



Extreme Application Sensor Engine Guide

10/2022

PN: 9037636-00

Subject to Change Without Notice

Table of Contents

Extreme Application SensorEngine Guide	1
Table of Contents	2
About This Guide	6
Configuration	7
Extreme Application Sensor Engine Virtual Engine Requirements	7
Adding the Extreme Application Sensor Engine Engine	7
Changing Extreme Application Sensor Engine Engine Settings	8
Changing Basic Network Configuration	8
Changing Date and Time Settings	8
Changing the ExtremeCloud IQ - Site Engine Server IP Address	9
Changing SNMP Configuration	9
Configuring the Application Sensor & Analytics Engine	9
Deploying the Extreme Application Sensor Engine Engine	13
VMWare TAP Interface	13
Deploying Extreme Application Sensor Engine Engine in an MSP or MSSP Environment	14
Configuring ExtremeCloud IQ - Site Engine Behind a NAT Router	14
Launching ExtremeAnalytics	15
Upgrading Extreme Application Sensor Engine Engine Software	15
Getting Started with Extreme Application Sensor Engine	16
ExtremeAnalytics Access Requirements	16
ExtremeAnalytics Engine Configuration	17
ExtremeAnalytics Tab Overview	17
Dashboard	17
Browser	17
Application Flows	18
Fingerprints	18
Packet Captures	18
Configuration	18

Reports	18
ExtremeAnalytics Dashboard Overview	19
Insights Dashboard Reports	19
Client/Server Dashboard Reports	19
Applications Browser Dashboard Report	20
Industry Dashboards	20
Enterprise Dashboard	20
Education Dashboard	20
Healthcare Dashboard	20
Venue Dashboard	20
Response Time Dashboard	20
Network Service Dashboard	21
Tracked Applications Dashboard	21
ExtremeAnalytics Insights Dashboard	21
Insights	22
Ring Chart	22
Custom Dashboard	23
ExtremeAnalytics Response Time Dashboard	23
Overview	24
Application	24
Top	25
Tracked Applications	25
Filters	25
Network Response Time Graph	25
Application Response Time Graph	26
ExtremeAnalytics Network Service Dashboard	27
Overview	27
Expected Response Time	28
Historical Response Time	29
ExtremeAnalytics Tracked Applications Dashboard	29

Overview	30
Expected Response Time	31
Historical Response Time	32
ExtremeAnalytics Browser Overview	32
Overview	32
Data Aggregation	33
Options	34
Data Table	34
Display Format	34
Target	34
Time Period	35
Statistic	35
Search Criteria	36
Bookmark	37
Save to Report Designer	37
Export to CSV	38
ExtremeAnalytics Application Flows	38
Overview	38
Application Flows Tables	40
Bidirectional Flows	41
Unidirectional Flows	41
Report Features	41
ExtremeAnalytics Bidirectional Flow Table	42
ExtremeAnalytics Unidirectional Flow Table	45
ExtremeAnalytics Fingerprints Overview	47
Analytics Application Data Collection	48
Data Collection Overview	48
Collection Targets	49
Collection Statistics	50
Collection Intervals	50

Using Sites to Collect In-Network Traffic	50
Data Collector Types	51
General Usage Collectors	51
Hourly General Usage Collectors	51
High-Rate General Usage Collectors	53
End-System Details Collector	54
Flow Information Sources	54
Enabling ExtremeControl Integration	55
Reports	56
Dashboard Report	56
Browser Reports	56
Glossary	58
Index	59

About This Guide

This document describes the installation and initial configuration of the Extreme Application Sensor Engine engine.

This document is intended for experienced network administrators who are responsible for implementing and maintaining communications networks.

Configuration

After you have installed the engine, power the engine on, and after the engine boots, you must go through the initial configuration process described in this chapter.

This chapter also includes information on how to change your engine settings following your initial configuration and how to upgrade the [Extreme Application Sensor Engine](#) engine software.

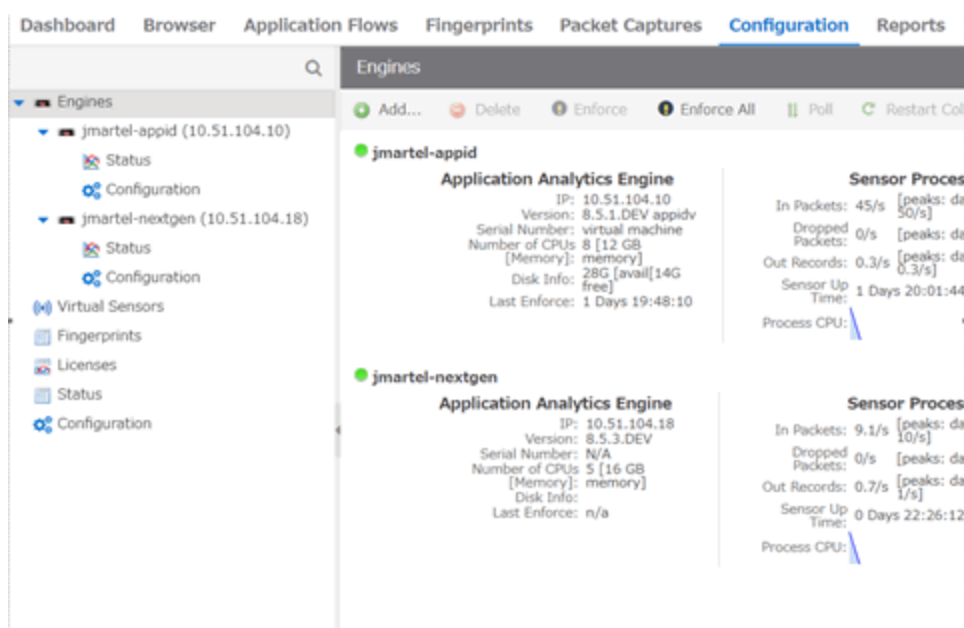
Extreme Application Sensor Engine Virtual Engine Requirements

OVA	CPUs	Memory (GB)	Disk (GB)	Maximum Number of Monitoring Interfaces Supported
Small	8	12	40	1
Medium	16	24	440	2
Large	24	36	960	3


Adding the Extreme Application Sensor Engine Engine

Use the following procedure to add the Extreme Application Sensor Engine engine to the **Analytics** tab in ExtremeCloud™ IQ - Site Engine:

1. Select the **Analytics > Configuration** tab.
2. Expand the **Engines** tab in the left-panel tree.



3. Select the Add button to open the **Add Application Sensor & Analytics Engine** window.



Add Application Analytics Engine ✕

IP Address:

Name:

Profile:

After adding an engine, go to the engine's Configuration page to add a wireless controller flow source, enable Access Control integration, or change the default web credentials.

OK Cancel

4. Enter the **IP Address** and the **Name** of the Extreme Application Sensor Engine engine.
5. Select the appropriate SNMP Profile from the **Profile** drop-down list.
6. Select **OK**.
7. Select **Enforce Engine** from the drop-down list.

The Extreme Application Sensor Engine engine is added to ExtremeCloud IQ - Site Engine.

Changing Extreme Application Sensor Engine Engine Settings

Use these steps if you need to change your Extreme Application Sensor Engine engine settings following your initial engine configuration.

Changing Basic Network Configuration

To change basic network configuration settings such as hostname and engine IP address, enter the following command at the engine CLI:

```
/usr/postinstall/dnetconfig
```

This starts the network configuration script and allows you to make the required changes. You must reboot the engine for the new settings to take effect.

Changing Date and Time Settings

To enable or disable using NTP to configure the engine date and time, or to manually set the date and time on the engine, enter the following command at the engine CLI:

```
/usr/postinstall/dateconfig
```

This starts the date and time configuration script and allows you to change the settings.

Changing the ExtremeCloud IQ - Site Engine Server IP Address

To change the IP address of the ExtremeCloud IQ - Site Engine server, enter the following command at the engine CLI:

```
/opt/appid/configMgmtIP <IP address>
```

Then, start using the new ExtremeCloud IQ - Site Engine server by typing:

```
systemctl restart analytics
```

Changing SNMP Configuration

To change SNMP configuration settings such as SNMP Trap Community String, SNMP User, SNMP Authentication, and SNMP Privacy credentials, enter the following command at the engine CLI:

```
/usr/postinstall/snmpconfig
```

This starts the SNMP configuration script and allows you to make the required changes.

Configuring the Application Sensor & Analytics Engine

After the initial engine installation is complete, use the following steps to configure the virtual engine to run the Extreme Application Sensor Engine engine:

1. Access the Application Sensor Analytics engine:

```
Welcome to the Extreme Networks Application Analytics Engine 8.5.3.22

applicationanalytics login: root
Last login: Wed Oct 28 07:48:39 EDT 2020 on tty1
*** Extreme Networks ***

This is the Application Analytics Engine 8.5.3.22. Alter files with caution.

WWW Site:      http://www.extremenetworks.com
Support Email: support@extremenetworks.com
Phone:         800-938-2408

*****
^[[P^[[P
```

2. Login as root with no password, and press **[Enter]**. The following screen displays:

```
=====
Extreme Networks - Analytics Server Appliance -
Welcome to the Analytics Server Engine 8.5.3.34 Setup
=====
Please enter the information as it is requested to continue with
the configuration. Typically a default value is displayed in brackets.
Pressing the [enter] key without entering a new value will use the
bracketed value and proceed to the next item.

If a default value cannot be provided, the prompt will indicate that the item
is either (Required) or (Optional). The [enter] key may be pressed without
entering data for (Optional) items. A value must be entered for (Required) items.

At the end of the setup process, the existing settings will be displayed
and opportunity will be provided to correct any errors.
=====

Press [enter] to begin setup or CTRL-C to exit: _
```

3. Press **[Enter]** to begin the setup. The **Root Password Configuration** screen displays:

```
=====
Root Password Configuration
=====
There is currently no password set on the system administrator account (root).
It is recommended that you set one so that it is active the first
time the machine is rebooted.
=====

Would you like to set a root password (y/n) [y]?

Enter new UNIX password:
Retype new UNIX password: _
```

NOTE: You must set a new root password. This new root password will be used by the initial user when logging in to the Extreme Application Sensor Engine engine.

4. Enter **y** to set the new root password.
5. Press **[Enter]**. Enter and retype the new password as prompted.
6. The **Extreme Application Sensor and Analytics Engine Network Configuration** screen displays. For each line, type the requested configuration information and press **[Enter]**.

```

=====
Analytics Server Appliance Network Configuration
=====
Enter the hostname for the appliance (Required): dtrue-nganalytics

Enter the IP address for dtrue-nganalytics [192.168.1.1]: 10.54.165.201
Enter the IP netmask [255.255.255.0]:

Enter the gateway address [10.54.165.1]:

Enter the IP address of the name server (Optional): 134.141.79.192
Enter the IP address of an alternate name server (Optional):

Enter the domain name for dtrue-nganalytics (Required): extremenetworks.com_

```

7. A screen asking you to confirm your network setting displays:

```

=====
Confirm Network Settings
=====
These are the settings you have entered. Enter 0 or any key other than a
valid selection to continue. If you need to make a change, enter the
appropriate number now or run the /usr/postinstall/dnetconfig script at a
later time.
=====

0. Accept settings and continue
1. Hostname:          dtrue-nganalytics
2. IP address:       10.54.165.201
3. Netmask:          255.255.255.0
4. Gateway:          10.54.165.1
5. Nameserver:       134.141.79.192
6. Domain name:      extremenetworks.com
7. NIS Server/Domain:

Enter selection [0]: _

```

8. In the SNMP Configuration screen, type the requested information for each line and press [Enter].

```

=====
SNMP Configuration
=====
These are the current SNMP V3 settings. To accept them and complete
SNMP configuration, enter 0 or any key other than the selection choices.
If you need to make a change, enter the appropriate number now or
run the /usr/postinstall/snmpconfig script at a later time.

0. Accept the current settings
1. SNMP User:          IP
2. SNMP Authentication Protocol: MD5
3. SNMP Authentication: snmpauthcred
4. SNMP Privacy Protocol: DES
5. SNMP Privacy:       snmpprivcred
6. Modify all settings
=====

Enter selection [0]: _

```

9. Enter 0 to accept your SNMP Configuration settings.

10. In the **Configure Date and Time Settings** screen, select whether you want to use an external Network Time Protocol (NTP) server. Enter **y** to use NTP, and enter your NTP server IP address(es). Enter **n** to configure the date and time manually.

```

=====
Configure Date And Time Settings
=====
The engine date and time can be set manually or using an external
Network Time Protocol (NTP) server. It is strongly recommended that
NTP is used to configure the date and time to ensure accuracy of time
values for SNMP communications and logged events. Up to 5
server IP addresses may be entered if NTP is used.
=====

Do you want to use NTP (y/n) [y]?

Please enter a NTP Server IP Address (Required): 134.141.79.191

Would you like to add another server (y/n) [n]? _

```

11. In the **NTP Servers validate selection** screen, enter **0** to accept the current settings.

```

=====
NTP Servers
=====
These are the currently specified NTP servers:

134.141.79.191

Enter 0 or any key other than a valid selection to complete NTP configuration and continue.
If you need to make a change, enter the appropriate number from the
choices listed below.

0. Accept the current settings and continue
1. Restart NTP server selection
2. Set date and time manually
=====

Enter selection [0]:

```

12. In the **Set Time Zone** screen, select the appropriate time zone and press **[Enter]**.

```

=====
Set Time Zone
=====
You will now be asked to enter the time zone information for this system.
Available time zones are stored in files in the /usr/share/zoneinfo directory.
Please select from one of the following example time zones:

1. US Eastern
2. US Central
3. US Mountain
4. US Pacific
5. Other - Shows a graphical list
=====

Enter selection [1]: _

```

13. The **Modify Settings** screen summarizes the settings you have entered and provides an opportunity to modify the settings, if desired. Enter **0** to accept the settings.

```

=====
Modify Settings
=====
All of the information needed to complete the installation of the Analytics Server
Engine has been entered. Enter 0 or any key other than a valid selection
to continue. If you need to make a change, enter the appropriate number from
the choices listed below.
=====

0. Accept settings and continue
1. Set the root user password
2. Set host name and network settings
3. Set SNMP settings
4. Set the system time
5. Modify all settings
Enter selection [0]:

```

The Extreme Application Sensor Engine engine software is automatically installed. This can take a few minutes. When the installation is complete, you see the following screen.

```

=====
Extreme Networks - Analytics Server Appliance - Setup Complete
=====
Setup of the Analytics Server Engine is now complete. Details of the engine setup
process are located in log files in the /var/log/install directory.
=====

root@true-nganalytics:~$

```

Deploying the Extreme Application Sensor Engine Engine

You will need to deploy the Extreme Application Sensor Engine engine in order to fully install it and use it with ExtremeCloud IQ - Site Engine's Analytics functionality.

VMWare TAP Interface

If you are using a vSwitch's TAP interface, you must enable promiscuous mode on the vSwitch to enable the Extreme Application Sensor Engine engine to capture packets. Promiscuous mode, which is typically used for packet sniffing, will enable the virtual Extreme Application Sensor Engine engine to see all of the mirrored traffic. By default, promiscuous mode for a newly created vSwitch is disabled (that is, set to Reject) as shown in the following figure.

For instructions on how to enable promiscuous mode on the vSwitch, visit <https://kb.vmware.com/s/article/1004099>.

Deploying Extreme Application Sensor Engine Engine in an MSP or MSSP Environment

This Help topic presents instructions for deploying Extreme Application Sensor Engine engine within an MSP (Managed Service Provider) or MSSP (Managed Security Service Provider) environment.

Configuring ExtremeCloud IQ - Site Engine Behind a NAT Router

If the ExtremeCloud IQ - Site Engine server is located behind a NAT (Network Address Translation) router, use the following steps to add an entry to the `nat_config.txt` file that defines the real IP address for the ExtremeCloud IQ - Site Engine server. This allows the ExtremeCloud IQ - Site Engine server to convert the NAT IP address received in the ExtremeAnalytics engine response to the real IP address used by the ExtremeCloud IQ - Site Engine server. Not adding the real IP address for the ExtremeCloud IQ - Site Engine server to the `nat_config.txt` file results in the Extreme Application Sensor Engine engine incorrectly displaying a state of **IMPAIRED** (orange) rather than **UP** (green).

NOTE: The text in the `nat_config.txt` file refers to a remote IP address and a local IP address. For this configuration, the NAT IP address is the remote IP address and the real IP address is the local IP address.

1. On the ExtremeCloud IQ - Site Engine server, add the following entry to the `<install directory>/appdata/nat_config.txt` file.
`<NAT IP address>=<real IP address>`
2. Save the file.
3. If the ExtremeCloud IQ - Site Engine Management server IP address is not configured to use the NAT IP address of the ExtremeCloud IQ - Site Engine server, perform the following steps:
 - a. Enter the following command at the engine CLI:
`/opt/appid/configMgmtIP <IP address>`
Where `<IP address>` is the NAT IP address of the ExtremeCloud IQ - Site Engine server.
Press **Enter**.
 - b. Restart the appidserver after the new IP address is configured by typing:
`appidctl restart`
Press **Enter**.
4. On the ExtremeCloud IQ - Site Engine server, add the following text to the `<install directory>/appdata/NSJBoss.properties` file. In the second to last line, specify the hostname of the ExtremeCloud IQ - Site Engine server.

NOTE: The Extreme Application Sensor Engine engine functions as a client computer independent of the server. Both engines and clients must be able to resolve the hostname you specify.

>In order to connect to a ExtremeCloud IQ - Site Engine server behind a NAT firewall or a ExtremeCloud

IQ - Site Engine server with multiple interfaces you must define these two variables on the ExtremeCloud IQ - Site Engine server. The `java.rmi.server.hostname` should be the hostname (not the IP) if multiple IPs are being used so that each client can resolve the hostname to the correct IP that they want to use as the IP to connect to.

```
java.rmi.server.hostname=<hostname of ExtremeCloud IQ - Site Engine server>
java.rmi.server.useLocalHostname=true
```

5. Save the file.
6. Add the ExtremeCloud IQ - Site Engine server hostname to your DNS server, if necessary.

NOTE: Extreme Application Sensor Engine engines, remote ExtremeCloud IQ - Site Engine clients, and any ExtremeControl engines must be able to connect to ExtremeCloud IQ - Site Engine using this hostname.

Launching ExtremeAnalytics

Now that you have configured the Extreme Application Sensor Engine engine, you are ready to access the ExtremeCloud IQ - Site Engine Launch Page and run the applications from a remote client machine.

1. Open a browser window on the remote client machine and enter the ExtremeCloud IQ - Site Engine Launch page URL in the following format:`http://<servername>:8080/`

where `<servername>` is the ExtremeCloud IQ - Site Engine engine IP address or hostname, and 8080 is the required port number. For example,
`http://10.20.30.40:8080/`

The ExtremeCloud IQ - Site Engine Launch Page opens.

2. Enter your ExtremeCloud IQ - Site Engine username and password and select **Login**.
3. Select the **Analytics** tab at the top of the window.

The **Analytics** tab displays.

Upgrading Extreme Application Sensor Engine Engine Software

Upgrades to the Extreme Application Sensor Engine engine software will be made available from the Network Management Suite (NMS) Download web page.

1. Download the Extreme Application Sensor Engine Engine Image 64bit (ZIP) file to your system.

To download an engine image:

1. Access the Extreme Portal at: <https://extremeportal.force.com/>.
2. After entering your email address and password, you are on the Support page.
3. Select the **Products** tab and select **Extreme Application Sensor and Analytics Engine**.
4. Select **Extreme Application Sensor and Analytics Engine** in the right-panel.
5. Select a version.
6. Download the following image file and extract the file to a directory on your system:
`Extreme Application Sensor Engine Engine Upgrade (BIN)`
2. Use FTP, SCP, or a shared mount point, to copy the upgrade file to the Extreme Application Sensor Engine engine.
3. SSH to the engine.
4. Cd to the directory where you downloaded the upgrade file. For example, enter the following to change to the /Users/jsmith directory: `cd /Users/jsmith`
5. Change the permissions on the upgrade file by entering the following command: `chmod 755 purview_appliance_upgrade_to_version.bin`
6. Run the install program by entering the following command: `./purview_appliance_upgrade_to_version.bin`

The upgrade begins automatically.

The ExtremeAnalytics engine restarts automatically when the upgrade is complete. Because your Extreme Application Sensor Engine engine settings were migrated, you are not required to perform any configuration on the engine following the upgrade.

Getting Started with Extreme Application Sensor Engine

This topic provides information to help you get started using Extreme Application Sensor Engine to view network application data in the **Analytics** tab. It includes information on ExtremeAnalytics access requirements, configuring the ExtremeAnalytics engine, enabling NetFlow flow collection, and configuring network locations.

ExtremeAnalytics Access Requirements

Both the Extreme Application Sensor Engine and the **Analytics** tab require the ExtremeCloud IQ - Site Engine Advanced (NMS-ADV) license. Contact your sales representative for information on obtaining an ExtremeCloud IQ - Site Engine Advanced license.

In order to view the **Analytics** tab, you must be a member of an authorization group assigned the ExtremeCloud IQ - Site Engine ExtremeAnalytics Read Access or Read/Write Access capability. The Read Access capability allows the ability to access the **Analytics** tab and view the ExtremeAnalytics reports. The Read/Write capability adds the ability to configure Extreme Application Sensor Engines and NetFlow Collecting devices. It also adds the ability to create and modify fingerprints.

ExtremeAnalytics Engine Configuration

The Extreme Application Sensor Engine engine provides the engine to monitor and classify layer 7 application information and reports that information to ExtremeCloud IQ - Site Engine, where it is managed and displayed in the **Analytics** tab.

The Extreme Application Sensor Engine engine must be installed and running on your network. Following installation, the Extreme Application Sensor Engine engine must be added to ExtremeCloud IQ - Site Engine and enforced via the **Configuration** tab in the **Analytics** tab.

- [Configuration - Analytics](#)

ExtremeAnalytics Tab Overview

The **ExtremeAnalytics** tab allows you to view and customize its [dashboard](#) and [browser](#), as well as ExtremeAnalytics [reports](#), [fingerprints](#), [packet captures](#), and [application flow](#) data. You can also manage and configure your ExtremeAnalytics engines.

NOTE: ExtremeAnalytics reports and application flow data is not available unless an ExtremeAnalytics engine is configured and you are a member of an [authorization group](#) assigned the ExtremeCloud IQ - Site Engine ExtremeAnalytics Read Access or Read/Write Access [capability](#). The Read Access capability allows the ability to access the **Analytics** tab and view the ExtremeAnalytics reports. The Read/Write capability adds the ability to configure ExtremeAnalytics engines and NetFlow Collecting devices. It also adds the ability to create and modify fingerprints.

Viewing ExtremeAnalytics application data requires certain [access requirements](#) and prerequisites. Both the ExtremeAnalytics feature and the **Analytics** tab require the ExtremeCloud IQ - Site Engine Advanced (NMS-ADV) license. Contact your sales representative for information on obtaining an ExtremeCloud IQ - Site Engine Advanced license.

Dashboard

The [Dashboard](#) tab displays an overview of application usage on your network through a series of graphs. It allows you to view network activity statistics based on client/server, application, industry, and response time for the specified ExtremeAnalytics engine. Many of the reports are links to more detailed pages.

Browser

The [Browser](#) tab lets you query information about recent network activity stored in the ExtremeCloud IQ - Site Engine database and display results in various grid and chart report formats. Using the Browser, you can create custom queries based on selected options including a data target, statistic type, and other search criteria.

Application Flows

You can choose from the **View** drop-down list to show you several options in the table on the Application Flows tab, including the latest flows from the specified ExtremeAnalytics engine, the worst network and application response times, classified and unclassified flows, and flows during a specified time frame. The table presents bidirectional flow data (aggregate flows) or unidirectional flow data (base flows).

Fingerprints

A [fingerprint](#) is a description of a pattern of network traffic which can be used to identify an application. The **Fingerprints** tab provides detailed information about fingerprints used by ExtremeAnalytics to identify application flows. You can choose to view in-use and customized fingerprint data.

Packet Captures

Use the [Packet Captures](#) tab to analyze the packets from the flows displayed on the **Application Flows** tab. The packet captures you create are presented in a table, which allows you to view details about the packet capture. Additionally, using this tab you can select a packet capture and view it in a packet analyzer.

Configuration

The [Configuration tab](#) provides detailed information on the ExtremeAnalytics engines you configure. It also lets you add and enforce your engines, and access engine reports and diagnostics. You must be a member of an authorization group assigned the ExtremeCloud IQ - Site Engine ExtremeAnalytics Read/Write Access capability to view the **Configuration** tab.

Reports

On the [Reports tab](#), you can access a selection of reports that provide detailed information on application usage on your network, as well as network activity statistics based on application, user name, client, and site. For many of the reports, you can select an item in the report to view details or right-click an item to select from other focused reports.

- [Dashboard Overview](#)
- [Browser](#)
- [Application Flows](#)
- [Fingerprints](#)
- [Packet Captures](#)
- [Configuration](#)
- [Reports](#)

ExtremeAnalytics Dashboard Overview

Accessible from the **Analytics** tab in ExtremeCloud IQ - Site Engine, the **Dashboard** tab displays an overview of application usage on your network, as well as network activity statistics through a series of real-time reports. The Dashboard is flexible and customizable - you can choose the reports and the design of the page to meet your specific needs. Many of the reports are links to more detailed pages.

The Dashboard includes a drop-down list with links to additional report dashboards:

- [Insights](#)
- [Client/Server](#)
- [Applications Browser](#)
- [Industry](#)
- [Response Time](#)
- [Network Service](#)
- [Tracked Applications](#)

Several report pages can be launched in the **Reports > Reports Designer** view in ExtremeCloud IQ - Site Engine by selecting the **Launch in Report Designer** icon ().

Insights Dashboard Reports

The Insights dashboard displays graphs with real-time network and application usage and service data, and tools that you can use to customize the dashboard using drag-and-drop capabilities.

Five ring charts display real-time Engine, Virtual Sensors, Disk Usage, License Usage, Network Response , and Application Response usage and service data. The ring charts are links to additional data. The Network Response and Application Response charts link to the [Network Service Response Time](#) and [Tracked Application Response Time](#) report dashboards, respectively.

Use the Custom Dashboard to drag and drop only the graphs you want on your dashboard. Each graph is a real-time preview and many are linked to additional detail reports. You can also choose whether the graphs in the Application Group area are organized in columns or rows in the Custom Dashboard area.

Client/Server Dashboard Reports

This dashboard displays reports on clients and servers seen on the network over the last 24 hours. It also displays reports on top clients by bandwidth, flow, or number of applications, and top servers by bandwidth or flow.

Select the **Info** icon () at the top right of the dashboard page to read a description of each report.

Applications Browser Dashboard Report

The Application Browser Dashboard displays bubble maps for top applications by bytes and flows, top profiles by bytes, and top sites by bytes. Place your cursor over a bubble to display bandwidth use or the number of flows. Use the drop-down menus to change the start date and time for the reports.

Drill-down for more information by selecting an application bubble to open a new graph of clients, flows, and usage data for that application. In that graph, select a client link to view application data for that client.

Industry Dashboards

Select the Industry Dashboard from the Dashboard drop-down list to access the following additional dashboards:

Enterprise Dashboard

The Enterprise Dashboard displays application information specific to the Enterprise network, including social applications, storage applications and cloud, business applications and email, and network applications and protocols.

Education Dashboard

The Education Dashboard displays application information specific to the campus network, including learning management systems, P2P, streaming, and social applications.

Healthcare Dashboard

The Healthcare Dashboard displays applications used in the healthcare environment, including patient care, medical applications, and HIPAA.

Venue Dashboard

The Venue Dashboard displays data grouped according to sports, social media, news and weather applications, as well as software update applications.

Response Time Dashboard

The Response Time Dashboard displays the response time in milliseconds of application data grouped by different criteria, selected from the drop-down list. The data is displayed as a line graph, which is updated periodically.

Network Service Dashboard

The Network Service Dashboard displays the response time of network services for the top five worst-performing sites as well as the overall average of all sites. The data for each network service at a site is displayed as a bar and line graph, which is updated periodically.

Tracked Applications Dashboard

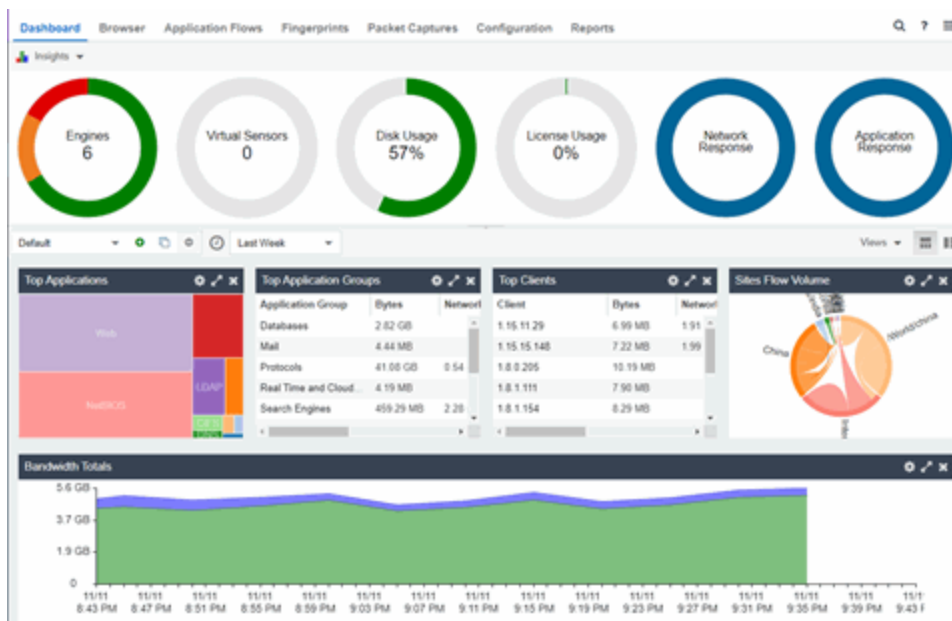
The Tracked Applications Dashboard displays the response time of the applications you configure in the **Tracked Applications** field on the **Analytics > Configuration > Configuration tab**. The data for each network service at a site is displayed as a bar and line graph, which is updated periodically. You can choose to organize the graphs in either columns (||||) or rows (■).

- [ExtremeAnalytics tab](#)

ExtremeAnalytics Insights Dashboard

Accessible from the **Analytics** tab in ExtremeCloud IQ - Site Engine, the Insights Dashboard displays an overview of application usage on your network, as well as network activity statistics based on client/server, application, industry, and response time.

Use the Insights Dashboard to view graphs that display real-time network and application usage and service data, and tools that you can use to customize your dashboard using drag-and-drop capabilities.




Insights

The Insights Dashboard displays ring charts and a customizable Application Group Dashboard. You can collapse and expand the ring charts and Application Group Dashboard for flexible display capabilities.

Ring Chart

Six ring charts display real-time Engines, Virtual Sensors, Disk Usage, Flow Rate, Network, and Application usage and service data:





- **Engines** — The number at the center of the ring chart indicates how many engines are represented by the chart. The colors in the graph indicate the states of the configured engines. Hover over a ring color to display a tooltip with the status of that engine. Select the graph to display overview and status details.
- **Virtual Sensors** — The number at the center of the ring chart indicates how many virtual sensors are represented by the chart. The colors in the graph indicate the states of the configured virtual sensors. Hover over a ring color to display a tooltip with the status of that virtual sensor. Select the graph to display overview and status details. Select the graph to open the [Virtual Sensors tab](#).
- **Disk Usage** — The number at the center of the ring chart indicates the percentage of Disk Usage. The colors in the graph display the percentage of disk usage being used. Hover over the ring color to display a tooltip with usage percentage and units of space details.

Select the graph to open the **Configuration** tab, where you can configure the information displayed in the Insights Dashboard.

- **Flow Rate** — The number at the center of the ring chart indicates the flow rate percentage. The colors in the graph indicate the flow rates for the different engines being used. Hover over a ring color to display a tooltip with status, percentage and rate details for each engine. Select the graph to open the **Licenses** tab.
- **Network Response** — The colors in the graph indicate the network response time for the application/site. Hover over a ring color to display a tooltip with status details and the number of networks at that status. Select a color in the graph to open the Network Service dashboard, which displays network service details.
- **Application Response** — The colors in the graph indicate the application response time for the application/site. Hover over a ring color to display a tooltip with response time details and the number of applications within the expected response time range. Select a color in the graph to open the Response Time dashboard, which displays network and application response time charts and details.

Custom Dashboard

The Custom Dashboard is a customizable space for viewing graphs that you select from the **Views** drop-down list. The buttons at the top right of the Applications Group dashboard ( ) enable you to save and copy your dashboard.



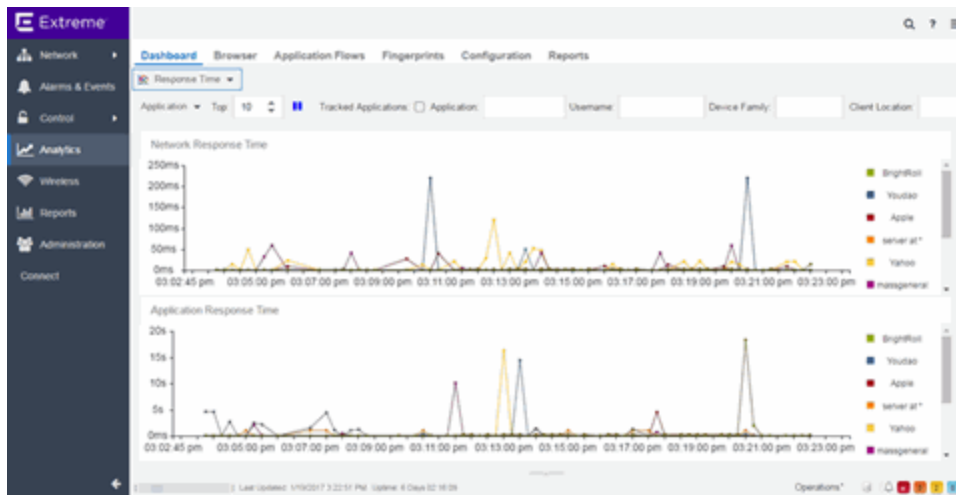
- [Analytics Tab](#)
- [How to Use the Application Group Dashboard](#)

ExtremeAnalytics Response Time Dashboard

The Response Time Dashboard displays the network and application response time data for the slowest targets on your network based on response time for the last 20 minutes. Use the graph to view response time data for a variety of filters, including application, device family, and username.

Additionally, you can use the dashboard to select the number of targets for which the response time is displayed and you can filter the information based on certain criteria and view flow data specific to the data you select.


To access the Response Time Dashboard, open the **Analytics > Dashboard** tab and select **Response Time** in the dashboard drop-down list.



Overview

The Response Time Dashboard contains two graphs, one displays the [network response time](#) and the other displays the [application response time](#). Data is updated every 15 seconds and displays data over the last 20 minutes.


If you have multiple ExtremeAnalytics engines, use the **Engine** drop-down list to select an engine to use as the source for the report data.



Use the toolbar at the top of the window to display data based on criteria you select and updates the two graphs.

Application

Use the **Application** drop-down list to group the data in the Response Time Dashboard by the following criteria:



Application ▼ Top: 10

- Total
- Application
- Application Group
- Client Site
- Server Site
- Device Family
- Username
- Client Address
- Server Address
- Input Interface
- Output Interface
- Cloud Region

Top

Use the **Top** field to limit the results in the graphs to display only the top results based on the number you enter.

For example, you can configure the graphs to display the top 3 slowest applications by response time.

Tracked Applications


Select the **Tracked Applications** box to add response time results for tracked applications to the Network Response Time and Application Response Time graphs.

Filters

You can also use the filter options at the top of the window to search for specific criteria. Using these fields limits the data to Tracked Applications, Application, Username, Device Family, Client Site, and Server Site. Entering a value in one of these fields filters the results displayed in the graphs below. Clear the data by selecting the **Clear** (⊗) button to the right of the filter options.

Network Response Time Graph

The Network Response Time graph displays the response time (in milliseconds) the TCP request took to complete for the Top N slowest Targets. The data in this graph depends on the criteria you select in the toolbar at the top of the window and can be [filtered](#) to match specific criteria. ExtremeCloud IQ - Site Engine displays data collected by the ExtremeAnalytics engine over the previous 20 minutes updated every 15 seconds. Use the **Pause** button in the toolbar to stop the graph from updating. Selecting the **Unpause** button resumes the updates and refreshes the graph with the most up-to-date data.



Place your cursor over a point in the graph to see a pop-up with details about that application at that moment in time.

Selecting a point opens a flow data table for that Target at that time at the bottom of the window, limited to match any [filters](#) you applied. Right-click a row in the flow to see additional options for working with that flow. Flows without an identified source are labeled with the device's IP Address.

Select the **Arrow** button (▼▼▼) at the top of the flow data table to collapse the table and select the **Arrow** button (▲▲▲) on the collapsed table to expand the table again.

Application Response Time Graph

The Application Response Time graph displays the response time (in milliseconds) the application request took to complete for the Top N slowest Targets. The data in this graph depends on the criteria you select in the toolbar at the top of the window and can be [filtered](#) to match specific criteria. ExtremeCloud IQ - Site Engine displays data collected by the ExtremeAnalytics engine over the previous 20 minutes updated every 15 seconds. Use the **Pause** button in the toolbar to stop the graph from updating. Selecting the **Unpause** button resumes the updates and refreshes the graph with the most up-to-date data.



Place your cursor over a point in the graph to see a pop-up with details about that application at that moment in time.

Selecting on a point opens a flow data table for that Target at that time at the bottom of the window, limited to match any [filters](#) you applied. Right-click a row in the flow to see additional options for working with that flow.

Select the **Arrow** button (▼▼▼) at the top of the flow data table to collapse the table and select the **Arrow** button (▲▲▲) on the collapsed table to expand the table again.

- [ExtremeAnalytics](#)

Network Service



Top: 5 Last Week Min. Data Required: 100 Views

ALL	DHCP	DNS	Kerber...	LDAP	RADIUS
ALL 156ms 80ms 157ms	ALL 153ms 116ms 186ms		ALL 148ms 90ms 193ms	ALL 156ms 80ms 157ms	ALL 156ms 107ms 193ms
Loc 150 185ms 99ms 184ms	Loc 200 145ms 116ms 155ms		Loc 160 161ms 140ms 171ms	Loc 1054 167ms 142ms 178ms	Loc 180 156ms 144ms 157ms
Loc 1054 174ms 80ms 191ms	Loc 160 143ms 121ms 151ms		PrivateA...pac 90ms 175ms	Loc 180 163ms 121ms 178ms	Loc 1054 174ms 150ms 174ms
PrivateA...pac 80ms 176ms	Loc 150 167ms 142ms 184ms		Loc 200 136ms 120ms 147ms	Loc 160 172ms 107ms 190ms	Loc 160 160ms 107ms 174ms
Loc 160 160ms 107ms 191ms	Loc 1054 169ms 143ms 182ms		Loc 170 173ms 150ms 195ms	PrivateA...pac 80ms 160ms	Loc 150 185ms 152ms 205ms
Loc 180 158ms 107ms 191ms	Loc 134141 11ms 107ms 191ms		Loc 140 155ms 107ms 191ms	Loc 170 163ms 107ms 191ms	PrivateA...pac 107ms 191ms

27 of 59

- DNS
- Kerberos
- LDAP
- RADIUS

The top graphs for each service displays the average response time of all of the sites for that service, while the following rows indicate the top worst performing sites for that service.

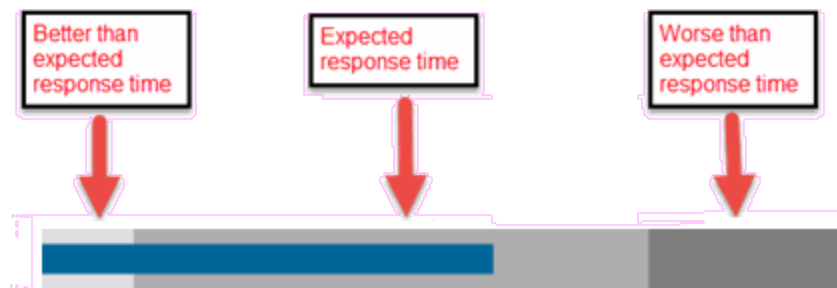
You can display or hide any of the application columns using the **Views** drop-down list. You can also select the **X** at the top of a column to hide the column from the dashboard. Select the **Single Row** icon () to display all columns in a single row, or select the **Double Row** icon () to


display the columns in two rows.

The worst performing sites are defined as those whose response time is the slowest when compared to the expected response time observed over the selected time period. For example, a site with an average RADIUS authentication response time of 40 ms over the past seven days that displayed a slowest response time of 50 ms would rank as a better performing site than a site with an average RADIUS authentication response time of 5 ms over the same period that displayed a slowest response time of 30 ms.

Expected Response Time

The Expected Response Time bar graph displays the range of response times, the most recently measured response time, and the expected response time for a network service a specific site during the date range you configure in the Date Range drop-down list. The value displayed on the far right of the graph is the slowest response time observed during the selected time period. The vertical green bar indicates the most recently observed response time for the network service.



Hover over the Expected Response Time graph to display a pop-up with the response time for the network service as well as the date and time the measurement occurred. The Expected Response Time bar graphs also display the client count, represented by a number and a monitor icon (), and a client byte count observed as of the most recent measured minute. The client count is the number of clients using the service at the site. The client byte count indicates the amount of storage being utilized by clients. The data used for the client count, the client byte count, and the reported response time are from the same recently observed minute.

NOTE: Client counts and client byte counts are not provided for the bar graphs that display the average response time of all the sites for that service.

ExtremeCloud IQ - Site Engine uses a standard deviation of the values gathered as response times to determine the expected response time for a network service at a site. In the bar graph, the medium gray color indicates a response time that falls within the "expected" range. A response time in the light gray range is better than expected, while a response time in the dark gray is worse than expected.

When a response time is determined to be worse than expected, the site name and the response time indicator turn red to flag the service.

Selecting the Expected Response Time bar graph opens the Response Time dashboard (which is also accessible from the **Analytics > Dashboard** tab) filtered to display the network service. If you select the network service for a particular site, the Response Time dashboard also filters to that site.

Historical Response Time

The Historical Response Time line graph shows all of the response times observed for the network service at a site.



Placing your cursor over a point in the graph causes a dot on the line graph to appear, indicating the point in the response time at which you are looking. Additionally, a pop-up with the date, time, and response time displays for that point.

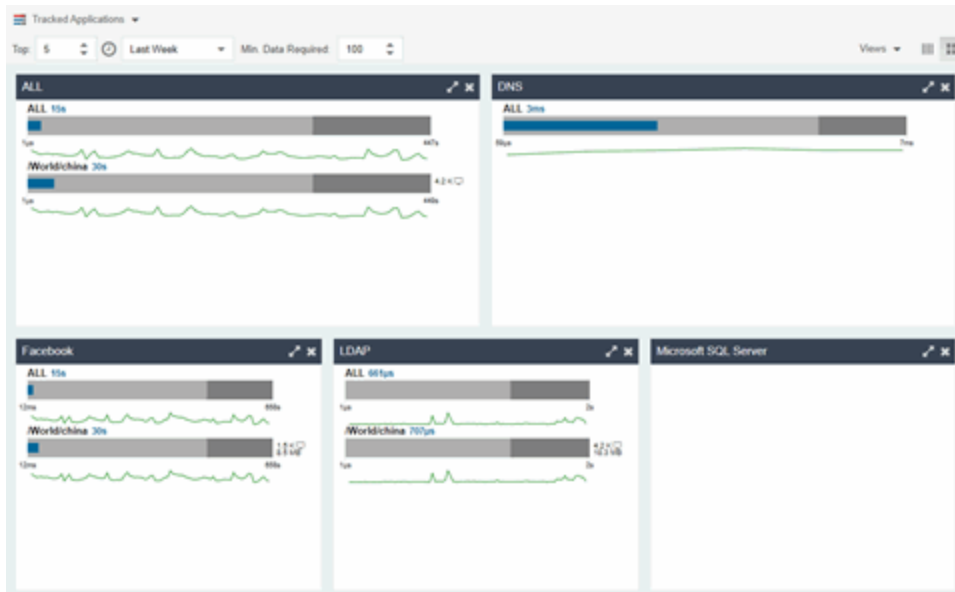
This is the data set from which ExtremeCloud IQ - Site Engine creates the Expected Response Time graph. The wider the expected response time range in the Expected Response Time graph (indicated by the medium gray color), the greater the variance in the values in this graph.

- [ExtremeAnalytics tab](#)

ExtremeAnalytics Tracked Applications Dashboard

The Tracked Application dashboard displays the performance (in response time) of your network for applications you configure in the **Tracked Applications** field on the **Analytics > Configuration > Configuration** tab.

To access the Tracked Application dashboard, open the **Analytics > Dashboard** tab and select **Tracked Applications** in the dashboard drop-down list.



Overview

The Tracked Applications dashboard contains two graphs for each application, one displays the average response time over the selected time period and the other displays the individual response times over that period for each site. Data is updated every minute and can be manually refreshed by selecting the **Refresh** button (↻).

Select the number of sites displayed in each column in the **Top** field. The Tracked Applications dashboard can display up to 25 sites.

Use the **Time Period** drop-down list to display the date and time range for which data is displayed. Selecting **Custom** displays additional fields allowing you to indicate a **Start Date** and time and an **End Date** and time.

Use the **Minimum Required Response Time Dashboard Data Points** to configure the minimum amount of data ExtremeCloud IQ - Site Engine requires before displaying a given application or site pair. The data below this threshold is not reliable and can set off a false alarm, however, you can adjust how much data is required based on the individual needs of your network.

Each column in the dashboard represents an application. The top row displays the average response time of all of the sites for that application, while the following rows indicate the top worst performing sites for that application.

You can display or hide any of the application columns using the **Views** drop-down list. You can also select the **X** at the top of a column to hide the column from the dashboard. Select the **Single Row** icon (☐) to display all columns in a single row, or select the **Double Row** icon (☐) to display the columns in two rows.


Select the **Maximize** icon (🔍) to expand a single application column.

The worst performing sites are defined as those whose response time is the slowest when compared to the expected response time observed over the selected time period. For example, a site with an average Microsoft Office 365 authentication response time of 40 ms over the past seven days that displayed a slowest response time of 50 ms would rank as a better performing site than a site with an average Microsoft Office 365 authentication response time of 5 ms over the same period that displayed a slowest response time of 30 ms.

Expected Response Time

The Expected Response Time bar graph displays the range of response times, the most recently measured response time, and the expected response time for an application a specific site during the date range you configure in the Date Range drop-down list. The value displayed on the far right of the graph is the slowest response time observed during the selected time period. The vertical blue or red bar indicates the most recently observed response time for the application.

NOTE: The values in this graph are an average of all response times observed every minute.



Hover over the Expected Response Time graph to display a pop-up with the most recent response time for the application as well as the date and time the measurement occurred. The Expected Response Time bar graphs also display the client count, represented by a number and a monitor icon (10 🖥️), and a client byte count observed as of the most recent measured minute. The client count is the number of clients using the service at the site. The client byte count indicates the amount of storage being utilized by clients. The data used for the client count, the client byte count, and the reported application response time are from the same recently observed minute.

NOTE: Client counts and client byte counts are not provided for the bar graphs that display the average application response time of all the sites for that service.

ExtremeCloud IQ - Site Engine uses the standard deviation of the values gathered as response times to determine the expected response time for an application at a site. In the bar graph, the medium gray color indicates a response time that falls within the "expected" range. This range is the average value of all observed response times plus or minus two standard deviations, or about 95 percent of all response time values. A response time in the light gray range is better than expected, while a response time in the dark gray is worse than expected.

When a response time is determined to be worse than expected, the site name and the response time indicator turn red to flag the application.



Selecting the Expected Response Time bar graph opens the Response Time dashboard filtered to display the application. If you select the application for a particular site, the Response Time dashboard also filters to that site.

Historical Response Time

The Historical Response Time line graph shows all of the response times observed for the application at a site.

NOTE: The values in this graph are an average of all response times observed every hour.



Hovering over a point in the graph causes a dot on the line graph to appear, indicating the point in the response time at which you are looking. Additionally, a pop-up with the date, time, and response time displays for that point.

This is the data set from which ExtremeCloud IQ - Site Engine creates the Expected Response Time graph. The wider the expected response time range in the Expected Response Time graph (indicated by the medium gray color), the greater the variance in the values in this graph.

- [ExtremeAnalytics tab](#)

ExtremeAnalytics Browser Overview

The **Browser** tab lets you query information about recent network activity stored in the ExtremeCloud IQ - Site Engine database and display results in various grid and chart report formats. Using the Browser, you can create custom queries that provide greater flexibility in defining what data to display and how to display it. You can access the Browser from the ExtremeCloud IQ - Site Engine **Analytics** tab.


Overview

The Browser allows you to generate reports in several different formats using data based on selected options including a data target, statistic type, start time, and other search criteria.

For example, you can display application response time for the last hour or the last three days. You can view the results as a grid or a chart. You can filter the results to display data for a specific application or site.

If you have multiple ExtremeAnalytics engines, use the **Engine** drop-down list to select an engine to use as the source for the report data. Then, select the desired options on the left side of the Browser view and select **Submit**. The report is displayed on the right side of the view. Select an item in the report to view details or right-click an item to select from other focused reports.

After you have generated a report, use the **Gear** menu () (at the bottom left of the options panel) to ( **Save**) [save it to the Report Designer](#) to use as a custom component, () [bookmark the report](#), or () [export it as a CSV file](#).



The screenshot displays the ExtremeAnalytics Browser interface. The top navigation bar includes links for Dashboard, Browser (selected), Application Flows, Fingerprints, Packet Captures, Configuration, and Reports. The left sidebar contains the Options and Search Criteria panels. The Options panel includes sections for Data Table (set to End-System Details - Hourly), Display Format (set to Grid), Target (set to Applications), Time Period (set to Last Interval), Statistic Type (set to Bytes), and Aggregation (radio buttons for Sum and Average). The Search Criteria panel includes dropdown menus for Site, Profile, Application Group, Device Family, and Engine (all set to All), and text input fields for User Name, Application, Client, and Limit (set to 10). A Submit button is located at the bottom right of the Search Criteria panel.

Data Aggregation

Network data displayed in a report is aggregated from your network by the ExtremeAnalytics engine and sent to ExtremeCloud IQ - Site Engine. The data gathering process begins with the ExtremeAnalytics engine, which monitors network activity on the switch or controller you configure using a traffic mirror and NetFlow or application telemetry. The traffic mirror gathers the first (N) packets of a flow to determine the application in use, while NetFlow (a flow-based

data collection protocol) provides information about the amount of data sent and received for the application. The engine holds this information in its cache and transmits the aggregated data to ExtremeCloud IQ - Site Engine every five minutes to update the High-Rate data table information and every hour to update the hourly data table information. Creating a report in the Applications Browser displays the information sent from the ExtremeAnalytics engine to ExtremeCloud IQ - Site Engine based on the criteria you select.

NOTE: Information held in the ExtremeAnalytics engine's cache is not saved. Restarting the ExtremeAnalytics engine before the data in the memory cache is sent to ExtremeCloud IQ - Site Engine results in the loss of that information.

Options

Following are definitions of the different options available when creating your custom query.

Data Table

Select which type of network activity data to query. The correct data table to use depends on the nature of the report.

- **End-System Details - Hourly** — End-system data collected every hour. Used when data for a specific client or server is needed, or when the information requested is highly specific, for example top applications used by Android devices in the London site.
- **Application Data - Hourly** — Application data collected every hour. Used for higher level information, such as top applications during an hour.
- **Application Data - High-Rate** — Application data collected at a higher rate (every five minutes). Used for a more detailed picture of how traffic changes over time.
- **Application Telemetry - Hourly** — Application Telemetry flow data collected every hour.

Display Format

Select the display format for the report: Grid, Chart Over Time, Word Cloud, Tree Map, or Bubble Map. If you select Chart Over Time as your report display format, you can select whether to display the data as a line or an area, and also select the color to use in the chart.

Target

Network traffic information is collected on objects in your network called targets. Some targets are physical, such as clients and servers, and some are logical, such as applications. Select the type of target that you want information about. Available targets vary depending on the selected data table. If you want information on a specific target, specify that target in the Search Criteria options.

- **Applications** — An application in ExtremeAnalytics is identified through layer 7 analysis of network traffic. For example, an application can be identified as Facebook.
- **Application/Client** — Information about applications used by clients, or about clients using an application.

- **Application/Device Family** — Information about applications used by device families, or about device families using an application.
- **Application/Interface** — Information about the applications used by interfaces.
- **Application/Profile** — Information about applications used by profiles, or about profiles using an application.
- **Application/Server** — Information about applications accessed on a particular server, or about servers using an application.
- **Application Groups** — Application categories, such as Cloud Computing or Social Networking, which are implied by the application.
- **Device Family** — The kind of device determined for a client, such as Windows or iOS. Device information is only available for some network traffic.
- **Interface/Applications** — Information about interfaces used by applications.
- **Application-Interface Pair/Client** — Displays the applications and interfaces used by clients.
- **Interface/Client** — Information about the interfaces used by clients.
- **Sites** — [Sites](#) are used by ExtremeAnalytics to identify the physical location for the client of an application flow. A site is a set of IP address ranges that identify a portion of your network. Multiple sites can be created to identify different buildings, sites, or geographical areas of your network.
- **Profiles** — A profile assigned to a client. Profile information is only collected under certain circumstances.
- **Threat** — Displays a list of the threat classifications that occurred during the **Time Period** you select.
- **Threat/Threat End-System Pair** — Displays a list of the threat classifications broken down by the IP addresses of the end-systems involved in the flow (the trusted and untrusted hosts) that occurred during the **Time Period** you select.
- **Clients** — The end-point of a flow which has the client role for that connection.
- **Servers** — The end-point of a flow which has the server role for that connection.
- **Total** — The total values for all detected traffic for the interval used by the data table (hourly or high-rate).

Time Period

Select the time duration for the report: Last Interval, Today, Yesterday, Last 24 Hours, Last 3 Days, or Last Week. You can also specify a custom start time and end time for the report. The Last Interval is the most recent recorded data covering a time period determined by the selected Data Table.

Statistic

Statistics are quantitative data that can be collected for the selected target. Available statistics vary depending on the selected target. Select the desired statistic for the report:

- **Bytes** — The number of bytes transferred in both directions, between the client and the server. Also known as bandwidth.

- **Flows** — The number of NetFlow records sent by the switch to report the traffic between the client and the server.
- **Application Response Time** — The average amount of time for a server to respond to a request.
- **Network Response Time** — The average amount of time to create a connection.
- **Received Bytes** — The number of bytes received by clients. This can be an estimated number of bytes if you are using an Application Telemetry flow.
- **Sent Bytes** — The number of bytes sent by clients. This can be an estimated number of bytes if you are using an Application Telemetry flow.
- **Inbound Flows** — The number of NetFlow records sent by the switch to report the server-to-client traffic. This is a rough indication of the duration of client connections.
- **Outbound Flows** — The number of NetFlow records sent by the switch to report the client-to-server traffic. This is a rough indication of the duration of client connections.
- **Clients** — The number of unique clients that have been seen associated with the target.
- **Servers** — The number of unique servers that have been seen associated with the target.
- **Application Count** — The number of unique applications seen for the selected target.

For byte, flow, and application count statistics, if you select a time range that is larger than the interval, specify whether you want the data aggregated as a summation of all the values for that statistic or as an average of all the values for that statistic.

Search Criteria

Defining search criteria allows you to further filter the report data. Available criteria will vary depending on the selected data table and target. If you select either of the Application Data tables, you can only filter based on the selected target. For example, if you select **Sites** as your target, you can only filter on defined sites. If you select the End-System Details data table, you can filter on additional criteria. For example, if you select **Sites** as your target, you can filter on defined sites as well as flows for iOS devices.


You can enter a partial term in the text field or use the SQL wildcard "%" (as a substitute for multiple characters) or "_" (as a substitute for a single character) for multiple matches. For example, for the Device Family name, you could enter "iPhone %" to match iPhone 3, 4, and 5.

NOTE: Values entered in the text fields that contain multiple, non-alphanumeric characters can cause issues with the returned results. If this happens, alternate values should be used.


- **Site** — Select a site to match or select World. If a site has been added to a map, you will also see a selection for that map. If you select custom, you can enter a partial site name or use the SQL wildcard characters to match one or more sites.
- **Profile** — Select an ExtremeControl profile to match or select All. If you select custom, you can enter a partial profile name or use the SQL wildcard characters to match one or more profiles. Profile information is only collected under certain circumstances.

- **Application Group** — Select an application group to match or select All. If you select custom, you can enter a partial application group name or use the SQL wildcard characters to match one or more groups.
- **Device Family** — Select the operating system family to match or select All. If you select custom, you can enter a partial device family name or use the SQL wildcard characters to match one or more families. Device information is only available for some network traffic.
- **User Name** — Enter a client's username to match. Username information is only available for some network traffic.
- **Application** — Enter an application name to match.
- **Client** — Enter a client's IP address or hostname to match.
- **Engine** — Select the ExtremeAnalytics engine for which you are generating the report.
- **Limit** — Select the number of results to return, for example, 10 clients.


Bookmark

After you have generated a report, select the Gear menu () in the lower left corner to save the options you have currently set. A new window opens for the current report with a link that can be bookmarked in your browser. You can then use the bookmark whenever you want the same search options.

Save to Report Designer

Select the Gear menu () in the lower left corner to access the Save to Report Designer window. This window lets you save the currently defined report to use as a custom component in the Report Designer. The custom component uses the target, statistic, and start time currently defined in the Browser.

Enter a name for the custom component and select any search criteria that you want displayed in the component panel. The search criteria is displayed as fields in the component panel, providing a custom interface that lets you further refine report data. If no search criteria are selected, the saved component only uses the target, statistic, and start time definitions when requesting data, creating a view-only report.

A dialog box titled "Save to Report Designer" with a close button (X) in the top right corner. The dialog contains instructional text, a note, a text input field for "Name:", and a list of checkboxes for various fields. At the bottom are buttons for "Select All", "Deselect All", "OK", and "Cancel".

Save to Report Designer

Create a custom report component for use in Report Designer. The saved definition will use the currently selected data target and statistic to get analytics data, and will use the fields selected below to provide a user-interactive way to refine the data results.



NOTE: Fields containing a selected or custom value will use that value as a default, but values containing multiple, non-alphanumeric characters may prevent a query from producing valid results. If that happens, alternate values should be used.

Name:

<input type="checkbox"/> Date/Time	<input type="checkbox"/> User Name
<input type="checkbox"/> Site	<input checked="" type="checkbox"/> Application
<input type="checkbox"/> Profile	<input type="checkbox"/> Client
<input type="checkbox"/> Application Group	<input type="checkbox"/> Engine
<input type="checkbox"/> Device Family	<input type="checkbox"/> Limit

Select All | Deselect All | OK | Cancel

Export to CSV

Select the Gear menu () in the lower left corner and select () to export the report data as a CSV file. The currently defined report opens in a spreadsheet, which can then be saved.

- [ExtremeAnalytics tab](#)

ExtremeAnalytics Application Flows

The **Application Flows** tab displays tables that present [Bidirectional](#) or [Unidirectional](#) client, server, and application flow data. To access the **Applications Flows** tab, open **Analytics > Application Flows**.

This Help topic provides information on the following topics:

- [Overview](#)
- [Application Flows Tables](#)
- [Report Features](#)

Overview

The **Application Flows** tab includes several functions that enable you to filter and customize your table data.

The screenshot shows the 'Application Flows' tab in the ExtremeAnalytics interface. The table displays network flow data with columns for Flows, Client Address, User, Server Address, Server Port, Application, and Application Group. The data is filtered to show 100 rows. The interface includes navigation tabs at the top (Dashboard, Browser, Application Flows, Fingerprints, Packet Captures, Configuration, Reports) and a search bar. The bottom status bar shows 'Max Rows: 100', 'Aggregate Flows: 0 Days 00:02:22 (Current Load: 0 fps / Peak 0 fps)', and 'Displaying 100 rows'.

Flows	Client Address	User	Server Address	Server Port	Application	Application Group
2				http	Undefined-1600005	Undefined
1				http	Google Maps	Location Services
2				http	Google Maps	Location Services
2				http	Ebay	E-commerce
1				http	Ebay	E-commerce
1				netbios-ssn	NetBIOS	Protocols
1				http	Undefined-1600005	Undefined
5				http	Undefined-1600005	Undefined
1				http	Google Maps	Location Services
1				http	Google Maps	Location Services
1				http	Ebay	E-commerce
1				http	Undefined-1600005	Undefined
1				http	Undefined-1600005	Undefined
2				http	Undefined-1600005	Undefined
2				http	Undefined-1600005	Undefined
1				http	Undefined-1600005	Undefined
1				http	Ebay	E-commerce
3				http	Undefined-1600005	Undefined
1				http	✓ Facebook	Social Networking
1				http	✓ Facebook	Social Networking
1				http	✓ Facebook	Social Networking
1				http	✓ Facebook	Social Networking
1				http	Undefined-1600005	Undefined
1				netbios-ssn	CIFS	Storage

Appliance Engine

If your network uses multiple ExtremeAnalytics engines, use the Engine menu to select an engine to use as the source for the flow data.

Bidirectional / Unidirectional

Select to display either [Bidirectional](#) (aggregate flows) or [Unidirectional](#) (base flows) flow data.

Show

Select from the drop-down list to filter flow data that displays. The available options vary depending the flow type (bidirectional or unidirectional) selected.

- All — Show all flows.
- Classified — Show only flows classified by an application fingerprint.
- Unclassified — Show only flows not classified by an application fingerprint.
- Unclassified Web Traffic — Show only web traffic that has not been classified by an application fingerprint.

Flows

By default, the table displays the latest flows collected. The available options vary depending the flow type (bidirectional or unidirectional) selected.

- Flows — Displays the latest flows collected by the specified engine.
- Flows After — Enables you to select a start date and time for the flows displayed.

- Worst Network Response Times — Sorts the flows based on the worst TCP response time and displays the flows with the worst time at the top of the chart.
- Worst Application Response Times — Sorts the flows based on the worst application response time and displays the flows with the worst time at the top of the chart.

Application Group

Use the **Application Group** menu to filter the table by application group.

Search

Use the **Search** field at the top right of the table to filter specific flow information. For example, searching on "snmp" or "10.20.30.131/24" filters the table so only flow data related to SNMP or the given subnet is displayed. You can enter one or more filters simultaneously, separated by semicolons. Individual components of a filter is separated by commas. For complete instructions on how to use the Flow Search, rest your cursor on the **Search** field and read the tooltip (select the "more" link in the tooltip). Press the **Reset** button at the bottom left of the window to clear the Search results and refresh the table.

You can also use the **Search** field to search for a specific application, user name, or IP address from your filtered results:

1. Select a user name or IP address from the filtered search results to launch PortView, which provides a detailed topology context for the user.
2. Enter **meta=** before the term for which you are searching includes all variations of that search term in the result set. For example, entering **meta=extreme** returns **extremenetworks.com**, **www.extremenetworks.com**, **extreme.boston.com**, and any other flows that include the word "extreme".
3. Right-click on a flow to access a menu of options including the ability to:
 - Add a new custom fingerprint based on the flow selected in the table.
 - Show all fingerprints associated with the application in the selected flow.
 - Create a UDP or TCP rule using the IP port.
 - Search ExtremeCloud IQ - Site Engine maps for the selected flow client.
 - Open a Flow Details report for the selected flow (bidirectional flows only).
 - Access a variety of reports for the flow.

Refresh

Use the **Refresh** drop-down list at the top right of the window to specify an interval (in seconds) at which the flows data automatically refreshes. To stop auto refresh, select the **Refresh Off** option.

Application Flows Tables

The columns included in the Application Flows tables vary, depending on the type of data flow you select (Historical, Bidirectional and Unidirectional). Additionally, right-click and select **Start Packet Capture** to save a packet capture of the flow on the **Packet Captures** tab.

Bidirectional Flows

The Bidirectional table displays bidirectional flow data stored in memory. It provides aggregated flow data for a given client, server, server port, application, and protocol. All matching flows are aggregated to show the flow count, total duration, amount of data transmitted, and additional information. The bidirectional report presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A check mark (✓) in the table denotes a tracked application or a tracked site.

Unidirectional Flows

The Unidirectional table displays unidirectional flow data stored in memory. It provides the raw non-aggregated flow data received from the flow sensors on the network. It presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A check mark (✓) in the table denotes a tracked application or a tracked site.

Report Features

The Application Flows tables include several report features and functions that enable you to drill down for more detailed application, site, response time, mapping and policy functions. The report features vary, depending on the type of data flow you select (Historical, Bidirectional and Unidirectional).

Interactive Tables

Manipulate table data in several ways to customize the view for your own needs:

- Select the column headings to **perform an ascending or descending sort** on the column data.
- **Hide or display different columns** by selecting a column heading drop-down arrow and selecting the column options from the menu.
- **Filter data in each column** by selecting a column heading drop-down arrow and using the Filters option on the menu.

The sort and filter functionality for these two tables behaves differently than for other ExtremeCloud IQ - Site Engine tables. In these tables, [Max Rows](#) are considered for display, and then sorting and filtering is applied to these rows. In other tables, sorting and filtering is applied to the entire table, and then Max Rows of the result is displayed. For example, if the Max Rows value is set to 50 and you create a filter for a specific IP address, only those 50 rows will be filtered for the IP, not all the flows maintained in memory on the server.

CSV Export

The [CSV Export button](#) enables you to save report data to a CSV file and to provide report data in table form.

Bookmark

Use the [Bookmark button](#) to save the search, sort, and filtering options you have currently set. It opens a new window for the current report with a link that can be bookmarked in your browser. You can then use the bookmark whenever you want the same search, sort, and filtering options.

Max Rows

By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.

Reset

The [Reset button](#) enables you to clear the search fields and all filters, and to refresh the table.

Aggregate / Base Flows

Aggregate Flows (bidirectional table) and Base Flows (unidirectional table) data uses an X number of days, hh:mm:ss format and includes Current Load and Peak Load calculations in flows per second.

- [ExtremeAnalytics tab](#)

ExtremeAnalytics Bidirectional Flow Table

This table on the **Application Flows** tab displays bidirectional flow data that is stored in memory. Use it to view aggregated flow data for a given client, server, server port, application, and protocol. All matching flows are aggregated to show the flow count, total duration, amount of data transmitted, and additional information. The bidirectional report presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A check mark (✓) in the table denotes a tracked application or a tracked site.



By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.

Text at the bottom of the table shows:

- The **CSV Export icon** - allows you to save report data to a CSV file and to provide report data in table form
- Aggregate Flows data - uses an X number of days, hh:mm:ss format and includes Current Load and Peak Load calculations in flows per second

Following are definitions for the table columns:

Flow Summary

Rest the cursor over the first column in the table and select the  arrow to open the **Flow Summary** window. Flow summary information can include response times, Uniform Resource Identifier, and header data for the flow. In the **Flow Summary** window, use the **Menu icon**  to access additional functionality, such as the ability to modify the application fingerprint or create a policy rule.

Flows

The number of base flows included in the aggregate flow. Select a link in the Flows column to open a **Flow Details** tab that displays the individual flows that contributed to the aggregate flow.

Client Address

The IP address or hostname of the system where the flow originated. Select the Client address link to open a **PortView** for the client (if it is in the database) or a **PortView** for the

switch configured as the NetFlow sensor.

Server Address

The IP address or hostname of the server handling the flow.

Server Port

Either the TCP or UDP port on the server handling the flow.

Application

The name of the application as identified by the ExtremeAnalytics engine using the Fingerprint database.

Application Group

The flow application group to which the application belongs.

Application Info

Additional information about the flow provided by the ExtremeAnalytics engine. Hover over the flow and a table of the information displays.

Type

The content type of a flow, such as sound, video, or text. Select the **Type** icon to open the flow's URI.

Network Response

The response time (in milliseconds) that it took for the TCP request to complete.

Application Response

The response time (in milliseconds) that it took the application request to complete.

Site

The name of the site that matches the client's IP address.

Detailed Site

The client's switch IP and switch port (wired), or controller IP, AP, and SSID (wireless).

Device Family

The operating system family for the client end-system.

User

The username used when the client system connected.

Profile

The ExtremeCloud IQ - Site Engine profile assigned to the client end-system.

Threat

Indicates if the flow contains potential threat activity from IP addresses known to be suspicious. IP addresses can be flagged as suspicious for a variety of reasons, including forced IP anonymity through the use of a Tor exit node, being listed as a threat by the Emerging Threats project, or classified as suspicious by internet users.

Protocol

The connection type protocol used by the flow.

Last Seen Time

The last time a unidirectional (base) flow was aggregated into this bidirectional flow.

Duration

The duration of a bidirectional (aggregate) flow is the sum of the durations of the unidirectional (base) flows that make up the bidirectional flow. The duration of a bidirectional flow can be greater than or less than the period of time indicated by the **First Seen** and **Last Seen Time**. This is because there can be times during that time period when no flow is active or when several flows are active at the same time.

NOTE: Bidirectional flows can be greater than the period of time between the **First Seen** and **Last Seen Time** columns because they display the sum of all flow records for a client and a server on a server port. For a flow that lasts for 60 seconds, there are two flow records (a client to server flow and a server to client flow), so the total duration can exceed 60 seconds. Multiple simultaneous connections from the client to the same server port (e.g. multiple browser windows open to a web-based email client) can also increase the duration.

Rate

The average bandwidth for the flow based on the total flow duration. Because bandwidth calculations are based on the total duration (not on the **First Seen** and **Last Seen Time**), they represent the average throughput for each flow considered separately, not as an aggregate.

Tx Packets

The number of packets transmitted for this flow. For flows collected via Application Telemetry, this number can be estimated.

Rx Packets

The number of packets received for this flow. For flows collected via Application Telemetry, this number can be estimated.

Tx Bytes

The number of bytes transmitted for this flow. For flows collected via Application Telemetry, this number can be estimated.

Rx Bytes

The number of bytes received for this flow. For flows collected via Application Telemetry, this number can be estimated.

Traffic Records

The number of records received in each flow.

Flow Source

The IP address of the NetFlow source switch, Application Telemetry source switch, or wireless controller sending the NetFlow data to the NetFlow collector.

Input Interface

The interface receiving the flow on the NetFlow sensor.

Output Interface

The interface transmitting the flow on the NetFlow sensor.

Client TOS

The DSCP (Diffserv Codepoint) value for the client to server flow. The TOS/DSCP value is used to configure quality of service for network traffic.

Server TOS

The DSCP (Diffserv Codepoint) value for the server to client flow. The TOS/DSCP value is used to configure quality of service for network traffic.

TTL

The TTL (IP Time to Live) value of the flow. The TTL field indicates the maximum number of router hops the packet can make before being discarded. The TTL field is set by the packet sender and reduced by every router on the route to its destination. When the value hits zero, the packet is dropped.

- [ExtremeAnalytics tab](#)
- [Application Flows tab](#)

ExtremeAnalytics Unidirectional Flow Table

This table on the **Application Flows** tab displays unidirectional flow data stored in memory. It displays the raw, non-aggregated flow data received from the flow sensors on the network. It presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A check mark (✓) in the table denotes a tracked application or a tracked site.



Hover over an application in the table to display switch data, which is an accumulation of multiple switches into single flow record, as well as the path that flow has taken.

By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.



Text at the bottom of the table shows Base Flows, using X number of days, hh:mm:ss format, and including Current Load and Peak Load calculations in flows per second.

Following are definitions for the table columns:

Flow Summary

Rest the cursor over the first column in the table and select the  arrow to open the **Flow Summary** window for a specific flow. Flow summary information can include response times, Uniform Resource Identifier, and header data for the flow. In the **Flow Summary** window, use the **Gear** menu  to access additional functionality such as the ability to modify the application fingerprint or create a policy rule.

Client/Server Flows

Identifies whether the flow is a Client Flow  or a Server Flow . The client/server direction of a flow is calculated by the ExtremeAnalytics engine. Hover over the icon to see a tooltip with more information.

Source Address

The IP address or hostname of the system where the flow originated. Select the Source address link to open a **PortView** for the client or server (if it is in the database) or a **PortView** for the switch configured as the NetFlow sensor.

Source Port

Either the TCP or UDP port on the client/server handling the flow.

Destination Address

The IP address or hostname of the system that received the flow.

Destination Port

Either the TCP or UDP port on the system that received the flow.

Application

The name of the application as identified by the ExtremeAnalytics engine using the Fingerprint database.

Application Group

The flow application group to which the application belongs.

Application Info

Additional information about the flow provided by the ExtremeAnalytics engine.

Type

The content type of a flow, such as sound, video, or text. Select the **Type** icon to open the flow's URI.

Network Response

The response time (in milliseconds) that it took for the TCP request to complete.

Application Response

The response time (in milliseconds) that it took the application request to complete.

Site

The site where the flow originated.

Detailed Site

The client's switch IP and switch port (wired), or controller IP, AP, and SSID (wireless).

Device Family

The operating system family for the client end-system.

User

The username used when the client system connected.

Profile

The ExtremeControl profile assigned to the client end-system.

Protocol

The connection type protocol used by the flow.

Last Seen Time

The last time the flow was seen.

Duration

The amount of time that the flow was active.

Rate

The average bandwidth for the flow based on the flow duration.

Packets

The number of packets in this flow. For flows collected via Application Telemetry, this number can be estimated.

Bytes

The number of bytes in this flow. For flows collected via Application Telemetry, this number can be estimated.

NetFlow Records

The number of NetFlow records for this flow.

Flow Source

The IP address of the NetFlow source switch, Application Telemetry source switch, or wireless controller sending the Flow data to the Flow collector.

Input Interface

The interface receiving the flow on the Flow sensor.

Output Interface

The interface transmitting the flow on the Flow sensor.

TOS

The DSCP (Diffserv Codepoint) value for the flow. The TOS/DSCP value is used to configure quality of service for network traffic.

TTL

The TTL (IP Time to Live) value of the flow. The TTL field indicates the maximum number of router hops the packet can make before being discarded. The TTL field is set by the packet sender and reduced by every router on the route to its destination. When the value hits zero, the packet is dropped.

- [ExtremeAnalytics tab](#)
- [Application Flows tab](#)

ExtremeAnalytics Fingerprints Overview

The **Fingerprints** tab provides detailed information about fingerprints used by ExtremeAnalytics to identify application flows. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. They can be created based on flow, application or

application group, or a destination address. For applications such as Facebook and Google, multiple fingerprints are included to capture the different ways these applications can be used.

Fingerprints are created and stored on the ExtremeCloud IQ - Site Engine server. When a fingerprint is changed or enabled, a flag is raised on the ExtremeAnalytics engine to show it needs enforcing. Access the Browser from the ExtremeCloud IQ - Site Engine **Analytics** tab.

There are two types of fingerprints: system fingerprints and custom fingerprints.

System fingerprints are provided by ExtremeCloud IQ - Site Engine. They cannot be deleted; however, they can be modified or disabled. When a system fingerprint is modified, it results in a new custom fingerprint that overrides the original system fingerprint.

Custom fingerprints are either new user-defined fingerprints or modifications of system fingerprints. Custom fingerprints can be deleted. If a custom fingerprint was overriding a system fingerprint, then deleting the custom fingerprint will reload the original system fingerprint.

Analytics Application Data Collection

The Extreme Application Sensor Engine and ExtremeAnalytics engines provide an application data collection function that collects and records information about network utilization. It includes:

- General Usage Collection — High-level application-centric data, collected hourly and in five-minute intervals.
- Extended Application Collection — Detailed data about all end-systems in the network, collected hourly.

Application data collection is based on network flow information. Network utilization for various objects in the network (called targets) is measured, collected, and used to create application data reports in ExtremeCloud IQ - Site Engine.

NOTE: Ensure at least 4GB of swap space is available for flow storage or impaired functionality can occur. Use the `free` command to verify the amount of available RAM on your Linux system.

This Help topic describes application data collection, including collection targets, statistics, and intervals. It also describes the different collectors used to perform the collection, as well as the sources for flow information.

Data Collection Overview

Application data collection is performed by the Extreme Application Sensor Engine and ExtremeAnalytics engines. The engines collect flow records from switches in your network. They then augment the collected flow data with detailed application information derived by network packet inspection, resulting in rich analytical data.

For example, if a NetFlow record reports 100 bytes transferred from client Workstation 1 to server Host A, then the collection process would add 100 bytes to the tally for Workstation 1, and 100 bytes to the separate tally for Host A. If the flow is identified as traffic for the Payroll

application, then 100 bytes would be added to another tally for Payroll as well. And finally, 100 bytes is added to another tally for the entire network. At the end of a collection interval, the totals for client Workstation 1, server Host A, the Payroll application, and the entire network are written to the database.

Data from network flows is collected in an aggregated form for a period of time (called a collection interval), and then stored in the ExtremeCloud IQ - Site Engine database. ExtremeCloud IQ - Site Engine uses this data to provide reports that show how your network is being utilized.

To conserve space on your ExtremeCloud IQ - Site Engine server hard drive, your Extreme Application Sensor Engine and ExtremeAnalytics engines only collect total flow records when the server hard drive drops below 10 GB of free space. If the ExtremeCloud IQ - Site Engine server hard drive drops an additional 1 GB (under 9 GB of free space), your Extreme Application Sensor Engine and ExtremeAnalytics engines stop collecting all flow data.

NOTE: To change the differential threshold (the additional amount of free space reduction after which all records stop being collected), edit the `RM_FREE_SPACE_MINIMUM_ALLOW_SUMMARY_KB` value in the `NSJBOSS.properties` file. The value is set to 1,000,000 KB by default, so the engine stops collecting all records when free space reaches 10GB - 1,000,000 KB = 9 GB.

Collection Targets

Flow data is collected on objects in your network called targets. Some targets are physical, such as clients and servers, and some are logical, such as applications.

The Extreme Application Sensor Engine and ExtremeAnalytics engines can track the following target types:

- Client — The end-point of a flow that has the client role for that connection.
- Server — The end-point of a flow that has the server role for that connection.
- Application — An application in ExtremeAnalytics, identified through layer 7 analysis (for example, Facebook).
- Application Group — Application categories, such as Cloud Computing or Social Networking.
- Site — The client's physical location on the network, based on its IP address. [Sites](#) are used by ExtremeAnalytics to identify the physical location for the client of an application flow.
- Device Family — The kind of device determined for a client, such as Windows or iOS.
- Profile — An ExtremeControl profile assigned to a client.

In some cases, the engines can also track combinations of targets. For example, it can track the total number of bytes transferred from Workstation 1 for the Payroll application separately from Workstation 2 for Payroll, and from Workstation 1 for Facebook. These target and sub-target pairs provide for ExtremeCloud IQ - Site Engine drill-down reports, for example, reports to show the top Payroll clients or the top applications for Workstation 1.

Collection Statistics

Collection statistics are quantitative data that can be collected for a target. This includes statistics directly reported in NetFlow records, such as bytes transferred, as well as information that can be derived indirectly, such as the number of unique clients seen using an application.

The engines can track the following statistics:

- Bytes — The number of bytes transferred in both directions, between the client and the server. Also known as bandwidth. You can track sent and received bytes as well as total bytes.
- Flows — The number of NetFlow records sent by the switch to report the traffic between the client and the server. You can track inbound and outbound flows as well as total flows.
- Clients — The number of unique clients associated with the target.
- Applications — The number of unique applications associated with the target.
- Network Response Time — The average amount of time to create a connection.
- Application Response Time — The average amount of time for a server to respond to a request.

Collection Intervals

The Extreme Application Sensor Engine and ExtremeAnalytics engines collect and aggregate flow data for a period of time called an interval. At the end of the interval, the engines write the totals to the ExtremeCloud IQ - Site Engine database and a new interval begins, with new totals collected starting at zero.

Some statistics are collected and written to the database on an hourly interval. Other statistics are collected at a high-rate interval of every five minutes, providing for a more detailed picture of how traffic changes over time.

All statistics can be collected over multiple intervals and averaged. When viewing report data, it is important to know the interval used for any average that is displayed.

Certain statistics, such as bytes and flows, can be collected over multiple intervals to provide a total over time, while other statistics, such as client count, cannot. To illustrate, the number of bytes seen in two hours would be the total of the number of bytes seen in each hour. However, the number of unique clients seen in two hours would not be the total of the number of unique clients seen in each hour, as some clients were probably seen in both hours.

Using Sites to Collect In-Network Traffic

While flow data collection can aggregate data for all flow traffic that is visible, it can be more useful to aggregate data for *in-network* flows only. These are flows used by clients that are located in your internal network. By collecting data for only in-network flows, the overhead of aggregating data over an interval can be reduced.

You can define your internal network by configuring [sites](#). A site is a set of IP masks that defines a well-known portion of your internal network. You can use the World site to identify your entire internal network. If you have already reserved certain IP address ranges for certain physical sites on your network, you can create multiple sites that correspond to these reserved IP ranges.

Multiple sites can be created to identify different buildings, sites, or geographical areas of your network. Any IP that matches any site is considered to be in-network. If you define multiple sites, you will be able to analyze data broken down by site.

Data Collector Types

There are two kinds of data collectors used in Extreme Application Sensor Engine and ExtremeAnalytics.

- General Usage Collectors — These are hourly and high-rate collectors that record the top targets during an interval. Many types of targets and target-pairs are supported.
- End-System Details Collector — This is an hourly collector that attempts to capture and record data for all in-network clients and servers that it detects. All traffic collected is tagged with site, profile, device family, and other attributes.

Data from these collectors is stored separately in the database. The collector data used in a report depends on the nature of the report. Higher-level information, such as top applications during an hour, will be based on general usage collector data, since it is relatively inexpensive to access. End-system details data might be used when data for a specific client or server is needed, or when the information requested is highly specific, for example, top applications used by Android devices in the London site.

General Usage Collectors

General usage collectors collect data about all instances of a target for the interval, and then record only the most significant targets (typically, the 100 most significant targets).

When the top targets are calculated for a collection interval, several different statistics can be used as a basis for choosing the most significant entries. For example, collectors can record the top applications based on bytes, and also record the top applications based on number of clients. For each type of target collected, there are different sets of bases used.

General usage collectors operate at both hourly and high-rate intervals. They can collect data from all flows or from in-network flows only.

Hourly General Usage Collectors

The following table describes the hourly data collected by the general usage collectors.

Target	Sub-Target	Bases	Traffic Used
Total			In-Network Flows/ All Flows

Target	Sub-Target	Bases	Traffic Used
Application		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Clients Network Response Time Application Response Time	In-Network Flows
Application	Client	Bytes	In-Network Flows
Application Group		Bytes Flows Clients	In-Network Flows
Client		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Applications Network Response Time Application Response Time	All Flows
Device Family		Bytes Flows Clients	In-Network Flows
Site		Bytes Flows Clients Network Response Time Application Response Time	In-Network Flows
Profile		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Network Response Time Application Response Time	In-Network Flows

Target	Sub-Target	Bases	Traffic Used
Threat		Bytes Flows Application Response Time Network Response Time Received Bytes Sent Bytes Inbound Flows Outbound Flows	In-Network Flows
Threat	Threat End-System Pair	Bytes Flows Application Response Time Network Response Time Received Bytes Sent Bytes Inbound Flows Outbound Flows	In-Network Flows
Server		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Network Response Time Application Response Time	All Flows
Application	Device Family	Bytes Flows Clients	In-Network Flows
Application	Profile	Bytes Flows Clients	In-Network Flows

High-Rate General Usage Collectors

The following table describes the high-rate data collected by the general usage collectors.

Target	Sub-Target	Bases	Traffic Used
Total			In-Network Flows/ All Flows
Application		Bytes Flows Clients	In-Network Flows

Target	Sub-Target	Bases	Traffic Used
Application Group		Bytes Flows Clients	In-Network Flows
Device Family		Bytes Flows Clients	In-Network Flows
Site		Bytes Flows Clients	In-Network Flows
Profile		Bytes Flows Clients	In-Network Flows

End-System Details Collector

The end-system details collector tracks client/application target pairs.

Unlike general usage collectors, this collector attempts to record data for all in-network clients and servers it sees during the hour. For each client or server, it records data for up to 10 applications, plus an "other" category to capture the remaining traffic. Information such as location, device family, and profile are also recorded for each end-system.

The large number of targets recorded each hour and the amount of detail recorded for each one, can result in a large volume of data being stored in the database. In order to prevent disk space from being over-utilized, there is a total limit of 50,000 clients which can be recorded each hour across all Extreme Application Sensor Engine and ExtremeAnalytics engines. There is also a 25,000 client limit per engine for most license types. However, if you have an NMS-ADV license without any ExtremeAnalytics license, the per-hour total limit is 100 clients across all Extreme Application Sensor Engine and ExtremeAnalytics engines.

Flow Information Sources

The ExtremeAnalytics engine uses NetFlow or SFlow records from the switches and wireless controllers in your network as a source for flow data. Information such as IP addresses, ports, and bytes transferred comes from this flow data source.

This data is augmented with additional layer 7 application information produced by the Extreme Application Sensor Engine and ExtremeAnalytics engines through deep packet inspection. Information such as application name and network response time comes from this source.

There is additional information that can be obtained from sources other than NetFlow/SFlow records and deep packet inspection.

NOTE: Most of these sources rely on ExtremeControl data. If ExtremeControl is part of your network configuration, then ExtremeControl integration can be enabled (see [instructions](#) below) to provide access to these sources. Site data is obtained from [sites](#) configured in ExtremeCloud IQ - Site Engine.

The following is a list of information that can be obtained from different sources:

- **Hostname** — The client or server's hostname can be derived using ExtremeControl. ExtremeControl integration must be enabled.
- **Site** — The site for a flow is the site of the client in the flow. Client and server sites are derived from the [sites](#) configured on the **Network** tab. If a client does not match a site, then the site is empty. If a flow has a site, the flow is considered to be in-network.
- **Detailed Site** — Detailed site information is derived from the switch and port information resolved for the client end-system. ExtremeControl Integration must be enabled.
- **Device Family** — The device family is a general description of the operating system detected in the client, for example, Windows, Linux, or Android. The device family is derived from network packet inspection. The device family can also be provided by ExtremeControl, if ExtremeControl integration is enabled.
- **Profile** — The client's profile is derived from the ExtremeControl profile assigned to the client end-system. ExtremeControl integration must be enabled.
- **Username** — The client's username is derived from network packet inspection. The username can also be provided by ExtremeControl, if ExtremeControl integration is enabled.

It is possible that different sources can provide different values for the same information. For example, network packet inspection can provide the device family name of Window 7, whereas ExtremeControl can provide the device family name of Windows.

Enabling ExtremeControl Integration

If your network configuration includes ExtremeControl, ExtremeControl data can be integrated with flow data to provide additional information. ExtremeControl integration is only useful if you are collecting flows for end-systems managed by ExtremeControl.

When ExtremeControl integration is enabled, if a client in a flow matches an end-system in ExtremeControl, then:

- The client hostname in the flow is derived from the end-system.
- The device family in the flow is derived from the end-system.
- The username in the flow is derived from the end-system.
- The profile in the flow is derived from the end-system's ExtremeControl profile.
- The detailed site in the flow is derived from end-system data.

If a server in a flow matches an end-system in ExtremeControl, then:

- The server hostname in the flow is derived from the end-system.

To enable ExtremeControl integration on the Extreme Application Sensor Engine and ExtremeAnalytics engines:

1. If the ExtremeControl distributed end-system cache is not enabled on the ExtremeCloud IQ - Site Engine server, you must enable it using the following steps.
 - a. Select **Administration > Options** from the menu bar to open the **Access Control Options** window.
 - b. Select **Advanced Settings**.
 - c. In the End-System Mobility section, select the **Enable distributed end-system cache** option.
 - d. Select the **Reload** button to reload the cache configuration on the ExtremeCloud IQ - Site Engine server. Select **OK**.
2. Enable ExtremeControl Integration on each ExtremeAnalytics engine where you want to use ExtremeControl data.
 - a. Access the **Analytics** tab.
 - b. Expand each Extreme Application Sensor Engine and ExtremeAnalytics engine and select **Advanced Configuration**. In the right panel under Configuration Options, select the **Enable ExtremeControl Integration** option.
 - c. If your ExtremeControl engines are using Communication Channels, you must select the **ExtremeControl Communication Channel** option and enter the channel name. The ExtremeAnalytics engine is only able to access end-systems in its channel.
 - d. Select **Save**.
 - e. Enforce your ExtremeAnalytics engines.

Reports

Data gathered from flow usage collection is the basis of many reports in the ExtremeCloud IQ - Site Engine's **Analytics** tab. When collection is enabled, these reports begin to exhibit data.

Dashboard Report

The main Dashboard report contains data produced by the hourly General Usage collectors, and displays data for a specific hour. Across the top are the hour's totals. Below them are Top Application Groups, as a chart, and Top Applications, as a table, for the same hour. There is also Application Group Usage over the last 3 days, as a chart and as a table.

Note that data from the Extreme Application Sensor Engine and different ExtremeAnalytics engines is maintained separately. If you have the Extreme Application Sensor Engine and more than one ExtremeAnalytics engine, you need to select which engine to view, using the engine menu in the top-left corner.

Browser Reports

The Browser provides special reports that lets you select the targets, statistics, and collection interval for your report, as well as define search criteria to further filter report data. Using the Browser, you can create custom queries that provide greater flexibility in defining what data to display and how to display it. When you create a Browser report, you select which type of

network activity data to use: end-system details (always hourly), application data hourly, or application data high-rate. For additional information, see [Applications Browser](#).

- [Getting Started with ExtremeAnalytics](#)
- [Analytics Tab](#)

Glossary

C

CVE-
CVE-

M

My Term
My definition

Index

A

Adding Application Analytics Engine 7