NetFlow Monitoring Using eG Enterprise

eG Innovations Product Documentation

www.eginnovations.com

eG
Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

NetFlow is a network protocol developed by Cisco in order to collect and monitor IP network traffic. It has now become the de-facto industry standard and is supported by platforms other than Cisco including; Juniper (Jflow); 3Com/HP, Dell and Netgear (sFlow); Huawei (NetStream); Alcatel-Lucent (Cflow); and Ericsson (Rflow).

NetFlow-enabled devices export traffic statistics as NetFlow records. Using a NetFlow Collector, these NetFlow records can be intercepted and processed. A NetFlow Analyzer/Aggregator can then be used to analyze the processed data and provide actionable information to administrators on traffic flow, source, destination, and traffic volume.

eG Enterprise supports Netflow monitoring. By providing proprietary tools for collecting, processing, and analyzing NetFlow records transmitted by NetFlow-enabled devices in the environment, eG Enterprise delivers in-depth insights into traffic sources, destinations, applications/protocols engaged in network conversations, the volume of data exchanged over the network, and bandwidth used. With the help of this information, administrators can find quick and accurate answers to the following performance queries:

- Which are the top sources in terms of traffic volume or bandwidth usage?
- Which are the top destinations in terms of traffic volume or bandwidth usage?
- Which are the top applications/protocols in terms of traffic volume or bandwidth usage?
- Which two hosts are engaged in a bandwidth-intensive conversation over the network?
- Is any site seeing abnormally high data traffic?
- Are key sources and destinations consuming bandwidth optimally?

Let us begin by understanding how the eG Enterprise performs NetFlow monitoring.

1.1 How eG Enterprise Performs NetFlow Monitoring

To intercept and process NetFlow records exported by NetFlow-enabled devices, eG Enterprise offers an eG NetFlow Collector. The collector runs as a Windows service. Once started, the collector service starts listening for NetFlow records on UDP port 9996 (by default). The NetFlow-enabled device in the target infrastructure should be configured to export NetFlow records to the collector. The collector receives these records, processes them, and stores the processed data in binary files.

To analyze and aggregate the data stored in the binary files, the eG external agent monitoring the NetFlow-enabled device serves as the eG NetFlow Analyzer/Aggregator/Collector. This agent periodically reads these binary files, pulls statistics on netflow, processes/aggregates these statistics on the basis of interfaces, sources, destinations, applications, conversations, sites, etc., and reports the aggregated data to the eG manager. The eG manager then stores this information in the eG database.

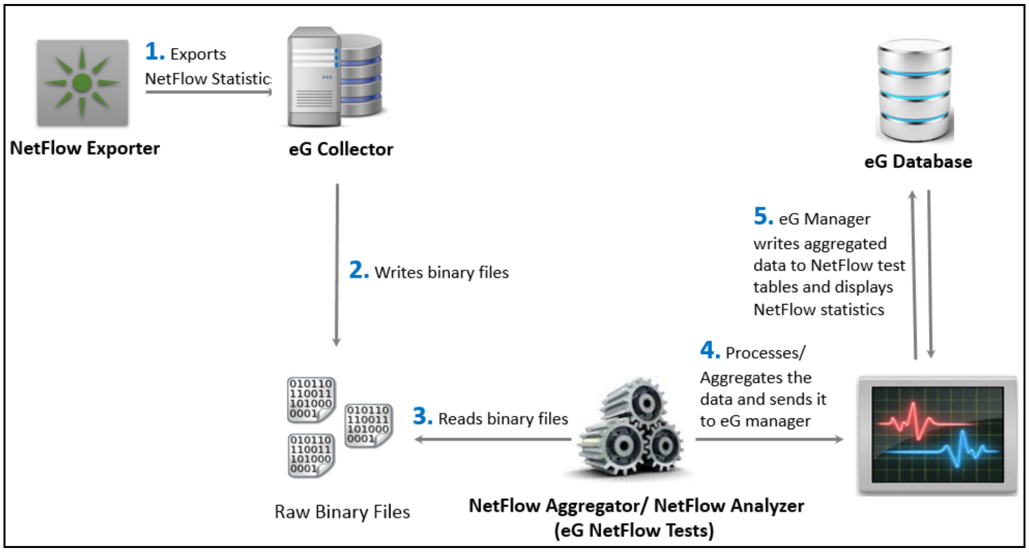
1. Exports NetFlow Statistic
NetFlow Exporter
eG Collector
eG Database
2. Writes binary files
5. eG Manager writes aggregated data to NetFlow test tables and displays NetFlow statistics
4. Processes/ Aggregates the data and sends it to eG manager
3. Reads binary files
010110 110011 101000 0001
Raw Binary Files
NetFlow Aggregator/ NetFlow Analyzer (eG NetFlow Tests)

Figure 1.1: How the eG NetFlow Collector Works

For a NetFlow-enabled device, the eG manager also presents real-time metrics on traffic and bandwidth in the eG monitoring console using a specialized NetFlow Device monitoring model. If abnormalities are spotted during netflow analysis, alerts are generated on this model.

Note:

eG Enterprise 6.3 supports collection of NetFlow v9 flow records from NetFlow-enabled devices (routers, switches, etc.).

The chapter that follows will deep-dive into how to configure NetFlow monitoring using eG Enterprise.

1.2 Licensing

NetFlow monitoring by eG Enterprise is licensed by the number of eG External Agents used for collecting flow data from NetFlow-enabled devices. Licensing is NOT restricted by the number of devices/interfaces exporting flow data.

Each eG External Agent includes one NetFlow Collector, which supports collection of up to 20,000 flows/second.

Chapter 2: Setting Up NetFlow Monitoring

The broad steps towards setting up NetFlow Monitoring are as follows:

1. Using the eG management console, manage the NetFlow- enabled device that will be generating NetFlow records, and assign an external agent to it.
2. Deploy the eG NetFlow Collector on the same system that hosts the external agent assigned to the NetFlow device at step 1 .
3. Open the UDP port (9996, by default) on the collector host, so that the collector can intercept and process NetFlow records; then, start the collector.
4. Configure the eG external agent to run tests that will read the binary files created by the collector and pull relevant data on network traffic and bandwidth usage.

This chapter discusses each of these steps in detail.

2.1 Managing a NetFlow Device in eG Enterprise

To do this, follow the steps below:

1. Login to the eG administrative interface.
2. Follow the Infrastructure -> Components -> Add/Modify menu sequence. In the page that appears next, select *NetFlow Device* as the **Component type** and click the **Add New Component** button to add a new component of that type. Figure 2.1 will then appear.

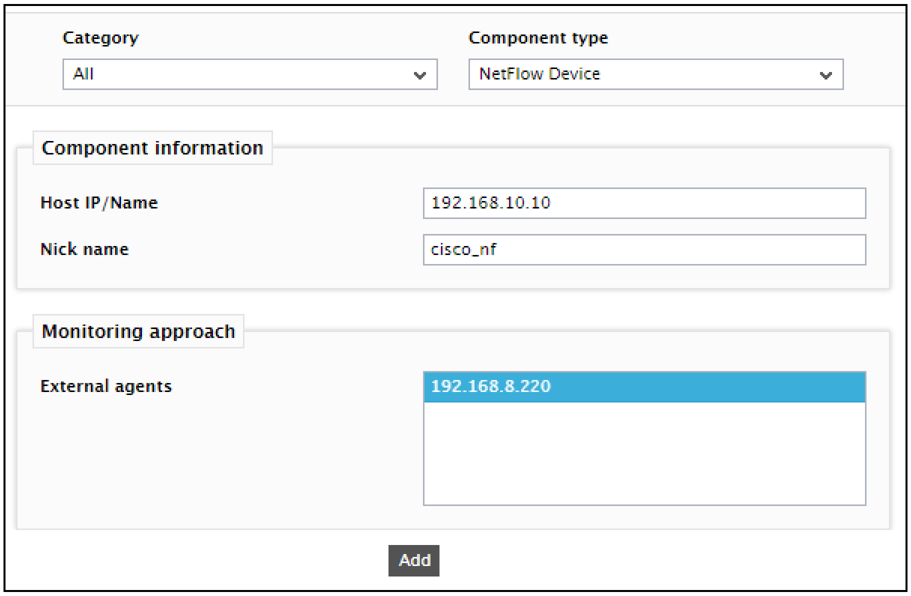

Category: All

Component type: NetFlow Device

Component information

Host IP/Name: 192.168.10.10

Nick name: cisco_nf

Monitoring approach

External agents: 192.168.8.220

Figure 2.1: Adding a NetFlow Device

3. Specify the IP address and **Nick name** of the NetFlow Device to be monitored in Figure 2.1. Then, pick an **External Agent** from Figure 2.1 to assign to the NetFlow Device.

   Note:

   Select an external agent that has been deployed on a Windows host for the purpose of NetFlow monitoring.

4. Finally, click the **Add** button.

2.2 Deploying the eG NetFlow Collector

The next step is to deploy the eG NetFlow Collector. As mentioned already, the collector has to be created as a Windows service. Note that the **collector service has to be created on the same Windows host on which the external agent monitoring the target NetFlow device has been deployed**.

To create the collector service on the external agent host, follow the steps below:

1. Login to the Windows system hosting the eG external agent.
2. Run the command prompt in elevated mode.
3. Go to the <EG_ INSTALL_ DIR>\Netflow\bin directory and execute the **CreateNetflowService.bat** batch file.

4. Successful execution of the batch file will result in the creation of a Windows service named **eGNetFlowAgent**. To confirm the creation of this service, open the **Services** window on the collector host. If you find the **eGNetflowAgent** service displayed therein (see Figure 2.2), it denotes that the collector has been successfully deployed.

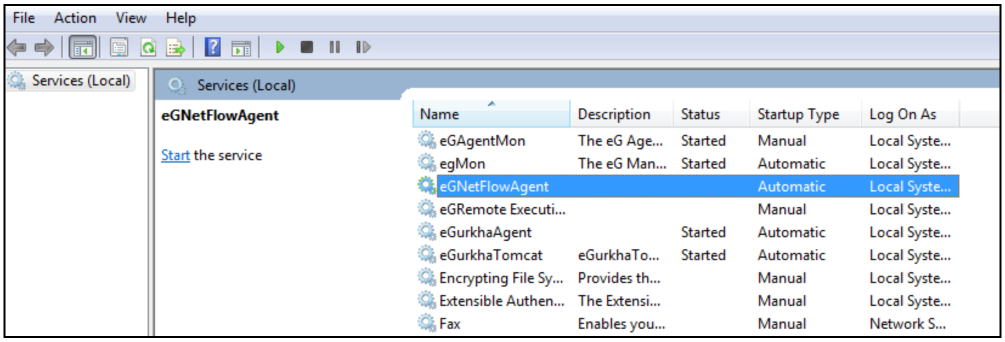

Name	Description	Status	Startup Type	Log On As
eGAgentMon	The eG Age...	Started	Manual	Local Syste...
egMon	The eG Man...	Started	Automatic	Local Syste...
eGNetFlowAgent			Automatic	Local Syste...
eGRemote Executi...			Manual	Local Syste...
eGurkhaAgent		Started	Automatic	Local Syste...
eGurkhaTomcat	eGurkhaTo...	Started	Automatic	Local Syste...
Encrypting File Sy...	Provides th...		Manual	Local Syste...
Extensible Authen...	The Extensi...		Manual	Local Syste...
Fax	Enables you...		Manual	Network S...

Figure 2.2: The eGNetflowAgent service displayed in the Services window

5. By default, the collector listens on UDP port 9996 for NetFlow records from the managed NetFlow device. To make sure that the NetFlow device communicates with the collector via this port, make sure you open this port on the firewall.

6. Once this is done, proceed to start the collector service. For this, right- click on the **eGNetFlowAgent** service in Figure 2.2 and select the **Start** option from the shortcut menu that appears.

7. On the other hand, if the default UDP port 9996 is already in use, then you will first have to change the listening port of the collector. In this case therefore, skip step 6 above; instead, follow the steps detailed in the Configuring the eG NetFlow Collector to Receive NetFlow Records topic.

2.3 Configuring the eG NetFlow Collector to Receive NetFlow Records

By default, the collector listens for NetFlow records on UDP port 9996. **Make sure you open that port on the firewall,** so that the managed NetFlow device is able to communicate with the collector via that port. If the default port is already in use in your environment, then you can change the listening port of the collector. For this, do the following:

1. Login to the system hosting the eG NetFlow Collector.

2. Edit the **NetFlow.properties** file in the <EG_INSTALL_DIR>\Netflow\config directory.

3. Look for the **net.bind.port** entry in that file. This will be set to *9996* by default. Change the port number against that entry and save the file.

4. Finally, start the collector service. For that open the Services window, right-click on the **eGNetFlowAgent** service therein, and select the **Start** option from the shortcut menu that appears.

2.4 Configuring the eG External Agent to Read Binary Files and Report Metrics

Now that the collector service is created and configured, proceed to configure the eG external agent to read the binary files in which the collector has stored the processed netflow data and pull metrics of interest. For this purpose, you need to configure the eG external agent to periodically run different netflow tests. To achieve this, follow the steps below:

1. Login to the eG administrative interface and click **Signout** to exit the interface.

2. Figure 2.3 will then appear listing all the tests that are still to be configured for the managed NetFlow device.

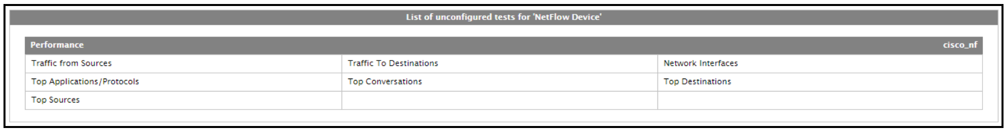

List of unconfigured tests for 'NetFlow Device'

Performance		cisco_nf
Traffic from Sources	Traffic To Destinations	Network Interfaces
Top Applications/Protocols	Top Conversations	Top Destinations
Top Sources		

Figure 2.3: List of unconfigured tests for the managed NetFlow device

3. Click on any test in Figure 2.3 to configure it. For instance, let us say you click on the **Top Sources** test inFigure 2.3. This will open Figure 2.4.

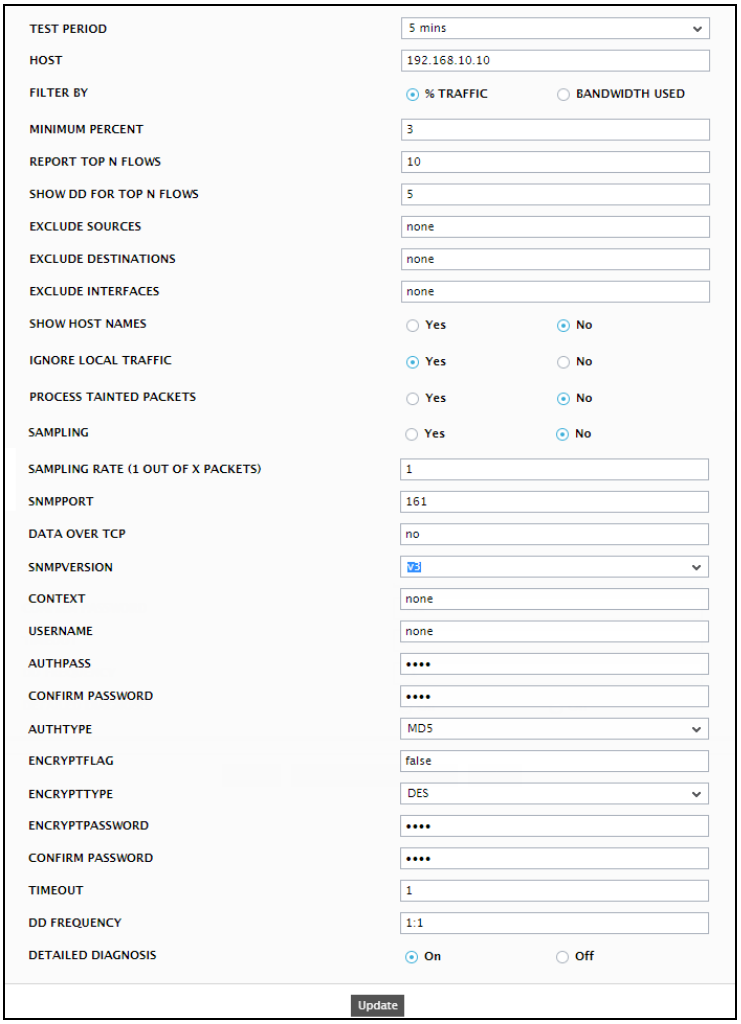

TEST PERIOD	5 mins	
HOST	192.168.10.10	
FILTER BY	● % TRAFFIC	○ BANDWIDTH USED
MINIMUM PERCENT	3	
REPORT TOP N FLOWS	10	
SHOW DD FOR TOP N FLOWS	5	
EXCLUDE SOURCES	none	
EXCLUDE DESTINATIONS	none	
EXCLUDE INTERFACES	none	
SHOW HOST NAMES	○ Yes	● No
IGNORE LOCAL TRAFFIC	● Yes	○ No
PROCESS TAINTED PACKETS	○ Yes	● No
SAMPLING	○ Yes	● No
SAMPLING RATE (1 OUT OF X PACKETS)	1	
SNMPPORT	161	
DATA OVER TCP	no	
SNMPVERSION	v3	
CONTEXT	none	
USERNAME	none	
AUTHPASS	••••	
CONFIRM PASSWORD	••••	
AUTHTYPE	MD5	
ENCRYPTFLAG	false	
ENCRYPTTYPE	DES	
ENCRYPTPASSWORD	••••	
CONFIRM PASSWORD	••••	
TIMEOUT	1	
DD FREQUENCY	1:1	
DETAILED DIAGNOSIS	● On	○ Off

Figure 2.4: Configuring the Top Sources test

4. To know how to configure the **Top Sources** test or any test mapped to the NetFlow Device component, refer to the Monitoring NetFlow topic.

Chapter 3: Monitoring NetFlow

eG Enterprise provides a specialized *NetFlow Device* monitoring model for monitoring NetFlow.

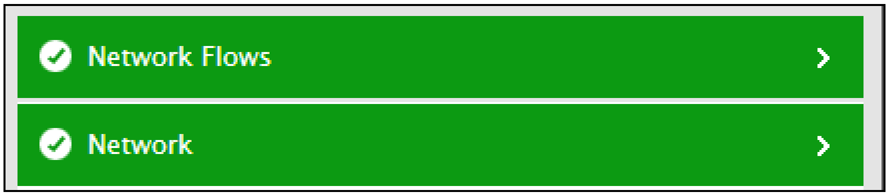

Figure 3.1: The layer model of the NetFlow Device

Tests mapped to the Network layer of this model measure the availability and responsiveness of the managed NetFlow Device. Tests mapped to the Network Flows record read the data stored by the collector in binary files and reports statistics on top flows. This chapter discusses the tests mapped to the Network Flows layer only.

To get up, close with each test mapped to the Network Flows layer, click on the test links below:

- Top Sources By Interfaces Test
- Top Destinations By Interfaces Test
- Top Applications/Protocols Test
- Top Conversations Test
- Key Sites Test
- Traffic from Sources Test
- Traffic to Destinations Test

3.1 The Network Flows Layer

The tests mapped to this layer read the aggregated netflow statistics that the collector stores in the binary files, groups these statistics by network interface, and for each interface, reports the network traffic and bandwidth usage of:

- Top-N sources
- Top-N destinations
- Top-N conversations
- Top-N applications/protocols

You can also optionally configure this layer to report on the network activity of key sites and specific sources/destinations.

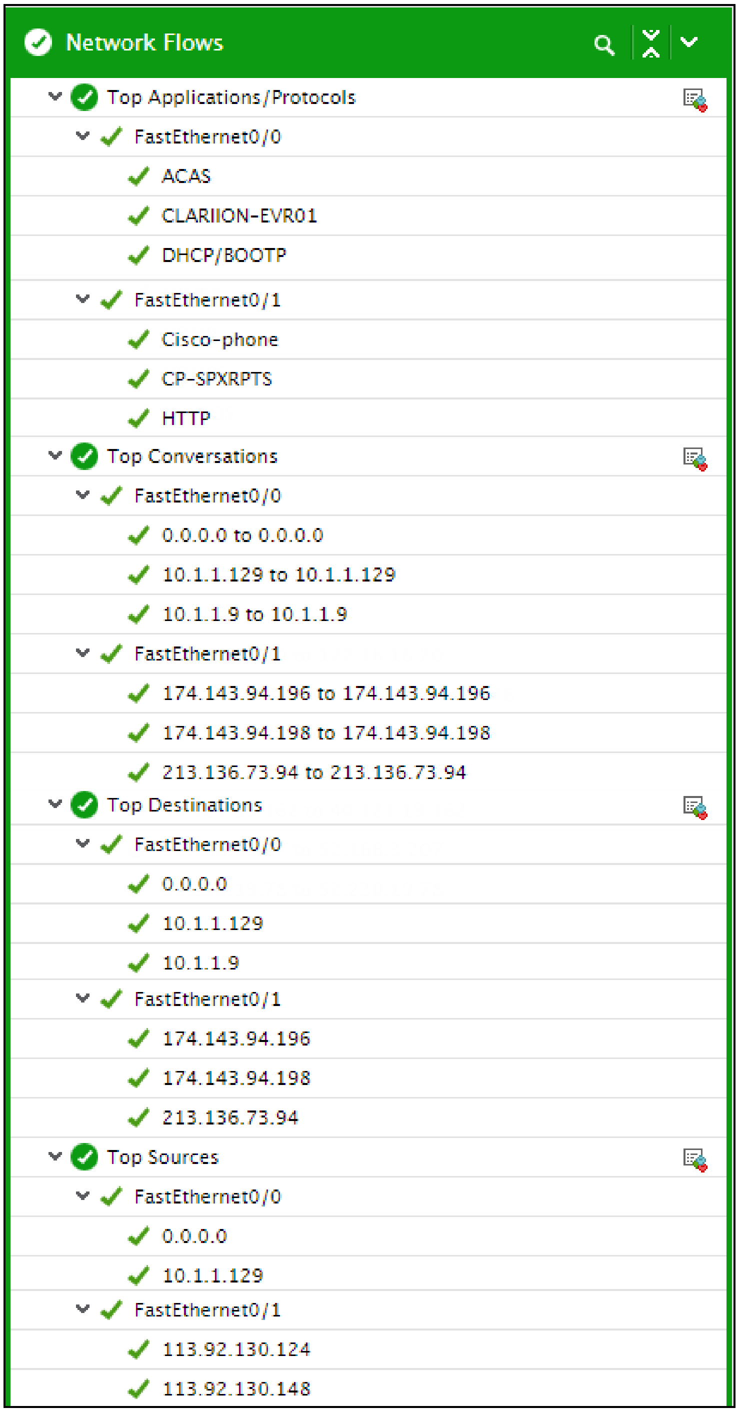

Network Flows

- Top Applications/Protocols
  - FastEthernet0/0
    - ACAS
    - CLARIION-EVR01
    - DHCP/BOOTP
  - FastEthernet0/1
    - Cisco-phone
    - CP-SPXRPTS
    - HTTP
- Top Conversations
  - FastEthernet0/0
    - 0.0.0.0 to 0.0.0.0
    - 10.1.1.129 to 10.1.1.129
    - 10.1.1.9 to 10.1.1.9
  - FastEthernet0/1
    - 174.143.94.196 to 174.143.94.196
    - 174.143.94.198 to 174.143.94.198
    - 213.136.73.94 to 213.136.73.94
- Top Destinations
  - FastEthernet0/0
    - 0.0.0.0
    - 10.1.1.129
    - 10.1.1.9
  - FastEthernet0/1
    - 174.143.94.196
    - 174.143.94.198
    - 213.136.73.94
- Top Sources
  - FastEthernet0/0
    - 0.0.0.0
    - 10.1.1.129
  - FastEthernet0/1
    - 113.92.130.124
    - 113.92.130.148

Figure 3.2: The tests mapped to the Network Flows layer

To know what each test does and what metrics it reports, click on the corresponding test link below:

- Top Sources By Interfaces Test
- Top Destinations By Interfaces Test
- Top Applications/Protocols Test
- Top Conversations Test
- Key Sites Test
- Traffic from Sources Test
- Traffic to Destinations Test

3.1.1 Top Sources By Interfaces Test

If the Network Interfaces test reveals that a specific interface is handling an abnormally high volume of traffic or is consuming bandwidth resources excessively, you may want to know where this traffic is originating from. This is where the **Top Sources** test helps!

The **Top Sources** test reveals those hosts whose interactions with other hosts in the environment are resulting in large volumes of data being transmitted over a network interface. In the event of a network slowdown, you can use this test to accurately identify hosts whose current network activities are 'suspect' - i.e., you can isolate those hosts that may be engaged in bandwidth-intensive transactions with other hosts, and could hence be contributing to the slowdown.

Target of the test : A Netflow Device

Agent deploying the test : An external agent

Outputs of the test : One set of results for each of the top-n sources (in terms of traffic volume) communicating through each interface of the target network device.

First level descriptor: Interface

Second level descriptor: Source host's IP address

Metrics are also reported for an **Others** descriptor. Netflow statistics pertaining to all sources that are not the top-n sources are aggregated and presented against the **Others** descriptor for each interface.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Filter By, Minimum Flow Percent, Report Top N Flows, Show DD for Top N Flows	By default, for each interface that the target Netflow device supports, this test reports metrics for only the top-10 sources that handle over 3% of the total traffic on that interface. Also, by default, the test will report detailed diagnostics for only the top-5 sources that handle over 3% of the total traffic on that interface. This default setting enables network administrators to focus on the abnormal flows alone. This behavior is governed by the default setting of the Filter By, Minimum Flow Percent, Report Top N Flows, and Show DD for Top N Flows parameters. By default, the Filter By option is set to %Traffic and the Minimum Flow Percent is set to 3 (%). You can increase or decrease the Minimum Flow Percent depending upon the level of traffic that you deem as abnormal. Also, by default, the Report Top N Flows is set to 10 and Show DD for Top N Flows is set to 5. You can increase or decrease the value of these two parameters depending upon the level of visibility you require. You can also override the Filter By default setting and have this test report real-time metrics and detailed diagnostics for only those sources (per interface) that use over a configured percentage of bandwith resources when transmitting data. For this, set Filter By to Bandwidth Used and configure a bandwidth usage limit (in %) against Minimum Flow Percent. This way, administrators can focus on only those sources that generate bandwidth-intensive traffic.
Exclude Sources	Provide a comma-separated list of sources that you want excluded from monitoring. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of source IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of source IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** Sources can be represented using IP addresses only and not host names.
Exclude Destinations	Provide a comma-separated list of destinations that you want excluded from monitoring. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of destination IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of source IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:**

Parameter	Description
	Destinations can be represented using IP addresses only and not host names.
Exclude Interfaces	Provide a comma-separated list of interfaces that you want excluded from monitoring. For example, your specification can be, FastEtherNet0/0, FastEthernet0/1,FastEtherNet0/2. You can even provide patterns of interfaces as a comma-separated list - for instance, *ethernet*,Fast*.
Show Host Names	This test captures statistics on traffic that originates from source hosts, where each host is by default represented by its IP address in the eG monitoring console. Accordingly, this flag is set to **No** by default. You can set this flag to **Yes** so that the names of the individual hosts are displayed in the eG monitoring console instead of their IP addresses.
Ignore Local Traffic	By default, this flag is set to **Yes**, indicating that the test will ignore the sources of all the intranet traffic on the router. If you want the test to report metrics pertaining to the sources of local traffic as well, set this flag to **No**.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when identifying top sources and reporting their traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**. Once this is done, then each interface will process only 1 packet out of a configured number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling Rate you specify.
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

Parameter	Description
	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 SNMP. 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 the Username 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, this flag is set to **No** by

Parameter	Description
	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**.
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

Parameter	Description
	measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total data exchange	Indicates the total amount of data transmitted and received by this source during the last measurement period.	KB	Compare the value of this measure across sources to identify which source host is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) that originated from this source, and the amount of data transacted in bytes and packets in every flow.
Total packets exchange	Indicates the total number of packets transmitted and received by this source during the last measurement period.	Packets	Compare the value of this measure across sources to identify which source host is contributing to the high level of network traffic
Data exchange rate	Indicates the rate at which this source transmitted/received data.	Kbps	
Packets exchange rate	Indicates the rate at which this source transmitted/received packets.	Kbps	
Total traffic from this source	Indicates what percentage of the total traffic on this interface was from this source.	Percent	A value close to 100% for this measure indicates that traffic from this source is imposing the maximum load on the network. If users complain of a latent network, you can compare the value of this measure

Measurement	Description	Measurement Unit	Interpretation
			across sources to accurately identify the source that is responsible for any congestion on the network.
Total bandwidth utilization	Indicates the bandwidth utilized by this source.	Percent	A value close to 100% indicates that this source is indicative of excessive bandwidth utilization by this source when transmitting data. If users complain of a latent network, you can compare the value of this measure across sources to accurately identify the source that is responsible for any congestion on the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this source.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on this interface pertains to this source.	Percent	
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface when receiving data from this source and receiving data for this source from destinations.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface when transmitting data to this source and when transmitting data for this source to destinations.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the amount of data received by this source.	KB	Compare the value of this measure across sources to know which source is receiving maximum data over this interface.
Data sent	Indicates the amount of data sent by this source.	KB	Compare the value of this measure across sources to know which source is transmitting maximum data over this interface.
Packets received	Indicates the number of packets sent by this source.	Number	
Packets sent	Indicates the number of packets received by this source.	Number	
Data received rate	Indicates the rate at which data is received by this source.	Kbps	If the value of this measure consistently drops for this source, it could indicate an incoming traffic congestion.
Data transmitted rate	Indicates the rate at which data is sent by this source.	Kbps	If the value of this measure consistently drops for this source, it could indicate an outgoing traffic congestion.
Packets received rate	Indicates the rate at which packets are received by this source.	Packets/Sec	If the value of this measure consistently drops for this source, it could indicate an incoming traffic congestion.
Packets transmitted rate	Indicates the rate at which packets are transmitted by this source.	Packets/Sec	If the value of this measure consistently drops for this source, it could indicate an outgoing traffic congestion.

Use the detailed diagnosis of the *Total data exchange* measure to determine the top netflows (in terms of the volume of data transacted) that originated from a particular source, and the amount of data and packets received. With the help of this detailed diagnosis, you can quickly compare the top netflows, know which netflow generated the maximum traffic, and figure out which destination that traffic was headed to. Once the problem destination is isolated, you can then investigate why traffic to that destination was high - is it because of the type of application executing on that destination?

(eg., an online game or a movie that would typically consume a lot of bandwidth), or is it because of a poor network line connecting the source and the destination?

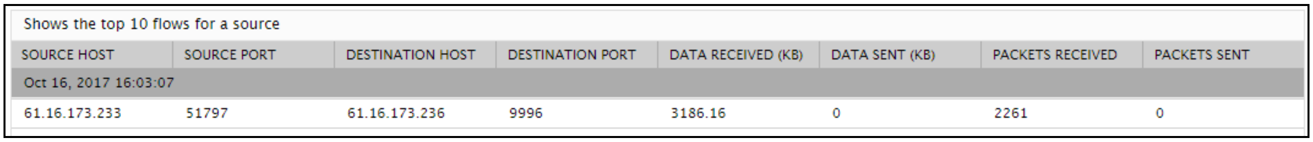

Shows the top 10 flows for a source

SOURCE HOST	SOURCE PORT	DESTINATION HOST	DESTINATION PORT	DATA RECEIVED (KB)	DATA SENT (KB)	PACKETS RECEIVED	PACKETS SENT
Oct 16, 2017 16:03:07							
61.16.173.233	51797	61.16.173.236	9996	3186.16	0	2261	0

Figure 3.3: The detailed diagnosis of the Total data exchange measure

3.1.2 Top Destinations By Interfaces Test

If the Network Interfaces test reveals that a specific interface is handling an abnormally high volume of traffic or is consuming bandwidth resources excessively, you may want to know where this traffic is headed. This is where the **Top Destinations** test helps!

This test discovers the top-n destinations (in terms of volume of traffic) of the netflows through each interface on a target device, and for each destination, reports the data traffic (in bytes and packets) leading to or coming from that destination. This way, the test points you to the destinations that generate the maximum traffic. The bandwidth used by each destination when receiving / transmitting data is also reported, so that you can quickly identify those destinations for which traffic has been consistently bandwidth-intensive.

Target of the test : A Netflow Device

Agent deploying the test : An external agent

Outputs of the test : One set of results for every top-n destinations (in terms of traffic volume) on each interface supported by the target device

First-level descriptor: Interface name

Second-level descriptor: Destination

Metrics are also reported for an **Others** descriptor. Netflow statistics pertaining to all destinations that are not the top-n destinations are aggregated and presented against the **Others** descriptor for each interface.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed

Parameter	Description
Host	The host for which the test is to be configured.
Filter By, Minimum Flow Percent, Report Top N Flows, Show DD for Top N Flows	By default, for each interface that the target Netflow device supports, this test reports metrics for only the top-10 destinations that handle over 3% of the total traffic on that interface. Also, by default, the test will report detailed diagnostics for only the top-5 destinations that handle over 3% of the total traffic on that interface. This default setting enables network administrators to focus on the abnormal flows alone. This behavior is governed by the default setting of the Filter By, Minimum Flow Percent, Report Top N Flows, and Show DD for Top N Flows parameters. By default, the Filter By option is set to %Traffic and the Minimum Flow Percent is set to 3 (%). You can increase or decrease the Minimum Flow Percent depending upon the level of traffic that you deem as abnormal. Also, by default, the Report Top N Flows is set to 10 and Show DD for Top N Flows is set to 5. You can increase or decrease the value of these two parameters depending upon the level of visibility you require. You can also override the Filter By default setting and have this test report real-time metrics and detailed diagnostics for only those destinations (per interface) that use over a configured percentage of bandwidth resources. For this, set Filter By to Bandwidth Used and configure a bandwidth usage limit (in %) against Minimum Flow Percent. This way, administrators can focus on only those destinations that generate bandwidth-intensive traffic.
Exclude Sources	Provide a comma-separated list of sources that you want excluded from monitoring. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of source IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of source IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** Sources can be represented using IP addresses only and not host names.
Exclude Destinations	Provide a comma-separated list of destinations that you want excluded from monitoring. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of destination IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of destination IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** Destinations can be represented using IP addresses only and not host names.
Exclude Interfaces	Provide a comma-separated list of interfaces that you want excluded from monitoring. For example, your specification can be, FastEtherNet0/0,

Parameter	Description
	FastEthernet0/1,FastEtherNet0/2. You can even provide patterns of interfaces as a comma-separated list - for instance, **ethernet*,Fast**.
Show Host Names	This test captures statistics on traffic to destination hosts, where each host is by default represented by its IP address in the eG monitoring console. Accordingly, this flag is set to **No** by default. You can set this flag to **Yes** so that the names of the individual hosts are displayed in the eG monitoring console instead of their IP addresses.
Ignore Local Traffic	By default, this flag is set to **Yes**, indicating that the test will ignore the destinations of all the intranet traffic on the target NetFlow device. If you want the test to report metrics pertaining to the destinations of local traffic as well, set this flag to **No**.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when identifying top sources and reporting their traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**. Once this is done, then each interface will process only 1 packet out of a configured number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling Rate you specify.
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.

Parameter	Description
Username	This parameter appears only when **v3** is selected as the SNMP. 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 the Username 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, 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:

Parameter	Description
	• **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**.
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 exchange	Indicates the total amount of data transmitted and received by this destination during the last measurement period.	KB	Compare the value of this measure across destinations to identify which destination host is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) to this destination, and the amount of data transacted in bytes and packets in every flow. With the help of this detailed diagnosis, you can quickly compare the top netflows, know which netflow generated the maximum traffic, and figure out to which destination that traffic was headed. Once the problem destination is isolated, you can then investigate why traffic to that destination is high.
Total packets exchange	Indicates the total number of packets transmitted and received by this destination during the last measurement period.	Packets	Compare the value of this measure across sources to identify which destination host is contributing to the high level of network traffic
Data exchange rate	Indicates the rate at which this destination transmitted/received data.	Kbps	
Packets exchange rate	Indicates the rate at which this destination transmitted/received packets.	Kbps	
Total traffic to this destination	Indicates what percentage of the total traffic on this interface was to this destination.	Percent	A value close to 100% for this measure indicates that traffic to this destination is imposing the maximum load on the network. If users complain of a latent network, you can

Measurement	Description	Measurement Unit	Interpretation
			compare the value of this measure across destinations to accurately identify the destination that is responsible for any congestion on the network.
Total bandwidth utilization	Indicates the bandwidth utilized by this destination.	Percent	A value close to 100% is indicative of excessive bandwidth utilization by this destination when receiving data. If users complain of a latent network, you can compare the value of this measure across destinations to accurately identify the destination that is responsible for any congestion on the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this destination.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on this interface pertains to this destination.	Percent	
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface when receiving data for this destination from different sources and when receiving data from this destination.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface when transmitting data to this destination and when transmitting data for this destination to a source.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.

Measurement	Description	Measurement Unit	Interpretation
Data received	Indicates the amount of data received by this destination.	KB	Compare the value of this measure across destinations to know which destination is receiving maximum data over this interface.
Data sent	Indicates the amount of data sent by this destination.	KB	Compare the value of this measure across destinations to know which destination is transmitting maximum data over this interface.
Packets received	Indicates the number of packets sent by this destination.	Number	
Packets sent	Indicates the number of packets received by this destination.	Number	
Data received rate	Indicates the rate at which data is received by this destination.	Kbps	If the value of this measure consistently drops for this destination, it could indicate an incoming traffic congestion.
Data transmitted rate	Indicates the rate at which data is sent by this destination.	Kbps	If the value of this measure consistently drops for this destination, it could indicate an outgoing traffic congestion.
Packets received rate	Indicates the rate at which packets are received by this destination.	Packets/Sec	If the value of this measure consistently drops for this destination, it could indicate an incoming traffic congestion.
Packets transmitted rate	Indicates the rate at which packets are transmitted by this destination.	Packets/Sec	If the value of this measure consistently drops for this destination, it could indicate an outgoing traffic congestion.

3.1.3 Top Applications/Protocols Test

One of the key challenges that network administrators face every day is identifying the applications/protocols that are the leading consumers of bandwidth resources and the root-cause for traffic congestion. The Top Applications/Protocols test helps administrators tackle this challenge!

This test automatically discovers the top-n applications/protocols in terms of the volume of traffic they generate, and reports the following:

- how much data is transmitted/received by each application/protocol through every interface on the target Netflow device;
- how much bandwidth resources are used by each application/protocol through every interface on the target Netflow device;

With the help of this information, administrators can quickly and accurately isolate those applications/protocols that are consuming maximum bandwidth and those that could probably be choking the network link. Sizing decisions can be taken based on the metrics reported.

Target of the test : A Netflow Device

Agent deploying the test : An external agent

Outputs of the test : One set of results for every top-n application/protocol (in terms of traffic volume) communicating with each interface supported by the target device

First-level descriptor: Interface name

Second-level descriptor: Application/Protocol

Metrics are also reported for an **Others** descriptor. Netflow statistics pertaining to all applications/protocols that are not the top-n applications/protocols are aggregated and presented against the **Others** descriptor for each interface.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed.
Host	The host for which the test is to be configured.
Filter By, Minimum Flow Percent, Report Top N Flows, Show DD for Top N Flows	By default, for each interface that the target Netflow device supports, this test reports metrics for only the top-10 applications/protocols that handle over 3% of the total traffic on that interface. Also, by default, the test will report detailed diagnostics for only the top-5 applications/protocols that handle over 3% of the total traffic on that interface. This default setting enables network administrators to focus on the abnormal flows alone. This behavior is governed by the default setting of the Filter By, Minimum Flow Percent, Report Top N Flows, and Show DD for Top N Flows parameters. By default, the Filter By option is set to %Traffic and the Minimum Flow Percent is set to 3 (%). You can increase or decrease the Minimum Flow Percent depending upon the level of

Parameter	Description
	traffic that you deem as abnormal. Also, by default, the Report Top N Flows is set to 10 and Show DD for Top N Flows is set to 5. You can increase or decrease the value of these two parameters depending upon the level of visibility you require. You can also override the Filter By default setting and have this test report real-time metrics and detailed diagnostics for only those applications/protocols (per interface) that use over a configured percentage of bandwidth resources. For this, set Filter By to Bandwidth Used and configure a bandwidth usage limit (in %) against Minimum Flow Percent. This way, administrators can focus on only those applications/protocols that generate bandwidth-intensive traffic.
Exclude Interfaces	Provide a comma-separated list of interfaces that you want excluded from monitoring. For example, your specification can be, FastEtherNet0/0, FastEthernet0/1,FastEtherNet0/2. You can even provide patterns of interfaces as a comma-separated list - for instance, *ethernet*,Fast*.
Show Host Names	This test captures statistics on traffic to/from applications, where each application is by default represented by its IP address in the eG monitoring console. Accordingly, this flag is set to **No** by default. You can set this flag to **Yes** so that the names of the individual hosts are displayed in the eG monitoring console instead of their IP addresses.
Ignore Local Traffic	By default, this flag is set to **Yes**, indicating that the test will ignore the applications/protocols responsible for all the intranet traffic on the target NetFlow device. If you want the test to report metrics pertaining to the applications/protocols responsible for local traffic as well, set this flag to **No**.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when identifying top sources and reporting their traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**. Once this is done, then each interface will process only 1 packet out of a configured

Parameter	Description
	number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling Rate you specify.
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 SNMP. 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 the Username 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

Parameter	Description
	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, 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**.
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

Parameter	Description
	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 exchange	Indicates the total amount of data transmitted and received by this application/protocol during the last measurement period.	KB	Compare the value of this measure across applications/protocols to identify which application/protocol is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) from/to this application/protocol, and the amount of data transacted in bytes and packets in every flow. With the help of this detailed diagnosis, you can quickly compare the top netflows, know which netflow generated the maximum traffic, and figure out between which source and destination that flow occurred. Once the problem conversation is isolated, you can then investigate why traffic on that conversation was high.
Total packets exchange	Indicates the total number of packets transmitted and received by this	Packets	Compare the value of this measure across applications/protocols to identify which application/protocol is

Measurement	Description	Measurement Unit	Interpretation
	application/protocol during the last measurement period.		contributing to the high level of network traffic
Data exchange rate	Indicates the rate at which this application/protocol transmitted/received data.	Kbps	
Packets exchange rate	Indicates the rate at which this application/protocol transmitted/received packets.	Kbps	
Total traffic by this application	Indicates what percentage of the total traffic on this interface was generated by this application/protocol.	Percent	A value close to 100% for this measure indicates that traffic generated by this application/protocol is imposing the maximum load on the network. If users complain of a latent network, you can compare the value of this measure across applications/protocols to accurately identify the application/protocol that is responsible for any congestion on the network.
Total bandwidth utilization	Indicates the bandwidth utilized by this application/protocol.	Percent	A value close to 100% is indicative of excessive bandwidth utilization by this application/protocol when receiving data. If users complain of a latent network, you can compare the value of this measure across applications/protocols to accurately identify the application/protocol that is responsible for any congestion on the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this application/protocol.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on	Percent	

Measurement	Description	Measurement Unit	Interpretation
	this interface pertains to this application/protocol.		
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface when receiving data for this application/protocol from different sources and when receiving data from this application/protocol.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface when transmitting data to this application/protocol and when transmitting data for this application/protocol.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.
Data received	Indicates the amount of data received by this application/protocol.	KB	Compare the value of this measure across applications/protocols to know which application/protocol is receiving maximum data over this interface.
Data sent	Indicates the amount of data sent by this application/protocol.	KB	Compare the value of this measure across applications/protocolsto know which application/protocol is transmitting maximum data over this interface.
Packets received	Indicates the number of packets sent by this application/protocol.	Number	
Packets sent	Indicates the number of packets received by this application/protocol.	Number	
Data received rate	Indicates the rate at which data is received by this application/protocol.	Kbps	If the value of this measure consistently drops for this application/protocol, it could indicate an incoming traffic congestion.

Measurement	Description	Measurement Unit	Interpretation
Data transmitted rate	Indicates the rate at which data is sent by this application/protocol.	Kbps	If the value of this measure consistently drops for this application/protocol, it could indicate an outgoing traffic congestion.
Packets received rate	Indicates the rate at which packets are received by this application/protocol.	Packets/Sec	If the value of this measure consistently drops for this application/protocol, it could indicate an incoming traffic congestion.
Packets transmitted rate	Indicates the rate at which packets are transmitted by this application/protocol.	Packets/Sec	If the value of this measure consistently drops for this application/protocol, it could indicate an outgoing traffic congestion.

3.1.4 Top Conversations Test

Bandwidth-intensive conversations between two hosts over the network is one of the reasons why users may experience slowness when communicating over the network. To resolve this, network administrators must first keep a close watch on the NetFlow between each pair of hosts that is interacting over the network, study the bandwidth usage of and data exchange between each pair, and identify the pair that is engaged in an bandwidth-intensive conversation. This is what the **Top Conversations** test helps administrators achieve!

This test auto-discovers the top-n conversations in terms of the amount of data exchanged over the network. For each conversation, this test reports the amount of data exchanged between the hosts engaged in that conversation and the bandwidth consumed in the process. This way, the test points administrators to the pair of hosts that are engaged in bandwidth-intensive communication.

Target of the test : A Netflow Device

Agent deploying the test : An external agent

Outputs of the test : One set of results for each of the top-n conversations (in terms of traffic volume) happening over every network interface supported by the target device

First-level descriptor: Interface name

Second-level descriptor: IP addresses of the hosts engaged in an conversation over the network. Format is <*SourceHost*> *to* <*DestinationHost*>

Metrics are also reported for an **Others** descriptor. Netflow statistics pertaining to all conversations that are not the top-n conversations are aggregated and presented against the **Others** descriptor for each interface.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Filter By, Minimum Flow Percent, Report Top N Flows, Show DD for Top N Flows	By default, for each interface that the target Netflow device supports, this test reports metrics for only the top-10 conversations that handle over 3% of the total traffic on that interface. Also, by default, the test will report detailed diagnostics for only the top-5 conversations that handle over 3% of the total traffic on that interface. This default setting enables network administrators to focus on the abnormal flows alone. This behavior is governed by the default setting of the Filter By, Minimum Flow Percent, Report Top N Flows, and Show DD for Top N Flows parameters. By default, the Filter By option is set to %Traffic and the Minimum Flow Percent is set to 3 (%). You can increase or decrease the Minimum Flow Percent depending upon the level of traffic that you deem as abnormal. Also, by default, the Report Top N Flows is set to 10 and Show DD for Top N Flows is set to 5. You can increase or decrease the value of these two parameters depending upon the level of visibility you require. You can also override the Filter By default setting and have this test report real-time metrics and detailed diagnostics for only those conversations (per interface) that use over a configured percentage of bandwidth resources. For this, set Filter By to Bandwidth Used and configure a bandwidth usage limit (in %) against Minimum Flow Percent. This way, administrators can focus on only those conversations that generate bandwidth-intensive traffic.
Exclude Sources	Provide a comma-separated list of sources that you want excluded from monitoring. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of source IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of source IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** • Sources can be represented using IP addresses only and not host names. • All conversations originating from the sources specified against Exclude Sources will be ignored during monitoring.

Parameter	Description
Exclude Destinations	Provide a comma-separated list of destinations that you want excluded from monitoring. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of destination IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of source IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** • Destinations can be represented using IP addresses only and not host names. • All conversations to the destinations specified against Exclude Destinations will be ignored during monitoring.
Exclude Interfaces	Provide a comma-separated list of interfaces that you want excluded from monitoring. For example, your specification can be, FastEtherNet0/0, FastEthernet0/1,FastEtherNet0/2. You can even provide patterns of interfaces as a comma-separated list - for instance, *ethernet*,Fast*.
Show Host Names	This test captures statistics on traffic to/from applications, where each application is by default represented by its IP address in the eG monitoring console. Accordingly, this flag is set to **No** by default. You can set this flag to **Yes** so that the names of the individual hosts are displayed in the eG monitoring console instead of their IP addresses.
Ignore Local Traffic	By default, this flag is set to **Yes**, indicating that the test will ignore the conversations over the intranet on the target NetFlow device. If you want the test to report metrics pertaining to the conversations over the intranet as well, set this flag to **No**.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when identifying top sources and reporting their traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**.

Parameter	Description
	Once this is done, then each interface will process only 1 packet out of a configured number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling Rate you specify.
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 SNMP. 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 the Username 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.

Parameter	Description
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, 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**.
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.

Parameter	Description
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 exchange	Indicates the total amount of data transmitted and received over this conversation during the last measurement period.	KB	Compare the value of this measure across conversations to identify which conversation is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) between the two hosts engaged in this conversation, and the amount of data transacted in bytes and packets in every flow. With the help of this detailed diagnosis, you can quickly compare the top netflows and know which netflow generated the maximum traffic. Once the problem flow is isolated, you can then investigate why traffic was abnormally high on that flow.
Total packets exchanged	Indicates the total number of packets transmitted and received over this conversation during the last	Packets	Compare the value of this measure across conversations to identify which conversation is contributing to the high level of network traffic.

Measurement	Description	Measurement Unit	Interpretation
	measurement period.		
Data exchange rate	Indicates the rate at which data was transmitted/received over this conversation.	Kbps	
Packets exchange rate	Indicates the rate at which packets were transmitted/received over this conversation.	Kbps	
Total traffic by this conversation	Indicates what percentage of the total traffic on this interface was generated by this conversation.	Percent	A value close to 100% for this measure indicates that traffic generated by this conversation is imposing the maximum load on the network. If users complain of a latent network, you can compare the value of this measure across conversations to accurately identify which two hosts are exchanging a large volume of data over the network and choking it.
Total bandwidth utilization	Indicates the bandwidth utilized by this conversation.	Percent	A value close to 100% is indicative of excessive bandwidth utilization by this conversation. If users complain of a latent network, you can compare the value of this measure across conversations to accurately identify which two hosts are engaged in a bandwidth-intensive conversation over the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this conversation.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on	Percent	

Measurement	Description	Measurement Unit	Interpretation
	this interface pertains to this conversation.		
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface during this conversation.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface during this conversation.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.
Data received	Indicates the amount of data received by this interface during this conversation.	KB	Compare the value of this measure across conversations to know which conversation is imposing the maximum load on the interface.
Data sent	Indicates the amount of data sent by this interface during this conversation.	KB	Compare the value of this measure across conversations to know which conversation is imposing the maximum load on the interface.
Packets received	Indicates the number of packets sent by this interface as part of this conversation.	Number	
Packets sent	Indicates the number of packets received by this interface during this conversation.	Number	
Data received rate	Indicates the rate at which data is received by this interface during this conversation.	Kbps	If the value of this measure consistently drops for this conversation, it could indicate an incoming traffic congestion.
Data transmitted rate	Indicates the rate at which data is sent by this interface during this conversation.	Kbps	If the value of this measure consistently drops for this conversation, it could indicate an outgoing traffic congestion.
Packets received	Indicates the rate at which	Packets/Sec	If the value of this measure

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>rate</td><td>packets are received by this interface during this conversation.</td><td></td><td>consistently drops for this conversation, it could indicate an incoming traffic congestion.</td></tr>
<tr><td>Packets transmitted rate</td><td>Indicates the rate at which packets are transmitted by this interface during this conversation.</td><td>Packets/Sec</td><td>If the value of this measure consistently drops for this conversation, it could indicate an outgoing traffic congestion.</td></tr>
<tr><td>Protocol</td><td>Indicates the protocol used for this conversation.</td><td></td><td>The values that this measure can report and their corresponding numeric values are listed in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>ICMP</td><td>1</td></tr>
<tr><td>IGMP</td><td>2</td></tr>
<tr><td>CGP</td><td>3</td></tr>
<tr><td>IPv4</td><td>4</td></tr>
<tr><td>ST</td><td>5</td></tr>
<tr><td>TCP</td><td>6</td></tr>
<tr><td>CBT</td><td>7</td></tr>
<tr><td>EGP</td><td>8</td></tr>
<tr><td>IGP</td><td>9</td></tr>
<tr><td>BBN-RCC-MON</td><td>10</td></tr>
<tr><td>NVP-II</td><td>11</td></tr>
<tr><td>PUP</td><td>12</td></tr>
<tr><td>ARGUS</td><td>13</td></tr>
<tr><td>EMCON</td><td>14</td></tr>
<tr><td>XNET</td><td>15</td></tr>
<tr><td>CHAOS</td><td>16</td></tr>
<tr><td>UDP</td><td>17</td></tr>
</table>
</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>
<table>
<tr><td>MUX</td><td>18</td></tr>
<tr><td>RDP</td><td>27</td></tr>
<tr><td>IPv6</td><td>41</td></tr>
<tr><td>IPv6-Route</td><td>43</td></tr>
<tr><td>IPv6-Frag</td><td>44</td></tr>
<tr><td>IDRP</td><td>45</td></tr>
<tr><td>RSVP</td><td>46</td></tr>
<tr><td>Encrypt</td><td>47</td></tr>
<tr><td>SWIPE</td><td>53</td></tr>
<tr><td>MOBILE</td><td>55</td></tr>
<tr><td>IPv6-ICMP</td><td>58</td></tr>
<tr><td>IPv6-NoNxt</td><td>59</td></tr>
<tr><td>IPv6-Opts</td><td>60</td></tr>
<tr><td>VISA</td><td>70</td></tr>
<tr><td>PVP</td><td>75</td></tr>
<tr><td>DGP</td><td>86</td></tr>
<tr><td>IPIP</td><td>94</td></tr>
<tr><td>PNNI</td><td>102</td></tr>
<tr><td>UDPLite</td><td>136</td></tr>
</table>
<b>Note:</b>
By default, this test reports the <b>Measure Value</b>s listed above to indicate the protocol used by this conversation. In the graph of this measure however, the same is represented using the numeric equivalents only.
</td></tr>
</table>

3.1.5 Key Sites Test

Studying network traffic to popular / frequently accessed web sites and measuring the bandwidth usage of this traffic is key to fine-tuning firewall policies in an enterprise and understanding the real bandwidth requirement. The **Key Sites** test simplifies this! For each web site that is configured for monitoring, this test reports the amount of data transmitted and received and the bandwidth utilized by that site. Web sites that consistently invite heavy traffic and consume excessive bandwidth can be identified in the process. If such sites are mission-critical business sites/applications, then this information will help you to determine the bandwidth required to ensure the peak performance of the sites and thus enable you to right-size your network. If such sites are inconsequential to your business, then this information will point you to where bandwidth is spent unnecessarily; this in turn will prompt you to initiate measures to control/regulate accesses to such sites.

To configure the web sites that this test should monitor, do the following:

1. Edit the eg_netflow.ini file in the <EG_AGENT_INSTALL_DIR>\agent\config folder (on Windows; on Unix installations of the eG agent, you will find this file in the /opt/egurkha/agent/config folder).
2. In the TOP SITES section of the file, create a sub-section for the managed NetFlow device. The IP address of the target NetFlow device should be the title of that sub-section.The sub-section title should be specified in square brackets. For instance, if you have managed the NetFlow device using the IP address *192.168.10.25* in your IT infrastructure, then, the specification in the eg_netflow.ini file will be:

```
=================================
TOP SITES
=================================
[192.168.10.25]
```

3. In this sub-section, specify the URL of the web sites to be monitored, one after another. Against each site URL, specify a comma-separated list of IP addresses of that web site. For example, if the web site www.xyz.com is associated with the IP addresses, 192.168.10.30, 192.168.10.35, 192.168.10.40, 192.168.10.90, then your specification will be as follows:

```
=================================
TOP SITES
=================================
[192.168.10.25]
www.xyz.com=192.168.10.30,192.168.10.35,192.168.10.40,192.168.10.90
```

   Where a site is associated with a specific range of IP addresses, you can even specify the IP range against the site URL, as shown below:

```
=================================
TOP SITES
=================================
[192.168.10.25]
www.xyz.com=192.168.10.25-192.168.10.45
```

4. Likewise, for a NetFlow device, you can configure multiple site URL specifications. For example:

```
=================================
TOP SITES
=================================
[192.168.10.25]
www.xyz.com=192.168.10.30-192.168.10.45
www.abc.com=192.168.10.125,192.168.10.121,192.168.10.130,192.168.10.90
```

5. If a single eG agent is monitoring multiple NetFlow devices, then in the eg_netflow.ini file of that eG agent, you can create multiple sub-sections - one each for every NetFlow device - and configure web sites to be monitored for each device. For example:

```
=================================
TOP SITES
=================================
[192.168.10.25]
www.xyz.com=192.168.10.30-192.168.10.45
www.abc.com=192.168.10.125,192.168.10.121,192.168.10.130,192.168.10.90
[192.168.10.200]
www.eazycart.com=192.168.10.1,192.168.10.2
www.fb.com=192.168.10.5,192.168.10.9
```

6. Finally, save the file.

Target of the test : A Netflow Device

Agent deploying the test : An external agent

Outputs of the test : One set of results for every site URL that is configured for monitoring in the eg_netflow.ini file

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when reporting traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**. Once this is done, then each interface will process only 1 packet out of a configured number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling Rate you specify.
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
Total data exchange	Indicates the total amount of data transmitted and received by this web site during the last measurement period.	KB	Compare the value of this measure across web sites to identify which web site is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) to or from this web site, and the amount of data transacted in bytes and packets in every flow. With the help of this detailed diagnosis, you can quickly compare the top netflows, know which netflow generated the maximum traffic, and figure out which source traffic originated from. Once the problem source is isolated, you can then investigate why traffic from that source is high.
Total packets exchanged	Indicates the total number of packets transmitted and received by this web site during the last measurement period.	Packets	Compare the value of this measure across sources to identify which web site is contributing to the high level of network traffic
Data exchange rate	Indicates the rate at which this web site transmitted/received data.	Kbps	
Packets exchange rate	Indicates the rate at which this web site	Kbps	

Measurement	Description	Measurement Unit	Interpretation
	transmitted/received packets.		
Total traffic to this web site	Indicates what percentage of the total traffic on this interface was to this web site.	Percent	A value close to 100% for this measure indicates that traffic to this web site is imposing the maximum load on the network. If users complain of a latent network, you can compare the value of this measure across web sites to accurately identify the web site that is responsible for any congestion on the network.
Total bandwidth utilization	Indicates the bandwidth utilized by this web site.	Percent	A value close to 100% is indicative of excessive bandwidth utilization by this web site when receiving data. If users complain of a latent network, you can compare the value of this measure across web sites to accurately identify the web site that is responsible for any congestion on the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this web site.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on this interface pertains to this web site.	Percent	
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface when receiving data for this web site from different sources and when receiving data from this web site.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.

Measurement	Description	Measurement Unit	Interpretation
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface when transmitting data to this web site and when transmitting data for this web site to a source.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.
Data received	Indicates the amount of data received by this web site.	KB	Compare the value of this measure across web sites to know which web site is receiving maximum data over this interface.
Data sent	Indicates the amount of data sent by this web site.	KB	Compare the value of this measure across web sites to know which web site is transmitting maximum data over this interface.
Packets received	Indicates the number of packets sent by this web site.	Number	
Packets sent	Indicates the number of packets received by this web site.	Number	
Data received rate	Indicates the rate at which data is received by this web site.	Kbps	If the value of this measure consistently drops for this web site, it could indicate an incoming traffic congestion.
Data transmitted rate	Indicates the rate at which data is sent by this web site.	Kbps	If the value of this measure consistently drops for this web site, it could indicate an outgoing traffic congestion.
Packets received rate	Indicates the rate at which packets are received by this web site.	Packets/Sec	If the value of this measure consistently drops for this web site, it could indicate an incoming traffic congestion.
Packets transmitted rate	Indicates the rate at which packets are transmitted by	Packets/Sec	If the value of this measure consistently drops for this web site, it

Measurement	Description	Measurement Unit	Interpretation
	this web site.		could indicate an outgoing traffic congestion.

3.1.6 Traffic to Destinations Test

Sometimes, administrators may mark certain destinations as critical and may want to closely watch the traffic flowing into those destinations alone and the bandwidth they use. To achieve this, administrators can use the Traffic to Destinations test. For each destination that is explicitly configured for monitoring, this test reports the traffic flowing into and out of every destination via each interface. Additionally, the test also reports the bandwidth used by this traffic. In the process, the test promptly alerts administrators to any abnormal increase in traffic volume to destinations and sudden/consistent rise in bandwidth usage by a destination.

Target of the test : A NetFlow device

Agent deploying the test : An external agent

Outputs of the test : One set of results for every destination configured for monitoring for each interface of the target device

First-level descriptor: Interface Name

Second-level descriptor: Destination IP address

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Report Top N Flows	This parameter is not applicable to the Traffic from Destinations Test.
Show DD for Top N Flows	By default, this test will report detailed diagnostics for only the top-5 destinations in terms of the volume of traffic they handle. If required , you can increase or decrease the value of this parameter to view detailed metrics for more or less (as the case may be) number of top destinations.
Include Destinations	Provide a comma-separated list of destinations that you want to monitor. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of destination IP addresses as a comma-separated list

Parameter	Description
	- for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of destination IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** Destinations can be represented using IP addresses only and not host names.
Exclude Interfaces	Provide a comma-separated list of interfaces that you want excluded from monitoring. For example, your specification can be, FastEtherNet0/0, FastEthernet0/1,FastEtherNet0/2. You can even provide patterns of interfaces as a comma-separated list - for instance, *ethernet*,Fast*.
Show Host Names	This test captures statistics on traffic that originates from source hosts, where each host is by default represented by its IP address in the eG monitoring console. Accordingly, this flag is set to **No** by default. You can set this flag to **Yes** so that the names of the individual hosts are displayed in the eG monitoring console instead of their IP addresses.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when identifying top sources and reporting their traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**. Once this is done, then each interface will process only 1 packet out of a configured number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling Rate you specify.
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

Parameter	Description
	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 SNMP. 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 the Username 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, this flag is set to **No** by default.

Parameter	Description
	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**.
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
Total data exchange	Indicates the total amount of data transmitted and received by this destination during the last measurement period.	KB	Compare the value of this measure across destinations to identify which destination host is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) that originated from this destination, and the amount of data transacted in bytes and packets in every flow. With the help of this detailed diagnosis, you can quickly compare the top netflows, know which netflow generated the maximum traffic, and figure out to which destination that traffic was headed. Once the problem destination is isolated, you can then investigate why traffic to that destination is high.
Total packets exchanged	Indicates the total number of packets transmitted and received by this destination during the last measurement period.	Packets	Compare the value of this measure across sources to identify which destination host is contributing to the high level of network traffic
Data exchange	Indicates the rate at which	Kbps	

Measurement	Description	Measurement Unit	Interpretation
rate	this destination transmitted/received data.		
Packets exchange rate	Indicates the rate at which this destination transmitted/received packets.	Kbps	
Total traffic to this destination	Indicates what percentage of the total traffic on this interface was to this destination.	Percent	A value close to 100% for this measure indicates that traffic to this destination is imposing the maximum load on the network. If users complain of a latent network, you can compare the value of this measure across destinations to accurately identify the destination that is responsible for any congestion on the network.
Total bandwidth utilization	Indicates the bandwidth utilized by this destination.	Percent	A value close to 100% is indicative of excessive bandwidth utilization by this destination when receiving data. If users complain of a latent network, you can compare the value of this measure across destinations to accurately identify the destination that is responsible for any congestion on the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this destination.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on this interface pertains to this destination.	Percent	

Measurement	Description	Measurement Unit	Interpretation
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface when receiving data for this destination from different sources and when receiving data from this destination.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface when transmitting data to this destination and when transmitting data for this destination to a source.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.
Data received	Indicates the amount of data received by this destination.	KB	Compare the value of this measure across destinations to know which destination is receiving maximum data over this interface.
Data sent	Indicates the amount of data sent by this destination.	KB	Compare the value of this measure across destinationsto know which destination is transmitting maximum data over this interface.
Packets received	Indicates the number of packets sent by this destination.	Number	
Packets sent	Indicates the number of packets received by this destination.	Number	
Data received rate	Indicates the rate at which data is received by this destination.	Kbps	If the value of this measure consistently drops for this destination, it could indicate an incoming traffic congestion.
Data transmitted rate	Indicates the rate at which data is sent by this destination.	Kbps	If the value of this measure consistently drops for this destination, it could indicate an outgoing traffic congestion.

Measurement	Description	Measurement Unit	Interpretation
Packets received rate	Indicates the rate at which packets are received by this destination.	Packets/Sec	If the value of this measure consistently drops for this destination, it could indicate an incoming traffic congestion.
Packets transmitted rate	Indicates the rate at which packets are transmitted by this destination.	Packets/Sec	If the value of this measure consistently drops for this destination, it could indicate an outgoing traffic congestion.

3.1.7 Traffic from Sources Test

Sometimes, administrators may mark certain sources as critical and may want to closely watch the traffic originating from those sources alone and the bandwidth they use. To achieve this, administrators can use the **Traffic from Sources** test. For each source that is explicitly configured for monitoring, this test reports the traffic flowing out and into every source via each interface. Additionally, the test also reports the bandwidth used by the data handled by every source. In the process, the test promptly alerts administrators to any abnormal increase in traffic volume on critical sources and sudden/consistent rise in bandwidth usage by a source.

Target of the test : A NetFlow device

Agent deploying the test : An external agent

Outputs of the test : One set of results for every source configured for monitoring for each interface of the target device

First-level descriptor: Interface Name

Second-level descriptor: Source IP address

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed
Host	The host for which the test is to be configured.
Report Top N Flows	This parameter is not applicable to the Traffic from Sources Test.
Show DD for Top N Flows	By default, this test will report detailed diagnostics for only the top-5 sources in terms of the volume of traffic they handle. If required , you can increase or decrease the value

Parameter	Description
	of this parameter to view detailed metrics for more or less (as the case may be) number of top sources.
Include Sources	Provide a comma-separated list of sources you want to monitor. For example, your specification can be, 192.168.10.45, 192.168.10.71,192.168.10.220. You can even provide patterns of source IP addresses as a comma-separated list - for instance, 192.168.10.*,192.168.8.1*1, *.168.9.45. Alternatively, you can provide a range of source IP addresses. For example, 192.168.10.32-192.168.10.45. **Note:** Sources can be represented using IP addresses only and not host names.
Exclude Interfaces	Provide a comma-separated list of interfaces that you want excluded from monitoring. For example, your specification can be, FastEtherNet0/0, FastEthernet0/1,FastEtherNet0/2. You can even provide patterns of interfaces as a comma-separated list - for instance, *ethernet*,Fast*.
Show Host Names	This test captures statistics on traffic that originates from source hosts, where each host is by default represented by its IP address in the eG monitoring console. Accordingly, this flag is set to **No** by default. You can set this flag to **Yes** so that the names of the individual hosts are displayed in the eG monitoring console instead of their IP addresses.
Process Tainted Packets	Network latencies and processing bottlenecks can sometimes cause netflow records to be transmitted slowly to the NetFlow collector. In such a situation, you can instruct the collector to either process or ignore the delayed NetFlow records. If you want the metrics reported to pertain to current NetFlow records only, then you may choose to ignore the delayed records. In this case, set this flag to **No**. If you want old NetFlow records to also be considered when identifying top sources and reporting their traffic statistics, then set this flag to **Yes**. Typically, any NetFlow record that takes 10 minutes or over to reach the NetFlow collector is deemed a tainted/delayed record.
Sampling, Sampling Rate	By default, NetFlow is designed to process all IP packets on an interface. This is why, the Sampling flag is set to **No** by default. In some environments however, e.g. on Internet backbones, processing all IP packets can be too costly, due to the extra processing required for each packet and large number of simultaneous flows. This is where sampling is useful. In such environments, set the Sampling flag to **Yes**. Once this is done, then each interface will process only 1 packet out of a configured number of packets. Specify the number of packets from which this 1 packet should be picked in the Sampling Rate text box. For instance, to pick 1 out of 1000 packets for sampling, set the Sampling Rate to 1000. Where Sampling is enabled, all NetFlow metrics - particularly, metrics on traffic volume - will be adjusted based on the Sampling

Parameter	Description
	Rate you specify.
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 SNMP. 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 the Username 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

Parameter	Description
	• **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, 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**.
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

Parameter	Description
	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 exchange	Indicates the total amount of data transmitted and received by this source during the last measurement period.	KB	Compare the value of this measure across sources to identify which source host is contributing to the high level of network traffic. Use the detailed diagnosis of this measure to determine the top netflows (in terms of the volume of data transacted) that originated from this source, and the amount of data transacted in bytes and packets in every flow. With the help of this detailed diagnosis, you can quickly compare the top netflows, know which netflow generated the maximum traffic, and figure out which destination that traffic was headed to. Once the problem destination is isolated, you can then investigate why traffic to that destination was high - is it because of the type of application executing on that destination? (eg., an online game or a movie that would typically consume a lot of bandwidth), or is it because of a poor network line connecting the source and the destination?
Total packets exchange	Indicates the total number of packets transmitted and	Packets	Compare the value of this measure across sources to identify which

Measurement	Description	Measurement Unit	Interpretation
	received by this source during the last measurement period.		source host is contributing to the high level of network traffic
Data exchange rate	Indicates the rate at which this source transmitted/received data.	Kbps	
Packets exchange rate	Indicates the rate at which this source transmitted/received packets.	Kbps	
Total traffic by this source	Indicates what percentage of the total traffic on this interface was from this source.	Percent	A value close to 100% for this measure indicates that traffic from this source is imposing the maximum load on the network. If users complain of a latent network, you can compare the value of this measure across sources to accurately identify the source that is responsible for any congestion on the network.
Total bandwidth utilization	Indicates the bandwidth utilized by this source.	Percent	A value close to 100% indicates that this source is indicative of excessive bandwidth utilization by this source when transmitting data. If users complain of a latent network, you can compare the value of this measure across sources to accurately identify the source that is responsible for any congestion on the network.
In traffic	Indicates what percentage of total incoming traffic on this interface pertains to this source.	Percent	
Out traffic	Indicates what percentage of total outgoing traffic on this interface pertains to this	Percent	

Measurement	Description	Measurement Unit	Interpretation
	source.		
Ingress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic coming into this interface when receiving data from this source and receiving data for this source from destinations.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in incoming traffic on this interface.
Egress bandwidth utilization	Indicates the percentage of bandwidth utilized by traffic going out of this interface when transmitting data to this source and when transmitting data for this source to destinations.	Percent	A value close to 100% is a cause for concern as it implies a potential congestion in outgoing traffic on this interface.
Data received	Indicates the amount of data received by this source.	KB	Compare the value of this measure across sources to know which source is receiving maximum data over this interface.
Data sent	Indicates the amount of data sent by this source.	KB	Compare the value of this measure across sources to know which source is transmitting maximum data over this interface.
Packets received	Indicates the number of packets sent by this source.	Number	
Packets sent	Indicates the number of packets received by this source.	Number	
Data received rate	Indicates the rate at which data is received by this source.	Kbps	If the value of this measure consistently drops for this source, it could indicate an incoming traffic congestion.
Data transmitted rate	Indicates the rate at which data is sent by this source.	Kbps	If the value of this measure consistently drops for this source, it could indicate an outgoing traffic

Measurement	Description	Measurement Unit	Interpretation
			congestion.
Packets received rate	Indicates the rate at which packets are received by this source.	Packets/Sec	If the value of this measure consistently drops for this source, it could indicate an incoming traffic congestion.
Packets transmitted rate	Indicates the rate at which packets are transmitted by this source.	Packets/Sec	If the value of this measure consistently drops for this source, it could indicate an outgoing traffic congestion.

3.2 NetFlow Dashboard

The NetFlow Dashboard meaningfully groups netflow statistics pertaining to a target NetFlow device and aesthetically presents these statistics using a variety of visual tools - eg., graphs, charts, grids, etc. - with the aim of enabling administrators rapidly spot network bottlenecks and easily and accurately isolate the root-cause of the bottlenecks.

Overview

The Overview tab page of the dashboard, as the name suggests, presents an overview of network health. The left panel of the Overview tab page displays dial charts representing the current health of the network link to the target NetFlow device. A quick look at these charts will tell you whether the device is currently available over the network, and if so, whether it is experiencing any network latency. In the right panel, you will find graphs that depict the level of traffic on and the bandwidth usage of each interface supported by the monitored NetFlow device, during the last hour (by default). With the help of these graphs, you can quickly spot those interfaces that have been consistently handling an abnormally high volume of traffic and those that have been consuming bandwidth excessively.

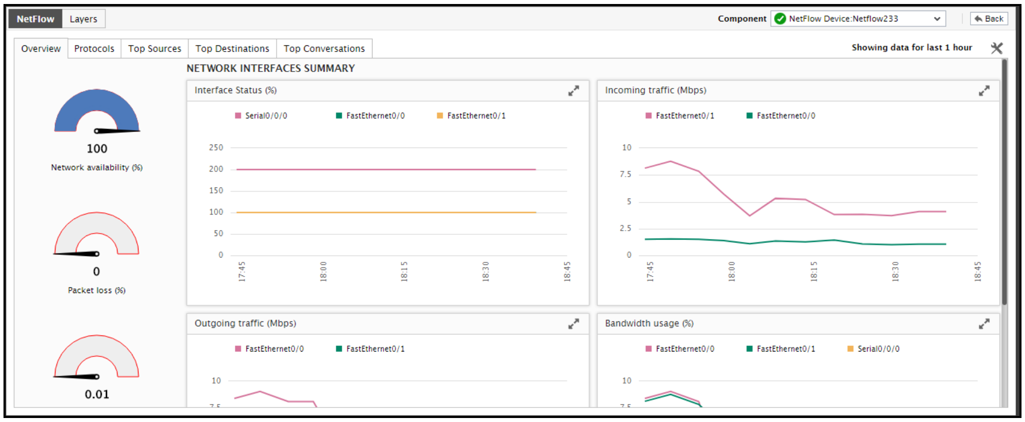

Figure 3.4: The Overview tab page

To increase or decrease the timeline of the graphs, click on the ⚒ icon at the right, top corner of Figure 3.4. This will open Figure 3.5, using which you can change the graph timeline. This way, you can alter the default timeline of all the graphs displayed in the Overview tab page.

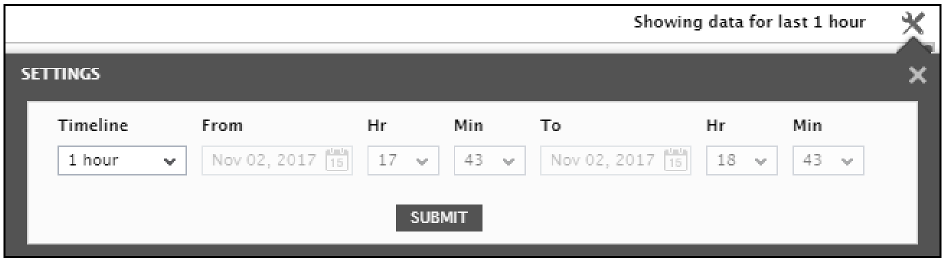
Showing data for last 1 hour

SETTINGS

Timeline: 1 hour
From: Nov 02, 2017 Hr: 17 Min: 43
To: Nov 02, 2017 Hr: 18 Min: 43

Figure 3.5: Changing the timeline of the graphs in the Overview tab page

On the other hand, if you want to change the timeline of a particular graph alone, then you first have to enlarge the graph. For that, click on it.Figure 3.6 will then appear. Using Figure 3.6, you can view the graph clearly and can also easily change the timeline of that graph alone.

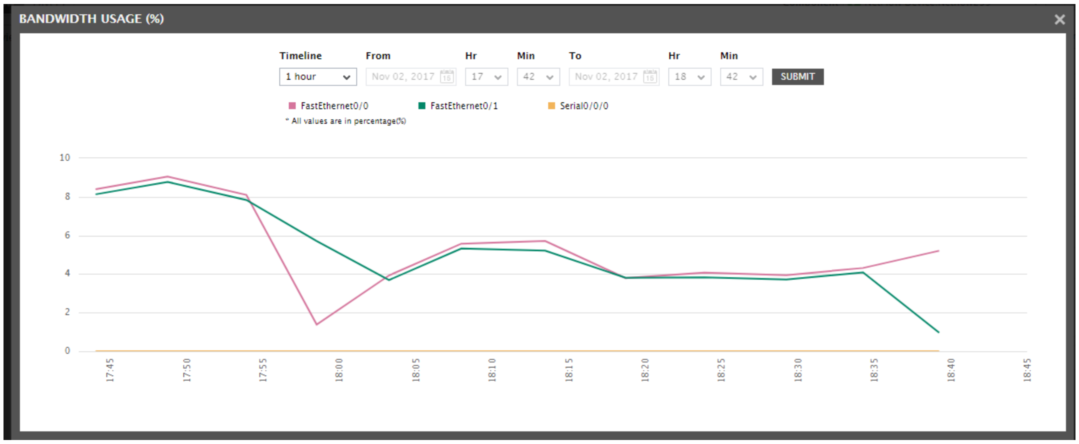
BANDWIDTH USAGE (%)

Timeline: 1 hour
From: Nov 02, 2017 Hr: 17 Min: 42
To: Nov 02, 2017 Hr: 18 Min: 42

Figure 3.6: Viewing the graph in the enlarged mode and changing its timeline

Protocols

Using the distribution charts in the Protocols tab page, you can easily analyze how traffic on each interface was distributed across applications/protocols, during the last hour (by default). If a particular interface is seeing an abnormally high volume of traffic, then a single glance at the distribution chart corresponding to that interface will lead you to the exact application/protocol that is contributing to that abnormality.

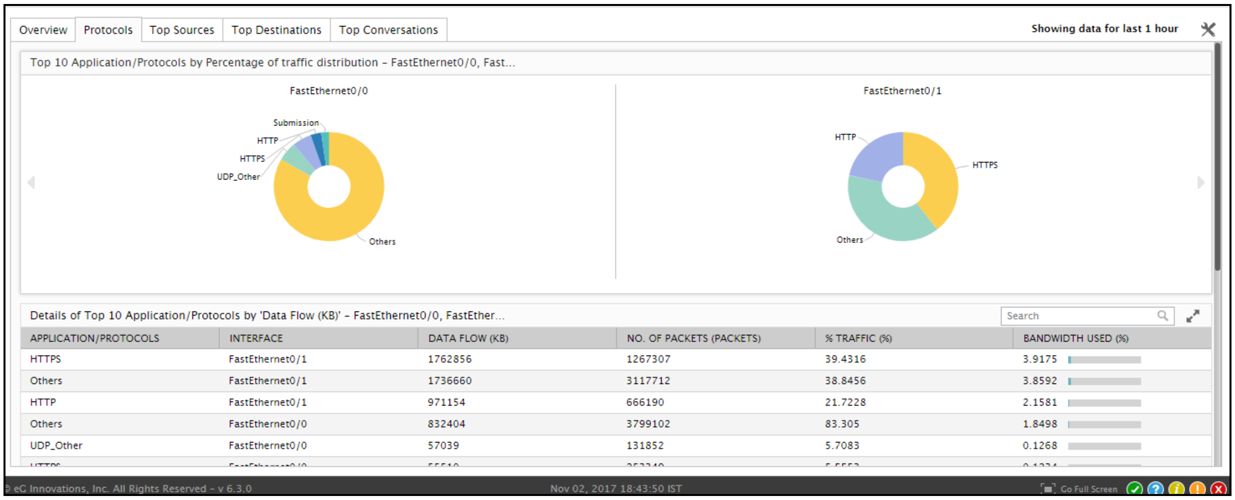
Overview
Protocols
Top Sources
Top Destinations
Top Conversations
Showing data for last 1 hour
Top 10 Application/Protocols by Percentage of traffic distribution – FastEthernet0/0, Fast...
FastEthernet0/0
Submission
HTTP
HTTPS
UDP_Other
Others
FastEthernet0/1
HTTP
HTTPS
Others
Search
© eG Innovations, Inc. All Rights Reserved – v 6.3.0
Nov 02, 2017 18:43:50 IST
Go Full Screen

Details of Top 10 Application/Protocols by 'Data Flow (KB)' – FastEthernet0/0, FastEther...

APPLICATION/PROTOCOLS	INTERFACE	DATA FLOW (KB)	NO. OF PACKETS (PACKETS)	% TRAFFIC (%)	BANDWIDTH USED (%)
HTTPS	FastEthernet0/1	1762856	1267307	39.4316	3.9175
Others	FastEthernet0/1	1736660	3117712	38.8456	3.8592
HTTP	FastEthernet0/1	971154	666190	21.7228	2.1581
Others	FastEthernet0/0	832404	3799102	83.305	1.8498
UDP_Other	FastEthernet0/0	57039	131852	5.7083	0.1268

Figure 3.7: Viewing traffic distribution charts and top-10 applications/protocols grid in the Protocols tab page

To analyze data flow across interfaces, use the Top 10 Applications/Protocols by Data Flow grid of Figure 3.7. This grid is by default sorted in the descending order of % Traffic, and hence, pinpoints the application/protocol that is contributing to high traffic volume on the network, regardless of interface. You can also use the Search option here to search for a particular record(s) in the grid. All you need to is type the string you want to search for in the Search text box and click the magnifying glass icon alongside. All records with a value that matches the search string in whole/part will be listed.

Scrolling down the dashboard will reveal the a time-of-day graph that compares bandwidth usage across applications/protocols during the last hour (by default). Using this graph, you can quickly and accurately isolate bandwidth-intensive applications/protocols.

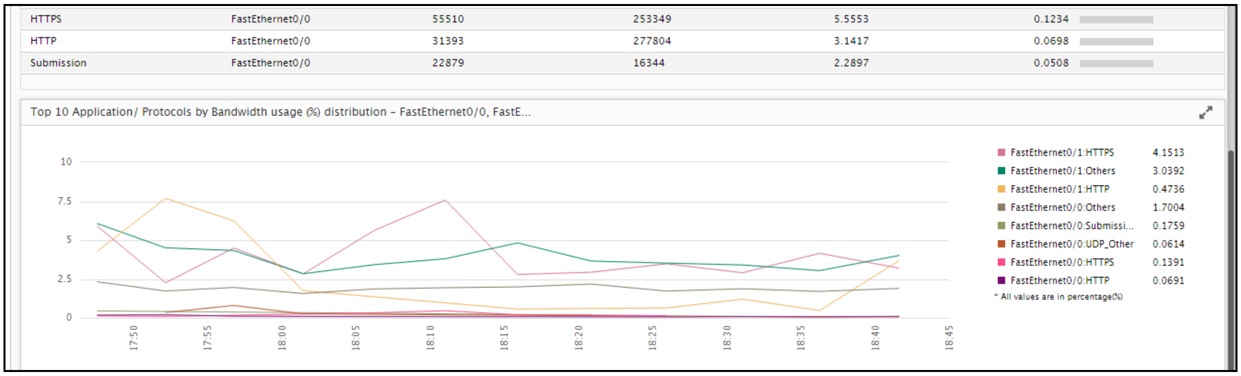

HTTPS	FastEthernet0/0	55510	253349	5.5553	0.1234
HTTP	FastEthernet0/0	31393	277804	3.1417	0.0698
Submission	FastEthernet0/0	22879	16344	2.2897	0.0508

Top 10 Application/ Protocols by Bandwidth usage (%) distribution - FastEthernet0/0, FastE...

🟥	FastEthernet0/1:HTTPS	4.1513
🟩	FastEthernet0/1:Others	3.0392
🟧	FastEthernet0/1:HTTP	0.4736
🟫	FastEthernet0/0:Others	1.7004
🟩	FastEthernet0/0:Submissi...	0.1759
🟥	FastEthernet0/0:UDP_Other	0.0614
🟥	FastEthernet0/0:HTTPS	0.1391
🟪	FastEthernet0/0:HTTP	0.0691

* All values are in percentage(%)

Figure 3.8: A graph comparing bandwidth usage across applications/protocols

You can make changes to the default settings governing this tab page by clicking on the ✕ icon at the right, top corner of Figure 3.7. Figure 3.9 will then appear.

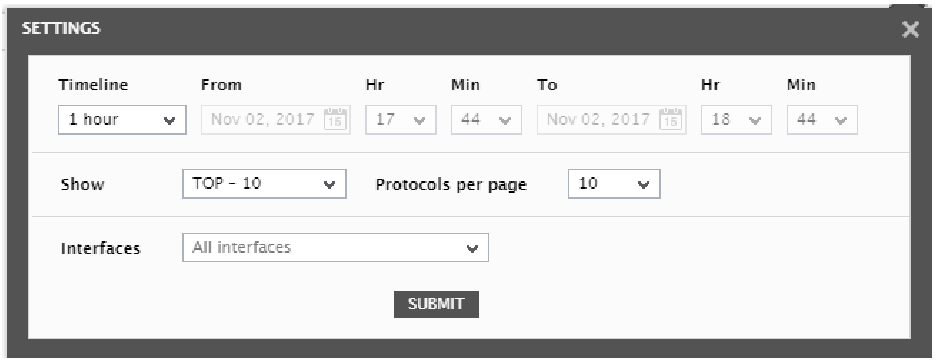

SETTINGS

Timeline: 1 hour
From: Nov 02, 2017 Hr: 17 Min: 44
To: Nov 02, 2017 Hr: 18 Min: 44

Show: TOP - 10
Protocols per page: 10

Interfaces: All interfaces

Figure 3.9: Overriding default settings of the Protocols tab page

Using Figure 3.9, you can change the default time line of 1 hour for the distribution charts and graphs in the Protocols tab page. Also, by default, the distribution charts, graphs, and the grid display the top-10 applications/protocols alone in terms of traffic or bandwidth usage (as the case may be). You can increase or decrease this number. Alternatively, you can even choose to show those applications/protocols that consume very less bandwidth or generate very low traffic on the interfaces. For this purpose, pick a different Top-N option or a Last-N option from the Show drop-down in Figure 3.9.

You can indicate how many protocols you want displayed per page by picking an option from the Protocols per page drop-down.

Moreover, by default, the charts, graphs, and grid in the Protocols tab page represent statistics pertaining to all interfaces. You can instead have them represent performance of a particular interface alone by picking an interface of your choice from the Interfaces drop-down. This will enable you to focus on applications/protocols pertaining to a specific interface.

Once you are done with all changes, click the Submit button in Figure 3.9 to register the changes.

Top Sources

To identify the top source in terms of data flow at a single glance, use the Top Sources tab page. The traffic distribution charts at the top of this tab page reveal which source imposed the maximum load on each interface during the last 30 minutes (by default). If a particular interface appears to be seeing an abnormally high level of traffic, then, you can use the distribution chart of that interface to accurately isolate the source that is contributing to the interface overload.

If you want to know which source, across interfaces, generated a very high volume of traffic during the same default timeline (30 minutes), use the grid below the distribution charts. This grid is by default sorted in the descending order of % Traffic, and hence, pinpoints the source that is contributing to high traffic volume on the network, regardless of interface.

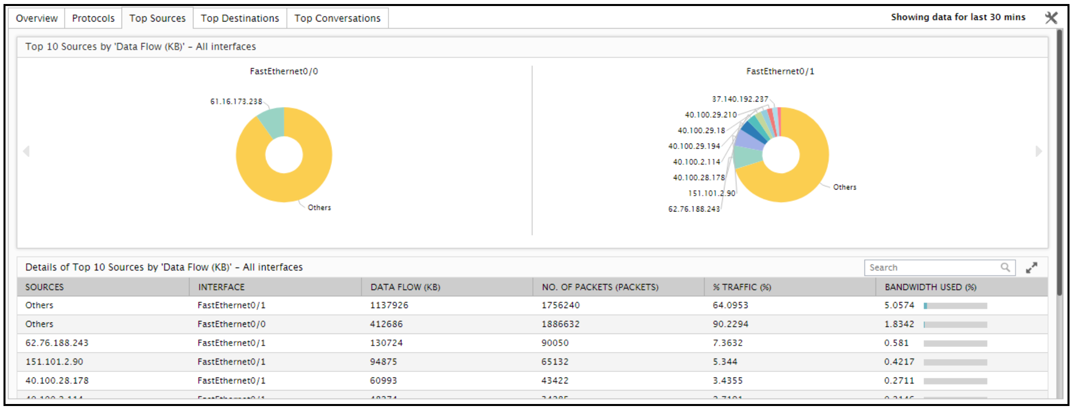

Top 10 Sources by 'Data Flow (KB)' – All interfaces

Details of Top 10 Sources by 'Data Flow (KB)' – All interfaces

SOURCES	INTERFACE	DATA FLOW (KB)	NO. OF PACKETS (PACKETS)	% TRAFFIC (%)	BANDWIDTH USED (%)
Others	FastEthernet0/1	1137926	1756240	64.0953	5.0574
Others	FastEthernet0/0	412686	1886632	90.2294	1.8342
62.76.188.243	FastEthernet0/1	130724	90050	7.3632	0.581
151.101.2.90	FastEthernet0/1	94875	65132	5.344	0.4217
40.100.28.178	FastEthernet0/1	60993	43422	3.4355	0.2711

Figure 3.10: The Top Sources tab page

You can also use the Search option here to search for a particular record(s) in the grid. All you need to do is type the string you want to search for in the Search text box and click the magnifying glass icon alongside. All records with a value that matches the search string in whole/part will be listed.

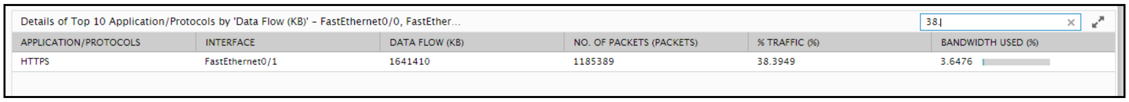

Details of Top 10 Application/Protocols by 'Data Flow (KB)' – FastEthernet0/0, FastEther...

38

APPLICATION/PROTOCOLS	INTERFACE	DATA FLOW (KB)	NO. OF PACKETS (PACKETS)	% TRAFFIC (%)	BANDWIDTH USED (%)
HTTPS	FastEthernet0/1	1641410	1185389	38.3949	3.6476

Figure 3.11: Using the Search option in the Top Sources tab page

You can make changes to the default settings governing this tab page by clicking on the ✕ icon at the right, top corner of Figure 3.11. Figure 3.12 will then appear.

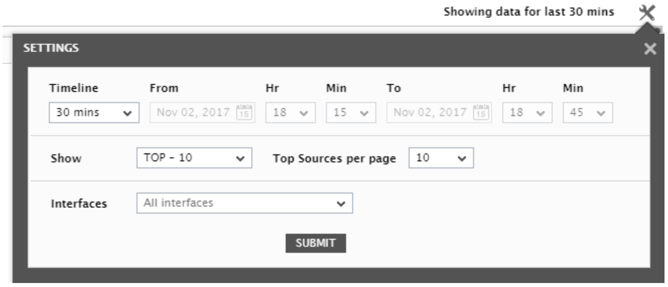

Showing data for last 30 mins

SETTINGS

Timeline	From	Hr	Min	To	Hr	Min
30 mins	Nov 02, 2017	18	15	Nov 02, 2017	18	45

Show: TOP - 10

Top Sources per page: 10

Interfaces: All interfaces

Figure 3.12: Overriding default settings of the Top Sources tab page

Using Figure 3.12, you can change the default time line of 30 minutes for the distribution charts and grid in the Top Sources tab page. Also, by default, the distribution charts and the grid display the top-10 sources alone in terms of traffic. You can increase or decrease this number. Alternatively, you can even choose to show those sources that consume very less bandwidth or generate very low traffic on the interfaces. For this purpose, pick a different Top-N option or a Last-N option from the Show drop-down in Figure 3.12.

You can indicate how many sources you want displayed per page by picking an option from the Top Sources per page drop-down.

Moreover, by default, the charts and grid in the Top Sources tab page represent statistics pertaining to all interfaces. You can instead have them represent performance of a particular interface alone by picking an interface of your choice from the Interfaces drop-down. This will enable you to focus on top sources pertaining to a specific interface.

Once you are done with all changes, click the Submit button in Figure 3.12 to register the changes.

Top Destinations

To identify the top destination in terms of data flow at a single glance, use the Top Destinations tab page. The traffic distribution charts at the top of this tab page reveal which destination imposed the maximum load on each interface during the last 30 minutes (by default). If a particular interface appears to be seeing an abnormally high level of traffic, then, you can use the distribution chart of that interface to accurately isolate the destination that is contributing to the interface overload.

If you want to know which destination, across interfaces, generated a very high volume of traffic during the same default timeline (30 minutes), use the grid below the distribution charts. This grid is by default sorted in the descending order of % Traffic, and hence, pinpoints the destination that is contributing to high traffic volume on the network, regardless of interface.

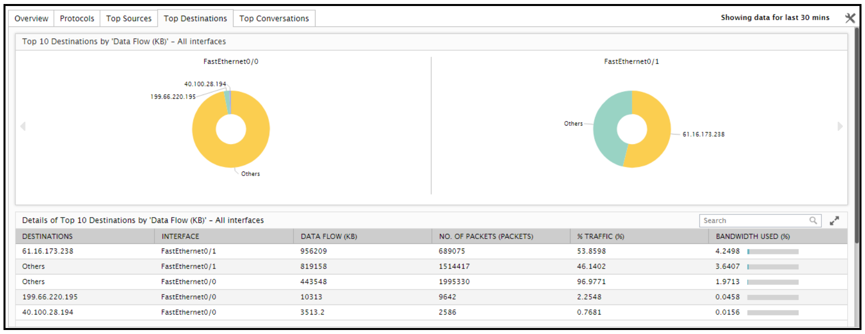

Top 10 Destinations by 'Data Flow (KB)' – All interfaces

Details of Top 10 Destinations by 'Data Flow (KB)' – All interfaces

DESTINATIONS	INTERFACE	DATA FLOW (KB)	NO. OF PACKETS (PACKETS)	% TRAFFIC (%)	BANDWIDTH USED (%)
61.16.173.238	FastEthernet0/1	956209	689075	53.8598	4.2498
Others	FastEthernet0/1	819158	1514417	46.1402	3.6407
Others	FastEthernet0/0	443548	1995330	96.9771	1.9713
199.66.220.195	FastEthernet0/0	10313	9642	2.2548	0.0458
40.100.28.194	FastEthernet0/0	3513.2	2586	0.7681	0.0156

Figure 3.13: The Top Destinations tab page

You can also use the Search option here to search for a particular record(s) in the grid. All you need to do is type the string you want to search for in the Search text box and click the magnifying glass icon alongside. All records with a value that matches the search string in whole/part will be listed.

You can make changes to the default settings governing this tab page by clicking on the ✕ icon at the right, top corner of Figure 3.13. Figure 3.14 will then appear.

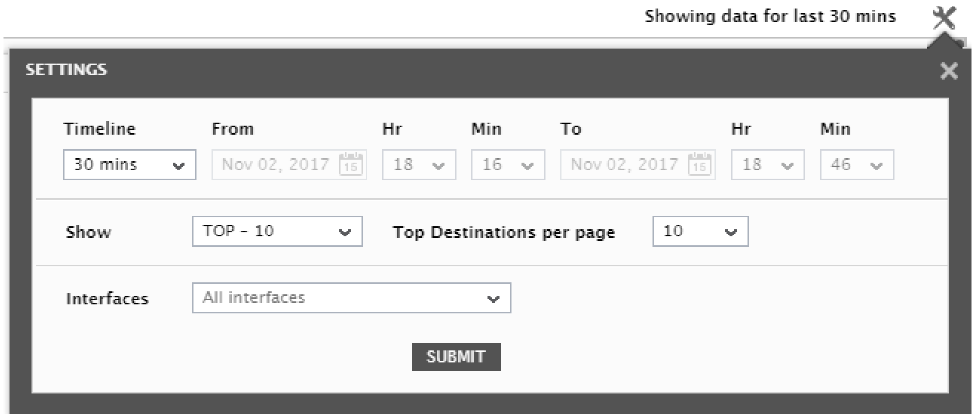

Timeline: 30 mins
From: Nov 02, 2017 Hr: 18 Min: 16
To: Nov 02, 2017 Hr: 18 Min: 46
Show: TOP - 10
Top Destinations per page: 10
Interfaces: All interfaces

Figure 3.14: Overriding default settings of the Top Destinations tab page

Using Figure 3.14, you can change the default time line of 30 minutes for the distribution charts and grid in the Top Destinations tab page. Also, by default, the distribution charts and the grid display the top-10 destinations alone in terms of traffic. You can increase or decrease this number. Alternatively, you can even choose to show those destinations that consume very less bandwidth or generate very low traffic on the interfaces. For this purpose, pick a different Top-N option or a Last-N option from the Show drop-down in Figure 3.14.

You can indicate how many sources you want displayed per page by picking an option from the Top Destinations per page drop-down.

Moreover, by default, the charts and grid in the Top Destinations tab page represent statistics pertaining to all interfaces. You can instead have them represent performance of a particular interface alone by picking an interface of your choice from the Interfaces drop-down. This will enable you to focus on top destinations pertaining to a specific interface.

Once you are done with all changes, click the Submit button in Figure 3.14 to register the changes.

Top Conversations

To identify the top conversations in terms of data flow at a single glance, use the Top Conversations tab page. The traffic distribution charts at the top of this tab page reveal which conversation imposed the maximum load on each interface during the last 30 minutes (by default). If a particular conversation appears to be seeing an abnormally high level of traffic, then, you can use the distribution chart of that interface to accurately isolate the conversation that is contributing to the interface overload.

If you want to know which conversation, across interfaces, generated a very high volume of traffic during the same default timeline (30 minutes), use the grid below the distribution charts. This grid is by default sorted in the descending order of % Traffic, and hence, pinpoints the conversation that is contributing to high traffic volume on the network, regardless of interface.

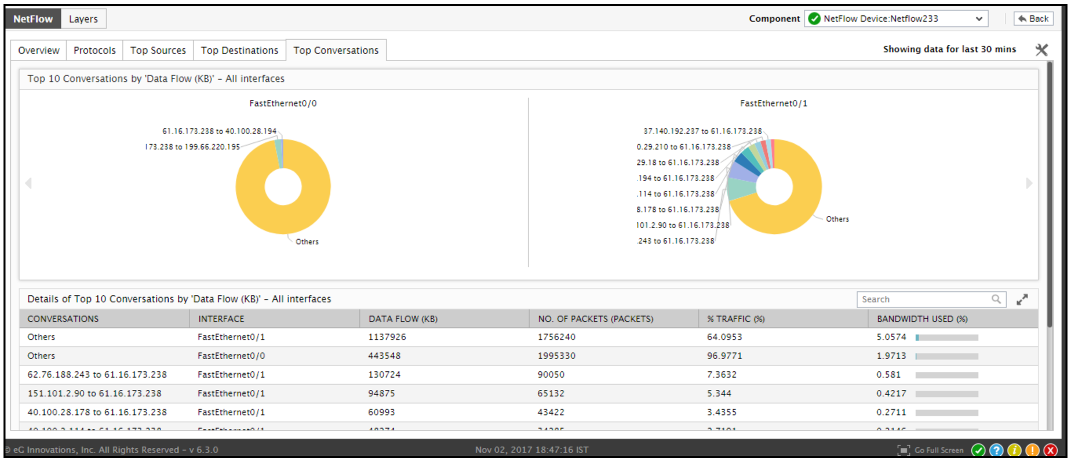
NetFlow
Layers
Component
NetFlow Device:Netflow233
Back
Overview
Protocols
Top Sources
Top Destinations
Top Conversations
Showing data for last 30 mins
FastEthernet0/0
61.16.173.238 to 40.100.28.194
173.238 to 199.66.220.195
Others
FastEthernet0/1
37.140.192.237 to 61.16.173.238
0.29.210 to 61.16.173.238
29.18 to 61.16.173.238
.194 to 61.16.173.238
.114 to 61.16.173.238
8.178 to 61.16.173.238
101.2.90 to 61.16.173.238
.243 to 61.16.173.238
Others
Search
© eG Innovations, Inc. All Rights Reserved - v 6.3.0
Nov 02, 2017 18:47:16 IST
Go Full Screen

Top 10 Conversations by 'Data Flow (KB)' – All interfaces

Details of Top 10 Conversations by 'Data Flow (KB)' – All interfaces

CONVERSATIONS	INTERFACE	DATA FLOW (KB)	NO. OF PACKETS (PACKETS)	% TRAFFIC (%)	BANDWIDTH USED (%)
Others	FastEthernet0/1	1137926	1756240	64.0953	5.0574
Others	FastEthernet0/0	443548	1995330	96.9771	1.9713
62.76.188.243 to 61.16.173.238	FastEthernet0/1	130724	90050	7.3632	0.581
151.101.2.90 to 61.16.173.238	FastEthernet0/1	94875	65132	5.344	0.4217
40.100.28.178 to 61.16.173.238	FastEthernet0/1	60993	43422	3.4355	0.2711

Figure 3.15: The Top Conversations tab page

You can also use the Search option here to search for a particular record(s) in the grid. All you need to do is type the string you want to search for in the Search text box and click the magnifying glass icon alongside. All records with a value that matches the search string in whole/part will be listed.

You can make changes to the default settings governing this tab page by clicking on the ✕ icon at the right, top corner of Figure 3.15. Figure 3.16 will then appear.

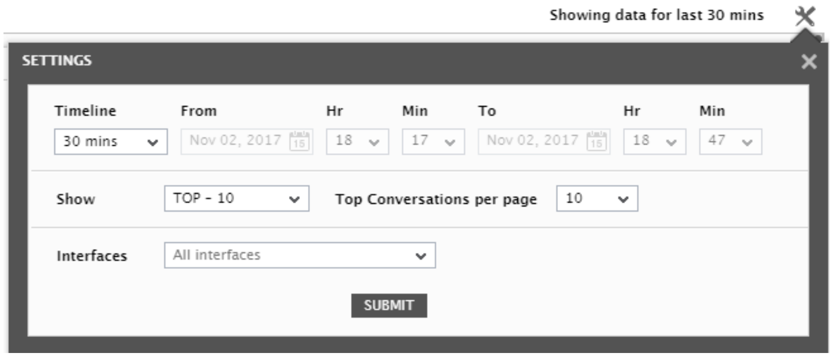

Showing data for last 30 mins

SETTINGS

Timeline	From	Hr	Min	To	Hr	Min
30 mins	Nov 02, 2017	18	17	Nov 02, 2017	18	47

Show: TOP - 10 Top Conversations per page: 10

Interfaces: All interfaces

Figure 3.16: Overriding default settings of the Top Conversations tab page

Using Figure 3.16, you can change the default time line of 30 minutes for the distribution charts and grid in the Top Conversations tab page. Also, by default, the distribution charts and the grid display the top- 10 conversations alone in terms of traffic. You can increase or decrease this number. Alternatively, you can even choose to show those conversations that consume very less bandwidth or generate very low traffic on the interfaces. For this purpose, pick a different Top-N option or a Last-N option from the Show drop-down in Figure 3.16.

You can indicate how many sources you want displayed per page by picking an option from the Top Conversations per page drop-down.

Moreover, by default, the charts and grid in the Top Conversations tab page represent statistics pertaining to all interfaces. You can instead have them represent performance of a particular interface alone by picking an interface of your choice from the Interfaces drop-down. This will enable you to focus on top conversations pertaining to a specific interface.

Once you are done with all changes, click the Submit button in Figure 3.16 to register the changes.

Chapter 4: Frequently Asked Questions

- **How is the NetFlow Device component licensed?**

  NetFlow monitoring by eG Enterprise is licensed by the number of eG External Agents used for collecting flow data from NetFlow-enabled devices. Licensing is NOT restricted by the number of devices/interfaces exporting flow data.

  Each eG External Agent includes one NetFlow Collector, which supports collection of up to 20,000 flows/second.

  Each external agent consumes a Premium Monitor license

- **Are the eG NetFlow tests available for the NetFlow Device component alone?**

  The eG NetFlow tests are by default enabled for the NetFlow Device component only. These tests are also mapped to the Fortigate Firewall and Network Node components, but are disabled by default. Using the eG admin interface, you can enable these tests for such components, if you so need.

- **On what port does the eG NetFlow Collector listen?**

  By default, the eG NetFlow Collector listens on UDP Port 9996.

- **Can I override the default port?**

  Yes, you can. For that, follow the steps below:

  - Login to the system on which the collector service is running.
  - Edit the Netflow.properties file in the <EG_INSTALL_DIR>\Netflow\config directory.
  - By default, the net.bind.port parameter in the file is set to 9996.
  - To change the binding port, provide a different port number against net.bind.port.
  - Then , save the file.

- **Are collector errors logged?**

  Yes; errors in the operations of the eG NetFlow Collector are logged in the collector.log file in the <EG_INSTALL_DIR>\NetFlow\logs directory.

- **Can a single eG NetFlow Collector receive NetFlow records from multiple NetFlow devices?**

  Yes; a single eG NetFlow Collector is capable of receiving and processing NetFlow records sent by multiple NetFlow devices.

- **What does the collector do if it receives NetFlow records from NetFlow devices that are not managed by eG Enterprise?**

  The collector ignores those NetFlow records that are received from NetFlow devices that are not managed by eG Enterprise. Also, in this case, the following message will be logged in the collector.log file in the <EG_INSTALL_DIR>\NetFlow\logs directory:

  ```
  30-Oct-2017 12:19:41 WARN Collector:498 - The /61.16.173.233 component is not
  managed yet. packets from this device is ignored
  ```

- **Where does the eG NetFlow Collector store the processed NetFlow data?**

  The eG NetFlow Collector processes the NetFlow records it receives from target NetFlow device and stores the processed data in binary files in its <EG_INSTALL_DIR>\NetFlow\data directory. In this directory, the collector creates a sub-folder each for every NetFlow device that it receives NetFlow data from. Upon receipt of NetFlow data from a target device, the collector processes that data, writes the data first to a .tmp file, and then moves it to a .dat file.

- **How frequently does the collector create binary files?**

  The collector writes NetFlow data to a binary file every 30 seconds by default. Data will first be written to a *.tmp and then moved to a *.dat file. If there is no data in the *.tmp file, then the *.dat file creation will be delayed by 30 seconds. If the .tmp file is empty for more than 30 minutes, it will be deleted and a new .tmp file will be created with current timestamp.

- **What happens to the binary files after the eG agent reads from them?**

  The eG agent deletes the files after reading them. If the eG agent is stopped or if files are not deleted by the agent, then the collector automatically deletes the files.

- **Which are the files the collector deletes?**

  Typically, the collector deletes all files that are of an age that is equal to or over 3 times the maximum test frequency of the eG NetFlow tests. For instance, assume that the frequency of the Top Sources, Top Destinations, Top Applications/Protocols, and Top Conversations test is set to 15 minutes, 10 minutes, 5 minutes, and 5 minutes respectively. The maximum frequency therefore is 15 minutes. In this case, the collector will delete all files that are over 45 minutes (3 * 15) old.

- **How frequently does the collector check for old files?**

  Every 30 seconds the collector will check the data folder for old files.

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.