Monitoring JetNEXUS Load Balancer

eG Innovations Product Documentation

www.eginnovations.com

eG
Total Performance Visibility

Table of Contents

CHAPTER 1: INTRODUCTION 1

CHAPTER 2: HOW TO MONITOR JETNEXUS LOAD BALANCER USING EG ENTERPRISE? 2

2.1 Managing JetNEXUS Load Balancer 2

CHAPTER 3: MONITORING JETNEXUS LOAD BALANCER 4

3.1 Hardware Layer 4

3.1.1 System Resource Test 5

3.2 Nexus LB Statistics Layer 7

3.2.1 Real Server Test 8

3.2.2 System Throughputs Test 12

3.2.3 Virtual Server Test 16

ABOUT EG INNOVATIONS 21

Table of Figures

Figure 2.1: Adding a JetNEXUS Load Balancer 3

Figure 2.2: List of Unconfigured tests to be configured for the JetNEXUS Load Balancer 3

Figure 3.1: The layer model of the JetNEXUS Load Balancer 4

Figure 3.2: The test associated with the Hardware layer 5

Figure 3.3: Tests mapping to the Nexus LB Statistics layer 7

Chapter 1: Introduction

JetNEXUS load balancers mitigate the risk of downtime, improving the performance, scalability and reliability of applications for a better end user experience. The JetNEXUS load balancers are best suited for Enterprise applications such as MS Exchange and Lync, or self-service websites or browser-based applications, particularly where user experience is key. The key Benefits offered by the JetNEXUS load balancers are:

- Load balance traffic across multiple application servers for high availability
- Reduce load on web servers and improve application performance
- Reverse Proxy functionality
- Advanced toolkit to help solve complex web application delivery problems

Since application delivery delays, inefficiencies, and failures can cause prolonged service outages and cost an enterprise money and reputation, the continuous operation and good health of the load balancer is of great importance. Therefore, it is imperative that the JetNEXUS Load Balancer should be continuously monitored to prevent such eventualities. This is where eG Enterprise helps administrators!

Chapter 2: How to Monitor JetNEXUS Load Balancer Using eG Enterprise?

eG Enterprise monitors the JetNEXUS load balancer using an eG external agent on any remote host in the environment. This agent is capable of monitoring the performance of the target load balancer appliance by polling the SNMP MIB of the load balancer. To enable the eG agent to collect performance metrics from the target load balancer, the following pre-requisites should be fulfilled:

- The JetNEXUS load balancer should be SNMP-enabled.
- The eG external agent should be able to access the target load balancer over the network.

Once the pre-requisites are fulfilled, manage the target load balancer using the eG admin interface. The procedure has been discussed in the following section.

2.1 Managing JetNEXUS Load Balancer

eG Enterprise can automatically discover the JetNEXUS Load Balancer, and also lets to manually add the component for monitoring using eG admin interface. To manage a JetNEXUS Load Balancer component, do the following:

1. Log into the eG admin interface.
2. If the target load alancer is already discovered, then directly proceed towards managing the broker using the **COMPONENTS – MANAGE/UNMANAGE** page.
3. However, if you are yet to discover the JetNEXUS Load Balancer, then run discovery (Infrastructure -> Components -> Discover) or follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu to manually add the component using the **Components** page.
4. Click on the **Add new Component** button after choosing the *JetNEXUS Load Balancer* from the **Component Type** drop down list in the **Components** page. This will lead you to the **Add Component** page (Figure 2.1). Remember that components manually added are managed automatically.

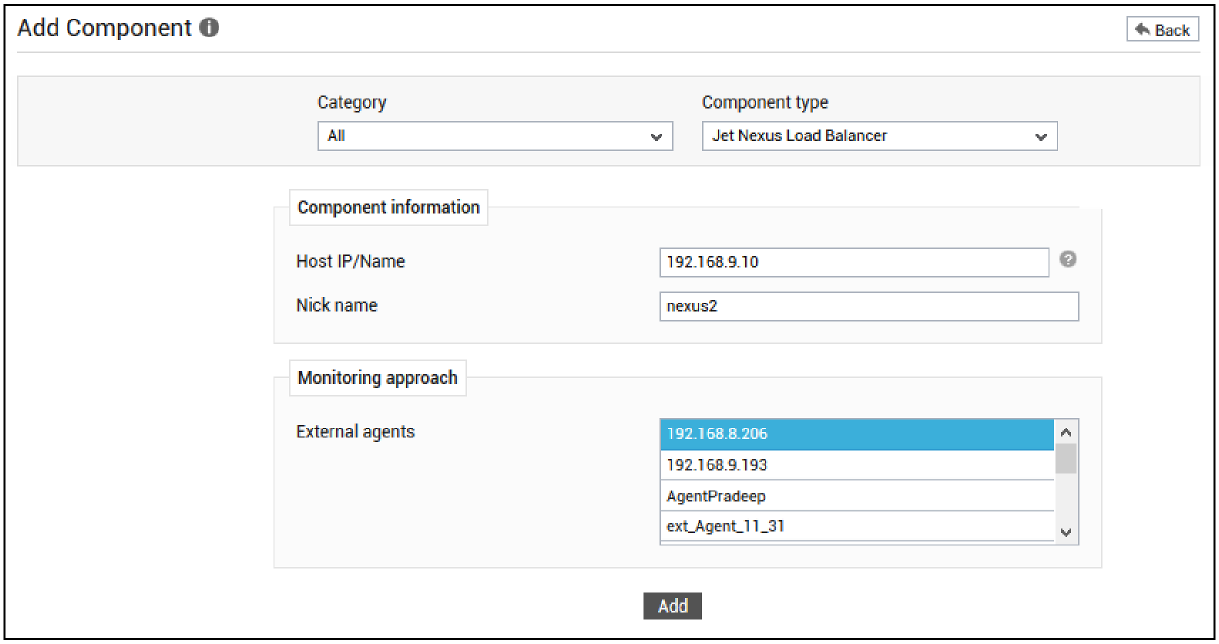

Add Component

Category: All

Component type: Jet Nexus Load Balancer

Component information

Host IP/Name: 192.168.9.10

Nick name: nexus2

Monitoring approach

External agents:
- 192.168.8.206
- 192.168.9.193
- AgentPradeep
- ext_Agent_11_31

Figure 2.1: Adding a JetNEXUS Load Balancer

5. Specify the **Host IP/Name** and the **Nick name** of the JetNEXUS Load Balancer in Figure 2.1.
6. Then, pick an external agent from the **External agents** list box and click the **Add** button to add the component for monitoring.
7. When you attempt to sign out, a list of unconfigured tests will appear as shown in Figure 2.2.

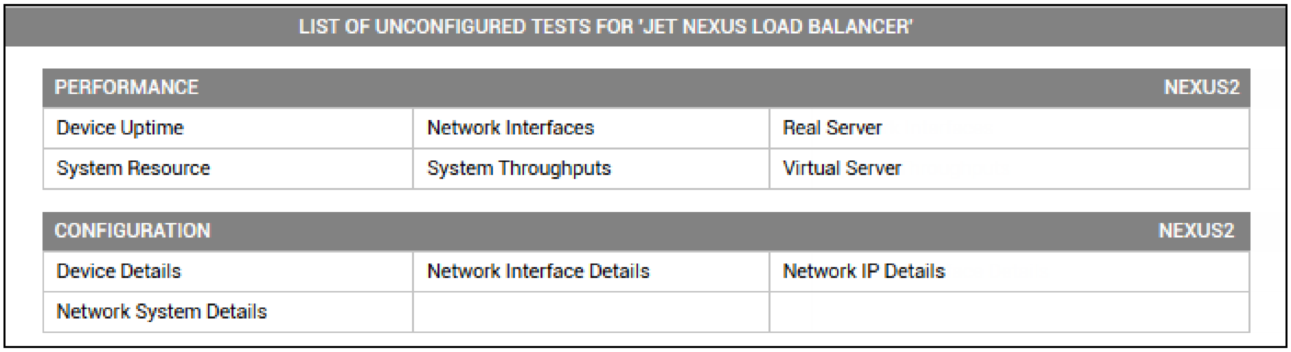

LIST OF UNCONFIGURED TESTS FOR 'JET NEXUS LOAD BALANCER'

PERFORMANCE		NEXUS2
Device Uptime	Network Interfaces	Real Server
System Resource	System Throughputs	Virtual Server

CONFIGURATION		NEXUS2
Device Details	Network Interface Details	Network IP Details
Network System Details		

Figure 2.2: List of Unconfigured tests to be configured for the JetNEXUS Load Balancer

8. To configure the tests, click on any of the tests for e.g., Real Server test in Figure 2.2. To know how to configure the test, refer to Section **3.2.1**.
9. Once the tests are configured, signout of the eG admin interface.

Chapter 3: Monitoring JetNEXUS Load Balancer

To ensure continuous operation and good health of the JetNEXUS Load Balancer and its core components, eG Enterprise provides a specialized JetNEXUS Load Balancer model (see Figure 3.1).

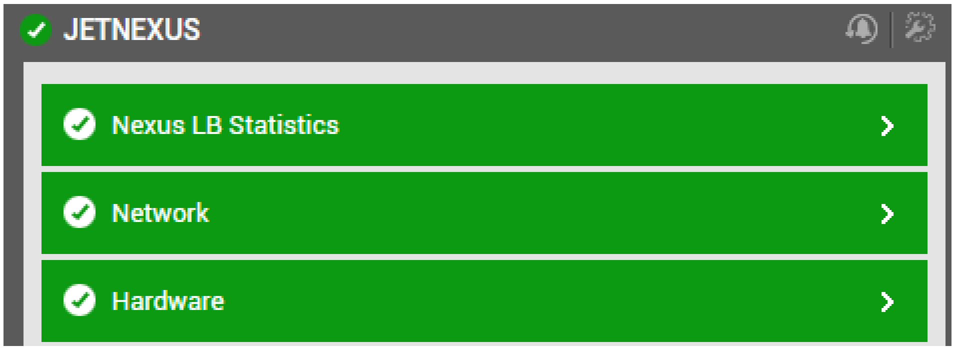

Figure 3.1: The layer model of the JetNEXUS Load Balancer

Every layer of Figure 3.1 is mapped to a variety of tests which connect to the SNMP MIB of the target load balancer to collect critical statistics pertaining to its performance. The metrics reported by these tests enable administrators to answer the following questions:

- Is CPU/memory/disk on the target load balancer over-utilized?
- Was the throughput of the target load balancer optimal while processing data/compressed data?
- How well the real server is processing client traffic? Which server is handling the maximum traffic?
- How well the virtual server is processing client traffic? Which virtual server is handling the maximum traffic?

Since the details about the **Network** test is available in the *Monitoring Unix and Windows Servers* document and the other tests in the Network layer is available in the *Monitoring Cisco Routers* document, the sections that follow will discuss the remaining layers in Chapter 3.

3.1 Hardware Layer

Using the System Resource test mapped to this layer, you can monitor the CPU, disk and memory utilization of the target JetNEXUS Load Balancer.

Figure 3.2: The test associated with the Hardware layer

3.1.1 System Resource Test

This test collects metrics pertaining to the CPU, disk and memory usage of the target load balancer.

Target of the test : A JetNEXUS Load Balancer

Agent deploying the test : An external agent

Outputs of the test : One set of results for the target load balancer 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 JetNEXUS Load Balancer device that is being monitored.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; the default is *161*.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the load balancer. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
Username	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP

Parameter	Description
	entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPVersion. From the AuthType list box, choose the authentication algorithm using which SNMP v3 converts the specified Username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPVersion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific

Parameter	Description
	components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
CPU usage	Indicates the percentage of CPU utilized on the target load balancer.	Percent	Ideally, the value of this measure should be low, A high value for this measure may indicate a CPU bottleneck.
Memory usage	Indicates the percentage of memory utilized on the target load balancer.	Percent	Ideally, the value of this measure should be low. A consistent increase in this value could be a cause for some serious concern, as it indicates a gradual, but steady erosion of the memory resources.
Disk usage	Indicates the percentage of disk space utilized on the target load balancer.	Percent	Ideally, the value of this measure should be low.

3.2 Nexus LB Statistics Layer

This layer tracks the statistics pertaining to the client traffic processing ability of each virtual server configured on the target load balancer and each real server associated to the virtual server. In addition, this layer also detects how well the target load balancer processes data/compressed data.

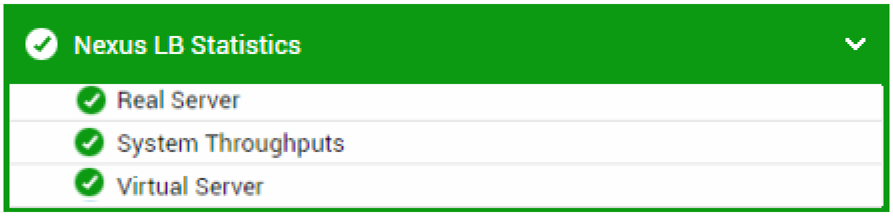

Figure 3.3: Tests mapping to the Nexus LB Statistics layer

3.2.1 Real Server Test

Real servers are dedicated physical servers that you typically configure in groups called real server groups. The JetNEXUS Load Balancer distributes incoming network traffic and session load across a real server group.Typically, the real servers in a group are bound to a virtual server. Whenever a client request is received, the virtual server bound to a real server group responds to those requests by routing the requests to those real servers in that group that are currently available. Load distribution can be done in different ways. The load balancer may use a round-robin method, where each server is used in turn. It can also use a weighted round robin system, where servers are assigned traffic based on their configured capabilities. Regardless of which load balancing technique/algorithm is used, the aim is to ensure that no single real server is overloaded with requests.

If one/more real servers in the group are offline/not connected/in standby mode for longer time, then the load should be shared by the other real servers in the group. However, if at any given point in time, one/more real servers in a group handle a significantly higher session and/or data load than the rest, it is a clear indicator of ineffective load-balancing! Under such circumstances, administrators may have to fine-tune/change the load-balancing algorithm.

To be able to quickly and accurately spot load-balancing irregularities and to initiate remedial measures, administrators should keep a close watch on the status of each real server in a group and the connections and data load handled by each server. This is what the **Real Servers** Test does!

For each real server in a group, this test reports the current status of the server. In addition, the test also enables you to analyze the data processing ability of each real server in the group by reporting the amount of data transmitted/received per second by each server. In the process, this test also reveals the data compression ratio, pool memory utilization and total number of clients connected to each server. This way, the test sheds light on issues in load-balancing, and thus urges administrators to take appropriate corrective action.

Target of the test : A JetNEXUS Load Balancer

Agent deploying the test : An external agent

Outputs of the test : One set of results for each *Virtual server:Real server* combination on the target load balancer being monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the JetNEXUS load balancer that is being monitored.
SSH Port	Besides SNMP, this test also uses the Radware CLI to pull metrics on real servers. To run the CLI commands, the test first needs to establish an SSH connection with the Radware Alteon load balancer. To enable the test to establish this connection, specify the **SSH PORT** here.
SSH Username, SSH Password, and Confirm Password	As stated earlier, this test also uses the Radware CLI to pull metrics on real servers. To use the CLI, the test first needs to connect to the Radware Alteon load balancer via SSH, and then run commands using CLI. For running the commands, this test requires the privileges of a valid SSH user with permission to run the CLI commands. Specify the user name and password of such a user against **SSH USERNAME** and **SSH PASSWORD** text boxes, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; the default is *161*.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the firewall. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
Username	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a

Parameter	Description
	contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPVersion. From the AuthType list box, choose the authentication algorithm using which SNMP v3 converts the specified Username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPVersion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By

Parameter	Description
	default, this flag is set 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

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Status</td><td>Indicates the current status of this real server.</td><td></td><td>The values that this measure can report and their corresponding numeric values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Not connected</td><td>0</td></tr>
<tr><td>Offline</td><td>1</td></tr>
<tr><td>Standby</td><td>2</td></tr>
<tr><td>Online</td><td>3</td></tr>
<tr><td>Draining</td><td>4</td></tr>
<tr><td>Not monitored</td><td>5</td></tr>
<tr><td>Finding status</td><td>6</td></tr>
<tr><td>Not licensed</td><td>7</td></tr>
<tr><td>Unknown</td><td>8</td></tr>
</table>
<b>Note:</b>
By default, this measure reports one of</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			the **Measure Values** listed in the table above to indicate the status of the real server. In the graph of the measure however, the real server status is indicated using the corresponding numeric equivalents only. Use the detailed diagnosis of this measure to determine the IP address of the real server.
Data received	Indicates the rate at which data was received by this real server during the last measurement period.	KB/sec	Compare the values of these measures across the real servers to identify the server that is handling maximum traffic.
Data transmitted	Indicates the rate at which data was transmitted from this real server during the last measurement period.	KB/sec	
Compression	Indicates the data compression percentage of this real server.	Percent	
Total clients	Indicates the total number of clients established connection on this real server.	Number	
Pool usage	Indicates the percentage of pool memory utilized by this real server.	Percent	
Hit count	Indicates the total number of requests received by this real server.	Number	

3.2.2 System Throughputs Test

Periodically, administrators should measure the workload on the JetNEXUS Load Balancer device and evaluate the load balancer's ability to handle the load, so that they can figure out whether the server's sized commensurate to its load or not. Administrators can perform this load analysis using the **System Throughputs** test.

This test tracks the load on the target load balancer and reports what type of data is contributing to the workload of the load balancer. The test additionally reports the number of connections established to the load balancer, so that you can figure out the count of connections that are generating the load on the load balancer.

Target of the test : A JetNEXUS Load Balancer

Agent deploying the test : An external agent

Outputs of the test : One set of results for the target load balancer 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 JetNEXUS Load Balancer device that is being monitored.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; the default is *161*.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the load balancer. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
Username	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the

Parameter	Description
	SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPVersion. From the AuthType list box, choose the authentication algorithm using which SNMP v3 converts the specified Username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPVersion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such

Parameter	Description
	environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates rate at which data was received by the load balancer during the last measurement period.	KB/sec	
Data transmitted	Indicates rate at which data was transmitted by the load balancer during the last measurement period.	KB/sec	
Compressed data received	Indicates rate at which compressed data was received by the load balancer during the last measurement period.	KB/sec	jetNEXUS is able to compress web content during its journey from server to client, accelerating delivery and dramatically improving user experience.
Compressed data transmitted	Indicates rate at which compressed data was transmitted from the load balancer during the last measurement period.	KB/sec	
Total connections	Indicates the total number of connections established to the load balancer during the last measurement period.	Number	
Current connections	Indicates the number of connections that are currently active on the load balancer.	Number	This measure is a good indicator of the current workload on the load balancer.

3.2.3 Virtual Server Test

Each real server group is mapped to a virtual server. The virtual server receives incoming client requests , uses a load-balancing algorithm to select an available real server in the group, and routes the requests to the selected real server.

Since application requests are front-ended by a virtual server, tracking the session load on a virtual server will reveal the load on the applications running on the real servers mapped to that virtual server. This also means that if the virtual server experiences a processing bottleneck, it is bound to impact load distribution and processing by the real servers, which will consequently affect user experience with the applications in the backend. To understand the load on your applications and to proactively detect any potential slowness that your applications may experience, it is good practice to continuously monitor each virtual server. This is exactly what the **Virtual Server** test does!

This test auto-discovers the virtual servers that are configured on the target load balancer and continuously monitors the current status of each virtual server and reveals how well each virtual server processes data. In addition, the test then reports how well the cache in each virtual server services requests. This way, you can figure out performance issues in the load-balancing virtual servers.

Target of the test : A JetNEXUS Load Balancer

Agent deploying the test : An external agent

Outputs of the test : One set of results for each *Virtual server*.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The IP address of the JetNEXUS load balancer that is being monitored.
SSH Port	Besides SNMP, this test also uses the Radware CLI to pull metrics on real servers. To run the CLI commands, the test first needs to establish an SSH connection with the Radware Alteon load balancer. To enable the test to establish this connection, specify the **SSH PORT** here.
SSH Username, SSH Password, and Confirm Password	As stated earlier, this test also uses the Radware CLI to pull metrics on real servers. To use the CLI, the test first needs to connect to the Radware Alteon load balancer via SSH, and then run commands using CLI. For running the commands, this test requires the privileges of a valid SSH user with permission to run the CLI commands. Specify

Parameter	Description
	the user name and password of such a user against **SSH USERNAME** and **SSH PASSWORD** text boxes, and confirm the password by retyping it in the **CONFIRM PASSWORD** text box.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; the default is *161*.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the firewall. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
Username	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPVersion. From the AuthType list box, choose the authentication algorithm using which SNMP v3 converts the specified Username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options:

Parameter	Description
	 - **MD5** – Message Digest Algorithm - **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPVersion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: - **DES** – Data Encryption Standard - **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set 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

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Status</td><td>Indicates the current status of this virtual server.</td><td></td><td>The values that this measure can report and their corresponding numeric values are as follows:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>Not connected</td><td>0</td></tr>
<tr><td>Offline</td><td>1</td></tr>
<tr><td>Standby</td><td>2</td></tr>
<tr><td>Online</td><td>3</td></tr>
<tr><td>Draining</td><td>4</td></tr>
<tr><td>Not monitored</td><td>5</td></tr>
<tr><td>Finding status</td><td>6</td></tr>
<tr><td>Not licensed</td><td>7</td></tr>
<tr><td>Unknown</td><td>8</td></tr>
</table>
<b>Note:</b>
By default, this measure reports one of the <b>Measure Value</b>s listed in the table above to indicate the status of the real server. In the graph of the measure however, the real server status is indicated using the corresponding numeric equivalents only.
Use the detailed diagnosis of this measure to determine the IP address of the real server.</td></tr>
<tr><td>Data received</td><td>Indicates the rate at which data was received by this virtual server during the last measurement period.</td><td>KB/sec</td><td rowspan="2">Compare the values of these measures across the real servers to identify the server that is handling maximum traffic.</td></tr>
<tr><td>Data transmitted</td><td>Indicates the rate at which data was transmitted from this virtual server during</td><td>KB/sec</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	the last measurement period.		
Data cached	Indicates the rate at which the data was cached in this virtual server during the last measurement period.	KB/sec	Content Caching is a performance optimization feature that reduces server load and significantly increases the scalability of web-based applications, particularly those with peak, lumpy or unpredictable traffic profiles. For sites or applications with a large proportion of cacheable content, Caching can dramatically reduce the amount and size of requests to back-end servers.
Compression	Indicates the data compression percentage of this virtual server.	Percent	
Total clients	Indicates the total number of clients established connection on this virtual server.	Number	
Hit count	Indicates the total number of requests received by this virtual server.	Number	
Cache hit count	Indicates the total number of requests serviced from the cache in this virtual server.	Number	
Cache hit	Indicates the percentage of requests serviced from the cache in this virtual server.	Percent	A high value is desired for this measure.

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.

This document may not be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the prior written consent of eG Innovations. eG Innovations makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The information contained in this document is subject to change without notice.

All right, title, and interest in and to the software and documentation are and shall remain the exclusive property of eG Innovations. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.