHED* state.	Number	
Server connections	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
in established state	connections to the actual server behind the virtual server that is bound to this service that are currently in an *ESTABLISHED* state.		
Request data received	Indicates the amount of request data received by the virtual server bound to this GSLB service during the last measurement period.	MB	These are good measures of the request and response load on a service. By comparing the value of each of these measures across services, you can instantly identify overloaded services and promptly detect irregularities in load-balancing. In the event that such irregularities come to light, you may want to consider fine-tuning the GSLB policies and/or GSLB methods and/or metric exchange policies supported by the system to ensure that the appliance takes intelligent load-balancing decisions.
Response data received	Indicates the amount of response data received by the virtual server bound to this GSLB service during the last measurement period.	MB	
Requests received	Indicates the number of requests received by the virtual server bound to this GSLB service during the last measurement period.	Number	
Responses received	Indicates the number of responses by the virtual server bound to this GSLB service during the last measurement period.	Number	
Virtual server hits	Indicates the number of times this service has been provided during the last measurement period.	Number	If the value of this measure is equal to the number of requests received for this service, then it indicates that all requests have been fulfilled. On the contrary, if the value of this measure is less than the number of requests received, it could indicate that one/more requests could not be serviced. This is a cause for concern.

3.8.7 GSLB Virtual Servers Test

A GSLB virtual server has one or more GSLB services bound to it, and load balances traffic among those services. It evaluates the configured GSLB methods (algorithms) to select the appropriate service to which to send a client request. Because the GSLB services can represent either local or remote servers, selecting the optimal GSLB service for a request has the effect of selecting the data center that should serve the client request.

The domain for which global server load balancing is configured must be bound to the GSLB virtual server, because one or more services bound to the virtual server will serve requests made for that domain. Unlike other virtual servers configured on a NetScaler appliance, a GSLB virtual server does not have its own virtual IP address (VIP).

To promptly detect irregularities in load-balancing and accurately point to the bottleneck points, you need to track requests to the sites, services, and virtual servers (that are bound to the services) and observe how every entity handles these requests. While the tests discussed previously monitor traffic to the sites and services, you can use the **GSLB Virtual Servers** test to figure out how the virtual servers load-balance the requests, responses, and data they receive. In the process, you can quickly identify overloaded virtual servers, and thus detect load-balancing irregularities.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each GSLB virtual server.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which the test is being configured.
NetScaler Username and NetScaler Password	To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the NetScaler Username and NetScaler Password text boxes.
SSL	The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why,

Parameter	Description
	the SSL flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the SSL flag to **No**.
Show Up Server Only	The default setting of this flag is **No**; this indicates that this test, by default, monitors all the GSLB virtual servers configured on the NetScaler appliance. If you want the test to monitor only those GSLB virtual servers that are up and running currently, then set this value to **Yes**.
Exclude Server	Provide a comma-separated list of GSLB virtual server names or name patterns that need to be excluded from monitoring. By default, this is set to *none*, indicating that all GSLB virtual servers are by default monitored
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
Server state	Indicates the current state of this virtual server.		If the virtual server is up, then the value of this measure is Up. If the virtual server is down, then the value of this measure is Down. The numeric values that correspond to these measure values have been listed in the table below:

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Numeric Value</th><th>Measure Value</th></tr>
<tr><td>0</td><td>Up</td></tr>
<tr><td>1</td><td>Down</td></tr>
<tr><td>2</td><td>Out of service</td></tr>
<tr><td>3</td><td>Transition out of service</td></tr>
<tr><td>4</td><td>Down when going out of service</td></tr>
<tr><td>-1</td><td>Unknown</td></tr>
</table>
Note:
By default, this measure reports the above-mentioned Measure Values while indicating whether a virtual server is up/down. However, in the graph of this measure, the Measure Values will be represented using their corresponding numeric equivalents only.
You can use the detailed diagnosis of this measure to determine the protocol type of each virtual server being monitored.
</td></tr>
<tr><td>Virtual server health</td><td>Indicates the current health of this virtual server.</td><td>Percent</td><td>While a high percentage is indicative of good health, a low percentage indicates otherwise.</td></tr>
<tr><td>Current client connections</td><td>Indicates the number of current client connections to this virtual server.</td><td>Number</td><td>These are good measures of the connection load on a virtual server. By comparing the value of each of these measures across servers, you can</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			instantly identify overloaded servers and promptly detect irregularities in load-balancing.
Current server connections	Indicates the number of connections to the actual servers behind this virtual server.	Number	
Client connections in established state	Indicates the number of client connections to this virtual server that are currently in an *ESTABLISHED* state.	Number	
Server connections in established state	Indicates the number of connections to the actual server behind this virtual server that are currently in an *ESTABLISHED* state.	Number	
Request data received:	Indicates the amount of request data received by this virtual server during the last measurement period.	MB	These are good measures of the request and response load on a virtual server. By comparing the value of each of these measures across virtual servers, you can instantly identify overloaded servers and promptly detect irregularities in load-balancing. In the event that such irregularities come to light, you may want to consider fine-tuning the GSLB policies and/or GSLB methods and/or metric exchange policies supported by the system to ensure that the appliance takes intelligent load-balancing decisions.
Response data received	Indicates the amount of response data received by this virtual server during the last measurement period.	MB	
Requests received	Indicates the number of requests received by this virtual server during the last measurement period.	Number	
Responses received	Indicates the number of responses received by this virtual server during the last measurement period.	Number	
Virtual server hits	Indicates the number of times this virtual server has serviced a request during the last	Number	If the value of this measure is equal to the number of requests received by this virtual server, then it indicates the virtual server was very busy during the

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		last measurement period. On the contrary, if the value of this measure is less than the number of requests received by this virtual server, it could indicate that the virtual server was well-loaded.

3.9 The NetScaler Users Layer

3.9.1 NetScaler Sessions Test

By tracking the user activity on the NetScaler appliance, administrators will be able to determine the session load on the appliance at any given point in time. Moreover, by observing user logins over time, administrators will be able to easily figure out if user logins are happening consistently or sporadically. The latter could hint at a connection issue, requiring immediate attention. Also, when monitoring user sessions, administrators will also be able to capture unexpected logouts or abrupt session terminations and investigate the reason for the same. Such useful session-based insights are provided by the **NetScaler Sessions** test. This test tracks ICA and VPN user logins to the NetScaler appliance, measures the current session load on the appliance, and reveals the communication channel (ICA or VPN) used by the maximum number of users. The test additionally reports the count and percentage of new logins through each channel, and also captures inexplicable session logouts. The detailed diagnostics provided by the test also point administrators to the users who logged in newly and the users whose sessions logged out.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every type - i.e., ICA and VPN - of user session that the NetScaler appliance handles

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **NETSCALER USERNAME** and **NETSCALER PASSWORD** - To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target

NetScaler device. Specify the credentials of such a user in the **NETSCALER USERNAME** and **NETSCALER PASSWORD** text boxes.

4. **CONFIRM PASSWORD** - Confirm the **NETSCALER PASSWORD** by retyping it here.
5. **SSL** - The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the **SSL** flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the **SSL** flag to **No**.
6. **DD FREQUENCY** – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.
7. **DETAILED DIAGNOSIS** – To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

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

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

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions :	Indicates the current number of sessions of this type (ICA or VPN) on the NetScaler appliance.	Number	This is a good indicator of the session load imposed on the NetScaler appliance by each communication channel (ICA or VPN).

Measurement	Description	Measurement Unit	Interpretation
New logins:	Indicates the number of new logins of this type to the NetScaler in the last measurement period.	Number	A consistent zero value could indicate a connection issue. You can use the detailed diagnosis of this test to know which users logged in recently.
Percent new logins:	Indicates the percentage of current sessions of this type that logged in during the last measurement period.	Percent	
Sessions logging out:	Indicates the number of sessions of this type that logged out.	Number	If all the current sessions of a type suddenly log out, it indicates a problem condition that requires investigation. The detailed diagnosis of this measure lists the sessions that logged out.

3.10 The NetScaler Users Layer

3.10.1 NetScaler Sessions Test

By tracking the user activity on the NetScaler appliance, administrators will be able to determine the session load on the appliance at any given point in time. Moreover, by observing user logins over time, administrators will be able to easily figure out if user logins are happening consistently or sporadically. The latter could hint at a connection issue, requiring immediate attention. Also, when monitoring user sessions, administrators will also be able to capture unexpected logouts or abrupt session terminations and investigate the reason for the same. Such useful session-based insights are provided by the **NetScaler Sessions** test. This test tracks ICA and VPN user logins to the NetScaler appliance, measures the current session load on the appliance, and reveals the communication channel (ICA or VPN) used by the maximum number of users. The test additionally reports the count and percentage of new logins through each channel, and also captures inexplicable session logouts. The detailed diagnostics provided by the test also point administrators to the users who logged in newly and the users whose sessions logged out.

Target of the test : A NetScaler VPX/MPX

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every type - i.e., ICA and VPN - of user session that the NetScaler appliance handles

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **NETSCALER USERNAME** and **NETSCALER PASSWORD** - To monitor a NetScaler device, the eG agent should be configured with the credentials of a user with *read-only privileges* to the target NetScaler device. Specify the credentials of such a user in the **NETSCALER USERNAME** and **NETSCALER PASSWORD** text boxes.
4. **CONFIRM PASSWORD** - Confirm the **NETSCALER PASSWORD** by retyping it here.
5. **SSL** - The eG agent collects performance metrics by invoking NITRO (NetScaler Interface Through Restful interfaces and Objects) APIs on the target NetScaler device. Typically, the NITRO APIs can be invoked through the HTTP or the HTTPS mode. By default, the eG agent invokes the NITRO APIs using the HTTPS mode. This is why, the **SSL** flag is set to **Yes** by default. If the target NetScaler device is not SSL-enabled, then the NITRO APIs can be accessed through the HTTP mode only. In this case, set the **SSL** flag to **No**.
6. **DD FREQUENCY** – Refers to the frequency with which detailed diagnosis measures are to be generated for this test. The default is *1:1*. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. You can modify this frequency, if you so desire. Also, if you intend to disable the detailed diagnosis capability for this test, you can do so by specifying *none* against **DD FREQUENCY**.
7. **DETAILED DIAGNOSIS** – To make diagnosis more efficient and accurate, the eG Enterprise suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

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

   - The eG manager license should allow the detailed diagnosis capability

- Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Current sessions :	Indicates the current number of sessions of this type (ICA or VPN) on the NetScaler appliance.	Number	This is a good indicator of the session load imposed on the NetScaler appliance by each communication channel (ICA or VPN).
New logins:	Indicates the number of new logins of this type to the NetScaler in the last measurement period.	Number	A consistent zero value could indicate a connection issue. You can use the detailed diagnosis of this test to know which users logged in recently.
Percent new logins:	Indicates the percentage of current sessions of this type that logged in during the last measurement period.	Percent	
Sessions logging out:	Indicates the number of sessions of this type that logged out.	Number	If all the current sessions of a type suddenly log out, it indicates a problem condition that requires investigation. The detailed diagnosis of this measure lists the sessions that logged out.

About eG Innovations

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

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

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

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