Monitoring Coyote Point Equalizer

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

Coyote Point Equalizer load balancers are a cost-effective appliance-based solution for managing the scalability, availability and performance requirements of any network infrastructure. By effectively managing Internet traffic, the Equalizer product line maximizes network potential by minimizing response times and ensuring site availability.

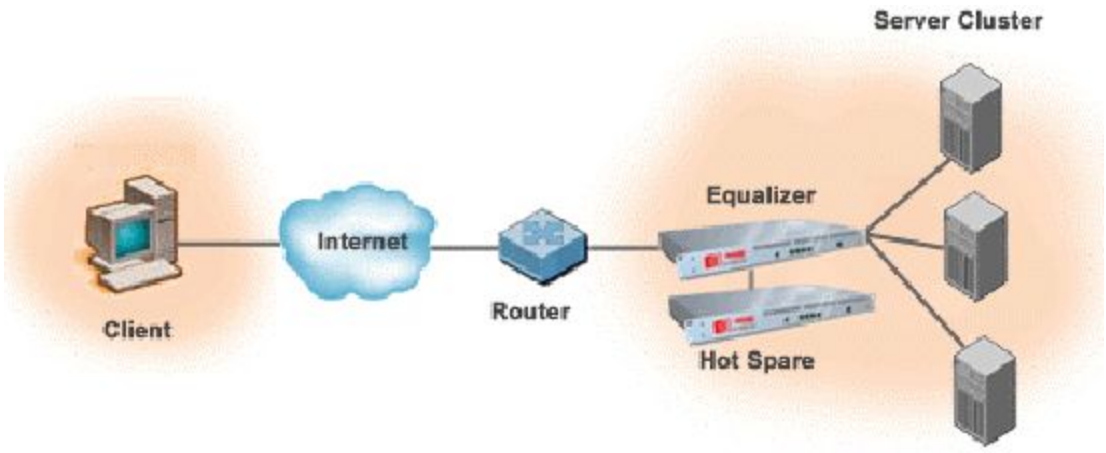

Figure 1.1: Typical deployment architecture of the Equalizer

As a gateway appliance, Coyote Point load balancers are typically deployed in a redundant configuration that includes a hot backup. Client requests are routed through the Equalizer to the appropriate server based on rules set by the administrator.

Since these load balancers are platform and (internet) protocol-independent, they are common-place in mission-critical business environments where maximum performance and high availability are key. Performance issues experienced by the equalizer can therefore adversely impact the availability of the critical services delivered by such environments, disrupting business and causing considerable revenue loss in the process. By continuously monitoring the operations and overall performance of the equalizer, such unpleasant eventualities can be avoided. The eG Enterprise helps network administrators in this regard!

Chapter 2: How does eG Enterprise Monitor Coyote Point Equalizer load balancers?

The eG Enterprise is capable of monitoring the Coyote Point Equalizer in an *agentless* manner using the using SNMP. The eG external agent periodically checks the SNMP MIB of the Coyote Point Equalizer for fetching metrics related to the performance of the Coyote Point Equalizer load balancer. This sections that follow describes how to manage and monitor the Coyote Point Equalizer load balancer.

2.1 Managing the Coyote Point Equalizer

The eG Enterprise cannot automatically discover a Coyote Point Equalizer so that you need to manually add the component for monitoring. To manage a Coyote Point Equalizer component, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENT** page that appears select Coyote Point Equalizer as the **Component** type. Then, click the **Add New Component** button. This will invoke Figure 2.1.

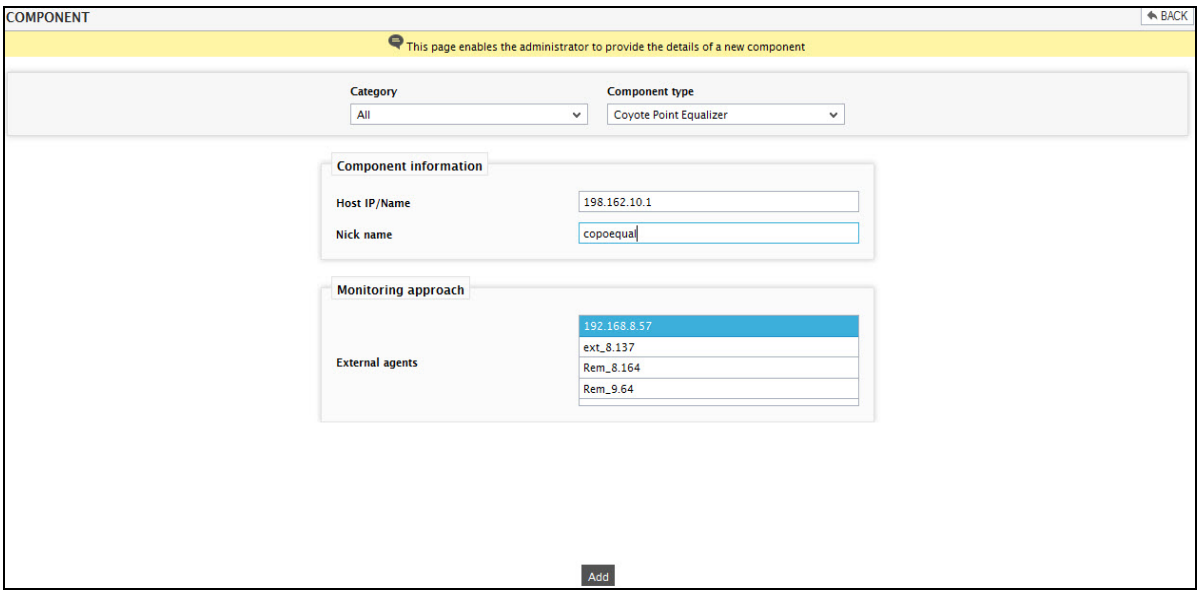

Figure 2.1: Adding the Coyote Point Equalizer

4. Specify the **Host IP/Name** and **Nick name o**f the Coyote Point Equalizer component to be monitored as shown in Chapter 2. Then, click **Add** button to register the changes.
5. When you attempt to sign out, a list of unconfigured tests appears.

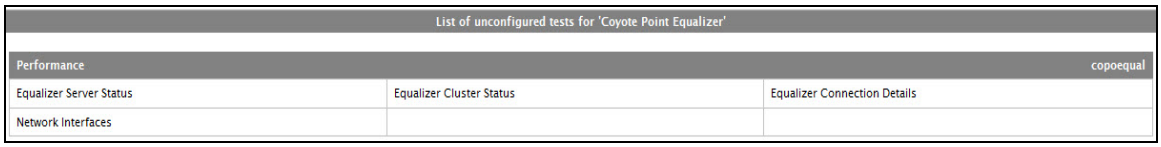

Figure 2.2: List of unconfigured tests to be configured for the Coyote Point Equalizer

6. Click on the **Equalizer Cluster Status** test to configure it. To know how to configure the test, refer to Section **3.2.1**.
7. Next, try to signout of the administrative interface. Now you will be prompted to configure the **Equalizer Server Status** test.
8. Click on the **Equalizer Server Status** test to configure it. To know how to configure the test, refer to Section **3.2.3**.
9. Finally signout of the eG administrative interface.

Chapter 3: Monitoring the Coyote Point Equalizer

eG Enterprise offers a specialized *Coyote Point Equalizer* (see Figure 3.1) monitoring model, which involves a single eG external agent that periodically polls the SNMP MIB of the equalizer, and collects a wide variety of performance information revealing the load on the device and the effectiveness with which the device balances this load across the servers in a farm. In the event of inconsistencies in load balancing, the agent proactively alerts administrators to the potential problem, so that he/she can initiate the relevant remedial action immediately.

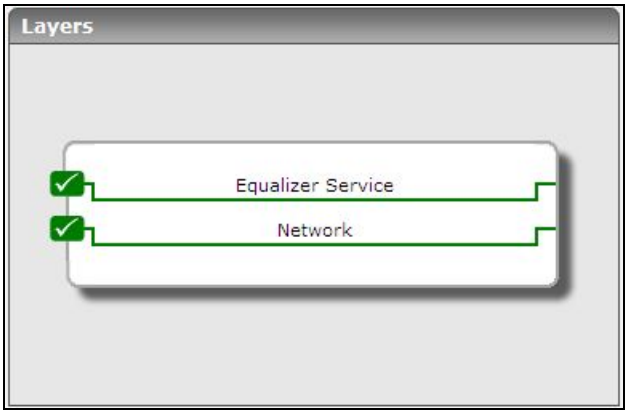

Figure 3.1: The layer model of the Coyote Point Equalizer

Each layer of Figure 3.1 above is mapped to tests that report the following:

- How many clusters are being managed by the equalizer and what are they? Is any cluster overloaded currently? If so, which one is it?
- Which cluster is currently handling the maximum number of connections?
- Which cluster is the busiest in terms of hits to its servers?
- How is the connection load on the equalizer? Is the equalizer able to handle the load?
- Which type of connections is the highest on the equalizer - Level-4 or Level-7?
- Did any connection to the equalizer time out?
- Is the equalizer evenly distributing load across all the servers in the cluster, or is any server currently overloaded?
- Is the equalizer able to assure requests of quick responses from the servers, or is any server in the cluster responding slowly to client requests? Is it owing to a badly tuned equalizer?

- Are client connections to a cluster uniformly distributed across all the servers in that cluster? If not, what is the reason for the imbalance?
- Is any server in the cluster idle?

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

3.1 The Network Layer

The tests mapped to the **Network** layer reveal the following:

- The availability of the equalizer and its responsiveness to requests
- The quality of network connections to the equalizer;
- The speed and bandwidth used by each of the network interfaces supported by the equalizer.

Figure 3.2: The tests mapped to the Network layer

Since all the tests displayed in Figure 3.2 have been dealt with extensively in the *Monitoring Cisco Routers* document, let us proceed to the next layer.

3.2 The Equalizer Service Layer

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

- The number and type of connections handled by the equalizer;
- The current load on the servers in the cluster and the server responsiveness;
- The load on the clusters managed by the equalizer, and the throughput of each cluster.

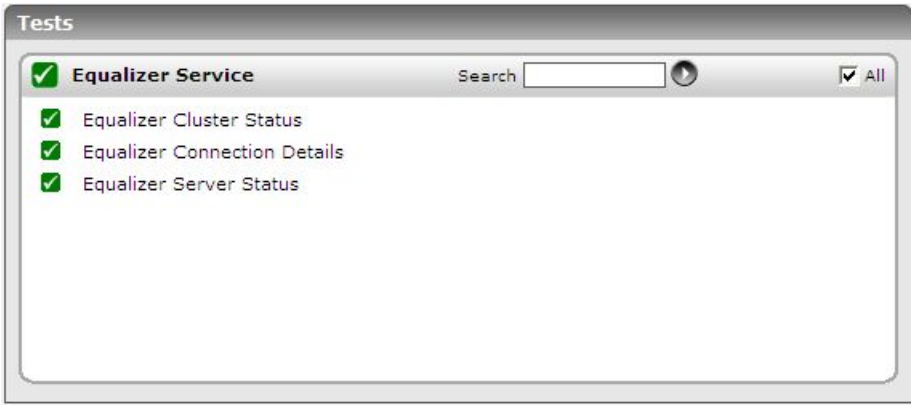

Figure 3.3: The tests mapped to the Equalizer Service layer

3.2.1 Equalizer Cluster Status Test

The Equalizer typically manages traffic to a group of servers in a server farm. While the servers in a farm can still be individually accessed, all traffic to the servers will be directed to a separate IP address, called a Virtual Cluster. The Virtual Cluster will accept traffic and distribute it to the available servers.

An Equalizer can be configured to manage multiple server farms/clusters. To be able to accurately assess the workload of the equalizer, you need to have a fair idea of the connection and data load on each of the clusters it manages. The Equalizer Cluster Status test enables you to ascertain the same. For each cluster, this test reports the current load on the cluster and indicates how busy the servers in the cluster are.

Target of the test : A Coyote Point Equalizer

Agent deploying the test : An external agent

Outputs of the test : One set of results for the each cluster managed by the target equalizer

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 this test is to be configured.
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

Parameter	Description
	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: • **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.

Parameter	Description
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**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Cluster load	Indicates the calculated load value for this cluster.	Number	This serves as a good indicator of the cluster workload. Comparing the value of this measure across clusters will enable you to identify those clusters that are overloaded.
Current connections	Indicates the number of connections currently active on this cluster.	Number	This again serves as a good indicator of the cluster workload.
Total connections	Indicates the total number of connections handled by this cluster.	Number	
Throughput	Indicates the rate of data traffic handled by this cluster over the last	MB/Sec	

Measurement	Description	Measurement Unit	Interpretation
	second.		
Hit rate	Indicates the rate at which servers in this cluster were accessed for performing transactions.	Mbps	Comparing the value of this measure across clusters will enable you to quickly spot the busiest clusters.

3.2.2 Equalizer Connection Details Test

This test not only reports the connection load on the equalizer in numbers, but also points to the nature of the workload by revealing the type of connections handled by the equalizer – this way, administrators can evaluate the workload of the device better. In addition, the test also turns the spotlight on inactive/idle connections, so that administrators can make sure that such connections are kept at a bare minimum.

Target of the test : A Coyote Point Equalizer

Agent deploying the test : An external agent

Outputs of the test : One set of results for the equalizer 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 this test is to be configured.
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

Parameter	Description
	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: • **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

Parameter	Description
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**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Level4 total connections	Indicates the number of L4 connections currently processed by the equalizer.	Number	• This serves as a good indicator of the Level-4 connection load on the equalizer. • Level-4 load balancing is to distribute requests to the servers at transport layer, such as TCP, UDP and SCTP transport protocol. The load balancer distributes network connections from clients who know a single IP address for a service, to a set of servers that actually perform the work. Since connection must be established between client and server in connection-oriented transport before sending the request content, the load balancer usually selects a server without looking at

Measurement	Description	Measurement Unit	Interpretation
			the content of the request.
Level4 peak connections	Indicates the high watermark of L4 connections processed by the equalizer.	Number	
Level4 idle timeout count	Indicates the number of L4 connections that timed out currently, because they were unused for a long time.	Number	Ideally, the value of this measure should be 0. A sudden/steady increase in this value could be a cause for concern.
Level7 active connections	Indicates the number of L7 connections currently active on the equalizer.	Number	Both these measures serve as effective pointers to the L7 connection workload on the equalizer. Layer-7 load balancing, also known as application-level load balancing, is to parse requests in application layer and distribute requests to servers based on different types of request contents, so that it can provide quality of service requirements for different types of contents and improve overall cluster performance. The overhead of parsing requests in application layer is high, thus its scalability is limited, compared to layer-4 load balancing. This is turn implies that a very high value for this measure will be accompanied by a significant increase in the processing overheads, but will ensure improved cluster performance.
Level7 total connections	Indicates the total number of L7 connections to the equalizer.	Number	
Level7 peak connections	Indicates the high watermark of L7 connections to the equalizer.	Number	

3.2.3 Equalizer Server Status Test

The real test of the efficiency of a load balancer lies in its ability to uniformly distribute load across the servers in a cluster, thereby ensuring the peak performance and continuous availability of the dependent services. Using the **Equalizer Server Status** test, administrators can accurately judge the efficiency and effectiveness of the equalizer. This test monitors the connection and calculated load on each server in a cluster, promptly detects load imbalances, and alerts administrators to them, so that they can quickly resolve the issue.

Target of the test : A Coyote Point Equalizer

Agent deploying the test : An external agent

Outputs of the test : One set of results for each server in each cluster managed by the equalizer

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 this test is to be configured.
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

Parameter	Description
	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: • **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

Parameter	Description
	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**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Server load	Indicates the current calculated load value for this server.	Number	This indicates the workload on the server. By comparing the value of this measure across all the servers in a cluster, you can instantly identify irregularities in load balancing. If found necessary, you can reconfigure the load balancing rules to ensure uniform load distribution across servers.
Response time	Indicates how quickly this server is currently responding to client requests.	ms	It is the job of a load balancer to ensure minimal response time for client requests. A high value for this measure could therefore indicate a defective load balancer or one that is improperly configured. Further investigation is hence necessary in this case to identify the root-cause of the anomaly.
Current connections	Indicates the number of connections that were active on this server during the last measurement period.	Number	The indicates the connection load on the server. By observing the graph of this measure over time, you can analyze the rate of growth of the load on the server. By comparing the value of this measure across all the servers in a cluster, you can instantly identify overloaded servers; this in turn brings irregularities in load balancing to light.
Total connections	Indicates the number of current connections to this server.	Number	If a sudden/consistent increase in the value of this measure is noticed, you might have to investigate further to identify the reason for this occurrence.

Measurement	Description	Measurement Unit	Interpretation
Idle time	Indicates the time for which this server was idle.	Secs	Ideally, the value of this measure should be low. A high value indicates that the server has remained unused for a long time. This could be owing to inconsistencies in load balancing or because the server is unavailable for use.

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.