ExtraHop 5.1
ExtraHop Web UI Guide

For more documentation, see https://docs.extrahop.com/.

Published: 2016-03-28

ExtraHop Networks
Seattle, WA 98101
877-333-9872 (US)
+44 (0)203 7016850 (EMEA)
+65-31585513 (APAC)
www.extrahop.com

Contents

Introduction to the ExtraHop Web UI

The ExtraHop Discover and Command appliances provide access to your network, application, client, and infrastructure data through a dynamic and highly customizable Web UI.

After you log into the ExtraHop appliance with a browser over HTTPS, you can immediately view your network activity through built-in system dashboards. If your environment includes a Command appliance, you can monitor all of the activity on your distributed Discover appliances from a single, centralized Command appliance.

The Web UI also enables you to:

- Create custom dashboard views of your network traffic in real time for the information that is most relevant to you. For more information, see Dashboards.
- Configure threshold and trend based alerts that notify you when there is a potential issue with a network device. For more information, see Alerts.
- View and drill down into real-time metrics about protocol activity on your network. For more information, see Metrics and Drill-down functionality.
- Generate reports on network metrics during a particular time interval, and export the information to PDFs or as CSV data. For more information, see Reports.
- Create custom applications and custom metrics, and track metrics for proprietary traffic. For more information, see Customizing ExtraHop appliances.

In addition, if your ExtraHop Discover appliance is paired to an ExtraHop Explore appliance, you can directly access stored transaction records through the Discover Web UI. Or, if you are monitoring multiple Discover appliances through a Command appliance, you can retrieve record information by node through the Command Web UI.

Discover and Command appliance administration tasks are available through the ExtraHop Admin UI for users with full administrator permissions. For more information, see the ExtraHop Admin UI Guide.

Explore appliance administration tasks are available through the ExtraHop Explore Admin UI. For more information, see the ExtraHop Explore Admin UI Guide.

The complete ExtraHop documentation set is available at https://docs.extrahop.com.

ExtraHop platform architecture

The ExtraHop platform comprises a suite of appliances that are designed to passively monitor the network traffic in your environment in real time. The ExtraHop system provides you with top-level and detailed metrics about the devices on your network, which you can analyze to determine where problems in your network might be developing.

ExtraHop Discover Appliance

The ExtraHop Discover appliance (EDA) provides the ability to analyze and visualize all of your network, application, client, infrastructure, and business data. The EDA passively collects unstructured wire data—all of the transactions on your network—and transforms this data into structured wire data.

Deploy a single EDA, either physical or virtual, anywhere in your network environment.

ExtraHop Explore Appliance

The ExtraHop Explore appliance (EXA) integrates with the ExtraHop Discover appliance to store transaction and flow records sent from the EDA. You can see, save, and search the structured flow and transaction information about events on your network with a simple, unified UI, with no modifications to your existing applications or infrastructure. Deploy a cluster of three or more EXA nodes to take advantage of data redundancy and performance improvements.

ExtraHop Command Appliance

The ExtraHop Command appliance (ECA) provides centralized management and reporting across multiple ExtraHop Discover appliances that are distributed across datacenters, branch offices, and the public cloud.

With an Explore appliance or cluster, you can pair the Explore appliance to multiple Discover appliances, and then query the records stored on each Discover node from the Command appliance.

For most large ExtraHop deployments, a dedicated ECA is the most efficient way to manage all of your remote appliances.

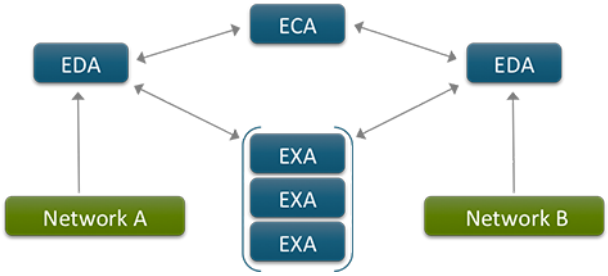

Metrics

The ExtraHop system provides a large number of protocol metrics. To search for a specific protocol metric, click **Metrics**, and then select metrics from the following sections:

Sources

View metrics for an application, device, or network. For more information, see the Sources section.

Groups

View metrics for a group of devices in activity groups or custom groups.

Records

View or query metrics associated with records. For more information, see the Records section.

You can also view metrics on the Dashboards page. For more information, see the Dashboards section.

Sources

You can view metrics that are associated with the following sources:

Applications

The ExtraHop system provides a default application, All Activity, which provides metrics for all transactions across all protocols for every device on your network. You can modify the default application template to suit the needs of your organization. You also can create your own application and assign specific devices and networks to that application. For more information, see the Applications section.

Devices

The ExtraHop system provides metrics about every discovered device on your network. For more information, see the Devices and Device Search sections.

Networks

A network is the entry point into the network capture, and metrics are collected for network capture attributes, network alerts, and network traffic details. These metrics provide a summary of all network activity retrieved in the capture. For more information, see the Networks section.

Device search

Understanding how a device is discovered by the ExtraHop system can help you search for a device.

The ExtraHop system automatically discovers devices based on what is happening on the network. ExtraHop has two device discovery modes: layer 2 (L2) discovery and layer 3 (L3) discovery.

L2 discovery

Detects an L2 address (MAC address).

L3 discovery

Detects an L3 address (IP address). This mode is the default and most commonly deployed discovery mode. The ExtraHop system discovers a device's IP address by monitoring ARP (Address Request Protocol) responses, and then matching the MAC address included in the response to a parent L2 device. The ExtraHop system maintains the linkage between L2 "parent" and L3 "child" devices.

After a device is discovered, the ExtraHop system tracks all of the wire data traffic associated with the device. Based on the type of traffic, the ExtraHop system assigns a device type to the device, such as a gateway, file server, database, or load balancer.

A device can have multiple names, which are all searchable. The ExtraHop system discovers device names by passively monitoring naming protocols, including DNS, DHCP, NETBIOS, and Cisco Discovery Protocol (CDP). If a name is not discovered through a naming protocol, the default name is derived from the device attributes (MAC address for L2 devices and IP address for L3 devices). You can also create a custom name for a device on the Devices page.

By default, all IP addresses that are observed outside of locally-monitored broadcast domains are aggregated at one of the incoming routers in your network. To identify and learn about individual devices outside of these routers, which are beyond your local network, you can create custom devices to enable reporting on these devices. For more information, see the Custom Devices section. You can also create a remote network that defines a range of IP addresses that are not on the local network. The ExtraHop system will then discover remote devices by their IP address.

Search for a device

You can filter searches to find on the Devices page with Find controls, which are located below the toolbar, to help you locate specific devices in the network capture.

By default, the search feature performs a substring search on the value entered in the **Find** text box. For example, if you submit the letter z for a name search, then the list of devices returned by the search includes all devices that have a letter z in the name, regardless of position.

If the search string value starts and ends with a forward slash (/), the portion of the input between the slashes is interpreted as a regular expression. The regular expression must be written in PostgreSQL syntax. Refer to PostgreSQL documentation for more information.

Searches can be filtered by the following device attributes:

any

Matches a substring in any device element.

ip address

Matches a substring in the device IP address. The IP address criteria can include CIDR notation in IP address/subnet prefix length format. For example, 10.10.0.0/16 for IPv4 networks or 2001:db8::/32 for IPv6 networks.

name

Matches a substring in the device name.

node

(Command appliance only) Matches a substring in the node name.

mac address

Matches a substring in the device MAC address.

tag

Matches a substring in the user-defined device tag.

type

Matches a substring to a specified device attribute type. When you select **type**, the **Find** text box becomes a drop-down list. In the **Find** drop-down list, select from the following:

Activity

Includes the metric types that were active in the selected time interval. For example, selecting **HTTP Server** returns devices with HTTP server metrics, and any other device with the custom type set to **HTTP Server**.

Device type

Includes Gateway, Firewall, Load Balancer, File Server, and Custom Device.

Class

Includes Node, Remote, Custom, and Pseudo.

vendor

Matches a substring in the device vendor name as determined by the Organizationally Unique Identifier (OUI) lookup.

vlan

Matches a substring in the device Virtual Local Area Network (VLAN) tag. VLAN information is extracted from VLAN tags, if the traffic mirroring process preserves them on the mirror port.

1. Navigate to a page that includes a device list, such as **Networks**, **Devices**, or **Groups**.
2. In the **Find** field, type search string characters.

   The device attribute filter is set to **any** by default, which applies the search string to all device attributes. You can adjust the search string to apply to a particular device attribute, such as the device name or the MAC address.
3. In the **by** drop-down list, select the device attribute that you want to search for.

   If you select the **type** attribute in your search, the **Find** field becomes a drop-down list that is populated with attributes to choose from.
4. Click **Search**.
   The device list is populated with the devices that match the search criteria.

Drill-down functionality

The ExtraHop Web UI enables you to drill down on metrics, so you can pivot from high-level information to identify root causes of interesting activity.

For example, metrics displayed on the Dashboards page show you the types of traffic in the network capture. If you see a spike in network traffic, you can isolate the protocol associated with the spike in a

chart, and then navigate to the Networks page to drill down on the protocol and examine which devices are associated with the spike in network traffic.

Note: If you have the ExtraHop Explore appliance connected to your network, you can also drill down on a protocol or metric using records. For more information, see the Records section.

Drill down on dashboard data

On the Dashboards page, if you see interesting activity in a chart, you can navigate from the chart data to the metric source page (which is an application, network, or device page), where you can drill down to see which devices are associated with the activity.

1. In a chart with interesting activity, do one of the following:
   - Click **Go>>** on a widget.
   - Click the command menu next to the **Go>>** button. Hover over **Go to *<metric source>***... and click the expanded menu option, which will take you to the associated metric source page.

     Note: For a chart with more than one metric source, click **Go to...** on the widget and select the associated metric source page from the drop-down menu. Or click the command menu next to the **Go to...** button and hover over **Go to…** in the drop-down menu. Click the expanded menu option, which will take you to the associated metric source page.
2. On the metric source page, hover over the charts to view any metrics that appear outside of the normal range or causing a spike in network traffic.
3. Click the protocol that is causing a spike in network traffic. Typically, there is a table beneath the charts that shows the list of devices associated with the protocol activity.
4. In the Device column of the table, click the name of the device that is causing the spike in network traffic. The Devices page appears with additional information.

Drill down on protocol page data

You can also drill down into metrics on individual metric source pages, such as applications, networks, or devices pages, to explore top-talking protocols and devices associated with them.

1. Click **Metrics**.
2. Click **Applications**, **Networks**, or **Devices**.
3. Select a specific application, network capture, or device from the center pane.
4. Click the protocol of interest in the left pane.
5. Hover over the charts to view any metrics that appear outside the normal range.
6. Click the data you find interesting in the chart to isolate data to a device or look for devices related to interesting data in a table beneath the chart.
7. Click the name of the device to view device details.

For example, to find the top-talkers for the L7 protocol:

1. Click **Metrics** and then click **Networks**.
2. Select a network in the center pane.
3. Click **L7 Protocols** in the left pane.
4. Hover over the charts to view any metrics that appear outside the normal range.
5. Click the legend on the graph to isolate the L7 protocol metrics. The Protocols table at the bottom of the page shows the list of applications associated with the selected protocol. From this list of devices, you can see which device is causing the spike in network traffic.
6. In the Device column of the table, click the name of the device that is causing the spike in network traffic. The Devices page appears with additional information.

Dashboards

Dashboards are customizable views of metrics that enable you to monitor and share data that are meaningful to you. On the Dashboards page, you can start by looking at general activity on your network through built-in system dashboards. Or you can build your own dashboard, with selected charts in a custom layout, to get a specific view of metrics for one or more devices, applications, or networks.

Dashboards are stored separately for each user that accesses the ExtraHop Discover appliance. When you log in, dashboards are organized in the following sections in the left pane, on the Dashboards page:

Dashboard Inbox

This section displays a list of dashboards that have been shared with you by other users. To share your dashboard with others, see the Share a Dashboard section.

My Dashboards

This section displays a list of dashboards that you created. These are the dashboards that you can keep private or share with other users. To create your own dashboard, see the Create a Dashboard and Configuring Dashboards sections. Editing access to your dashboard can be granted on a per-user basis. For more information, see Share a Dashboard section.

System Dashboards

This section displays the default built-in dashboards, which are the Activity dashboard and the Networks dashboard. These dashboards cannot be deleted, modified, or shared. When you visit the Dashboard page for the first time, the built-in Activity dashboard displays. This dashboard contains an overview of network traffic and a group of charts for active protocols. The Activity dashboard in the ExtraHop Command appliance contains a list of nodes ordered by device count. The active protocol charts in the Command appliance measure activity for the top seven nodes for each licensed protocol.

Essentials Dashboards

Essentials dashboards are created by ExtraHop staff to display common and related network metrics. A set of Essentials dashboards are available in the Essentials bundle on your ExtraHop appliance, but you can download additional dashboards from the ExtraHop Solution Bundle Gallery. For more information, see the Essentials Bundle section.

For more information, see the Navigating Dashboards section.

Create a dashboard

1. Click **New Dashboard** at the bottom of the left pane (dashboard dock) or from the command menu in the upper right corner of the page.
2. In the Dashboard Properties window, review the following:

   Title

   Type a name for the dashboard.

   Author

   Type your name.

   Description

   Type a brief description of the dashboard.

   Permalink

   (Optional) To change the five-character unique identifier in the permalink, click the link and type a meaningful name.

Note: The permalink name can have up to 100 characters combining letters, numbers, and the following symbols: ._-+)[]. The name cannot contain spaces.

Editors

Specifies the names of users that have editing access for the dashboard. The default editor is the author. Add editors to your dashboard by sharing your dashboard. For more information, see the Share a Dashboard section.

Theme

Select a radio button to specify a style for the dashboard.

3. Click **Create**.
   The new dashboard is populated with a region that contains an unconfigured chart and text box widget.
4. (Optional) To configure the first chart, click inside the Chart area.
   For more information, see the Configure a Chart section.
5. (Optional) To edit text in the first text box widget, click inside the Text Box area.
   For more information, see the Configure a Text Box Widget in Markdown section.
6. Click **Exit Layout Mode** when you are satisfied with your changes.

Configuring a dashboard

You can configure new and existing custom dashboards that you create, or have editing access to.

There are several ways to configure a dashboard. For more information, see the following related links:

- Change Dashboard Properties
- Edit a Dashboard Layout
- Add a Region
- Remove a Region
- Modify Sources in a Region
- Add a Widget
- Remove a Widget
- Copy a Widget
- Print a Widget
- Configure a Chart
- Configure a Text Box Widget in Markdown

Change dashboard properties

1. Click the command menu in the upper right corner of the page, and select **Dashboard Properties** to configure a new or existing custom dashboard.
2. In the Dashboard Properties window, review the following:

Title

Type a name for the dashboard.

Author

Type your name.

Description

Type a brief description of the dashboard.

Permalink

To change the five-character unique identifier in the permalink, click the link and type a meaningful name.

Note: The permalink name can have up to 100 characters combining letters, numbers, and the following symbols: ._-+)[]. The name cannot contain spaces.

Sharing

To share a dashboard with users who can view and edit, click the link. For more information, see the Share a Dashboard section.

Editors

Displays the names of users that you granted editing access to. To change the editors, click **Sharing**.

Theme

Select a radio button to specify a style for the dashboard.

3. Click **Save**.

Edit a dashboard layout

1. Click the command menu in the upper right corner of the screen, and select **Edit Layout** to add and arrange regions and widgets in a custom dashboard.

   For more information, see the Add a Region, Remove a Region, Add a Widget, and Remove a Widget sections.
2. After making changes, click **Exit Layout Mode**.

   Note: If an error message appears, another user might be making changes. It is best practice for each ExtraHop user to have an individual account.

Add a region

1. Click the command menu button in the upper right corner of the page and select **Edit Layout**.
2. From the bottom of the page, click and drag a region onto the dashboard.
3. Click **Region** in the upper left corner of the region and type a new name in the **Title** field.
4. Click **Save**.
5. Click the **Exit Layout Mode** button in the upper right corner of the dashboard to return to the Dashboards page.

Remove a region

1. Click the command menu in the upper right corner of the page and select **Edit Layout**.
2. Click the command menu in the upper right corner of the widget and select **Delete**.
3. Click **Delete**.
4. Click **Exit Layout Mode** in the upper right corner of the dashboard to return to the Dashboards page.

Modify sources in a region

1. Click the command menu in the upper right corner of the region and select **Modify Sources**.
2. In the Modify Sources window, select the object that you want to change from the list on the right and choose a new metric source.

   You can also change the title of the region by clicking on the region name to the right.
3. Click **Save Dashboard**.

Add a widget

1. Click the command menu button in the upper right corner of the page and select **Edit Layout**.
2. Drag-and-drop one of the following widget types onto the region.

Chart

This widget is user-defined. For information, see Configure a Chart section.

Alert History

This widget shows the alert history information about the objects in the list. Click **Add metric source** to customize the alert history.

Activity Groups

This widget shows a list of all activity during the specified time interval and cannot be configured.

Text Box

This widget provides a space for typing and displaying custom text in a dashboard region. You can format text with the Markdown syntax. For more information, see the Configuring a Text Box Widget in Markdown section.

Note: If you place a widget on top of another widget, it will appear red, indicating that widgets are overlapping and will not display properly when you click **Exit Layout Mode**. To create more space in the region for the new widget, expand the region size and then move the widget to a new location until it is no longer red.

3. Click **Save**.

Note: If an error message appears, another user might be making changes. It is best practice for each ExtraHop user to have an individual account.

4. Click **Exit Layout Mode** in the upper right corner of the dashboard to return to Dashboards.

Remove a widget

1. Click the command menu in the upper right corner of the page and select **Edit Layout**.
2. Click the command menu in the upper right corner of the widget and select **Delete**.
3. Click **Delete Widget**.
4. Click **Exit Layout Mode** in the upper right corner of the dashboard to return to the Dashboards page.

Copy a widget

To copy a widget to another dashboard:

1. Right-click any table, chart, or tile on the widget.
2. Select **Copy to...**.
3. Select the dashboard to place the widget or select **New Dashboard** to create a new dashboard.
   For more information, see the Create a Dashboard topic.

Print a widget

1. Right-click any table, chart, or tile on the widget and select Print.
   The print preview appears in a new window.
2. Click the **Theme** drop-down list and select a theme.
3. Click **Print Widget**.

Configure a chart

1. In your dashboard, click the command menu in the upper right corner of the page, and then select **Edit Layout**.
2. Click anywhere within the chart to open the Widget Configuration window.

Note: If you do not select **Edit Layout**, you can also open the Widget Configuration window by clicking the command menu in the chart header, and then clicking **Edit**.

3. Select a chart type from the Type section in the center pane.

   For more information about determining which chart to select, see the Chart Types section. The chart you selected will display in the Preview section of the center pane.
4. Click **Add metric source**.
5. Select a metric source by typing all or part of the name of the application, device, device group, or network.
   A filtered list of objects will appear, including suggestions listed by Recent Objects, Popular Objects, and Networks.
6. Click the object that you want to select as the metric source.
   The metric source is added to the chart in the Preview section.
7. Select a metric by typing all or part of a protocol or metric name and then selecting the metric from the filtered list of built-in custom metrics that appear.
   a) (Optional) Click **Drilldown** to search for detailed metrics by host, client IP, server IP, or by URI.
   b) (Optional) For frequency-related (dataset) metrics, click the drop-down list to display your data as the following statistic values:
      - Summary - 95th, 75th, 50th, 25th, and 5th percentiles
      - Percentile - three custom percentiles that you can specify
      - Minimum
      - Median - 50th percentile
      - Maximum
      - Mean
   c) (Optional) To remove a metric, click the **x** button in the upper left corner in the metric field. Or, to replace a metric, click the metric name to open a new search.
   d) (Optional) Click **Add Metric** to continue adding metrics to display on the chart.
   e) (Optional) Click **Add Metric Source** to continue adding metric sources to display on the chart.
   f) (Optional) Click the **Options** tab to change the units, show the metric as a rate, change the suffix notation, or change the labels.
8. Click **Save**.
9. Click **Exit Layout Mode**.

Chart types
Charts help you to quickly identify interesting data points or patterns in your network data. You can add and configure chart widgets in your own dashboards with the Metric Explorer.

The following table provides descriptions and compatible metrics for the different chart types that you can configure in the ExtraHop Web UI.

Chart type	Description	Compatible ExtraHop metrics
Area	Displays a metric as a line, which connects data points over time, with the area between the line and axis filled in with color. If your chart contains more than one metric (for example, **HTTP - Requests** and **HTTP - Responses**), each metric is stacked together to illustrate the cumulative value of all of the metrics. When you drill down by IP address on a metric, each IP address is stacked in the chart.	All metrics

Chart type	Description	Compatible ExtraHop metrics
	Click the legend to isolate individual metrics.	
Bar	Displays metrics as horizontal bars with the length representing the value.	All metrics
Candlestick	Displays the distribution of a metric over time in a box-and-whisker line. The box-and-whiskers are composed of a body (orange), a tick mark, and an upper and lower shadow. For dataset metrics, the percentiles are displayed. For sampleset metrics, the mean and standard deviation are displayed only. For dataset metrics: Select **Summary** to display the 95th, 75th, 50th, 25th, and 5th percentiles. The body represents the range from the 25th percentile to the 75th percentile. The tick mark represents the median (50th percentile). The upper shadow represents the 95th percentile. The lower shadow represents the 5th percentile. Select **Percentile** to type custom percentiles values. **Note:** You must type three percentiles, and each value must be separated by a comma. The upper shadow represents the top range for your selection. The tick mark is the middle value. The lower shadow is the bottom range of your selection.	Dataset and sampleset For example: • Dataset: HTTP server processing time • Sampleset: HTTP server processing time detail
Column	Displays metrics as vertical bars over time. If your chart contains more than one metric, percentiles, or drill-down IP addresses, the bars are stacked together to illustrate the cumulative value of the metrics. Click the legend to isolate individual metrics.	All metrics

Chart type	Description	Compatible ExtraHop metrics
Heatmap	Displays a matrix of frequency metrics over time, where the colors represent a percentile. For example, the 90th percentile is represented by a darker color on the heatmap, which indicates when a higher value occurred over time. The heatmap legend shows which color is mapped to specific percentiles: • Light: 0th percentile • Medium-Light: 50th percentile • Medium-Dark: 75th percentile • Dark: 90th percentile	Dataset For example: • HTTP server request transfer time • HTTP server processing time • HTTP server response transfer time
Line	Displays a metric in a line, which connects a series of data points over time. If your chart contains more than one metric (for example, **HTTP - Requests** and **HTTP - Responses**), percentiles, or drill-down IP addresses, a line for each metric will overlap to illustrate how the values compare over time. Click the legend to isolate individual metrics.	All metrics
Line & Column	Displays a metric in a line, which connects a series of data points over time. If your chart contains more than one metric (for example, **HTTP - Requests** and **HTTP - Responses**), you can select **Display as Columns** on one of the metrics to display those values as a column chart underneath the line chart. The default color for the column chart displays columns in red if there is an error associated with the metric. To change the color, click **Options** and click the **Display columns in red** checkbox. For more information see the Chart Options section.	All metrics
List	Displays metrics in a list.	All metrics

Chart type	Description	Compatible ExtraHop metrics
Pie	Displays metrics as a portion of a whole.	Count For example: • HTTP requests • DNS request timeouts
Single value	Displays the value for a single metric. You can select which value to display when you select your metric. Set up additional widgets side by side to create multiple tiles.	All metrics
Status	Displays metrics in a table with status information overlaid onto the data.	All metrics

Chart options
You can configure the following different chart options for each chart type.

Units

By default, metric units are displayed as bytes per a rate (such as per second) for count metrics. You can remove the rate conversion and convert bytes to bits. For some chart types, you can also select a logarithmic conversion for the y-axis or x-axis.

Suffix Notation

Select whether you would like the metric values to display as base 10 or base 2.

Legend and Labels

Select how to display information, such as metric source or IP address, in a label or legend. You also have the option to hide a legend.

Error Indication

(Line & Column only) Select whether to display errors associated with your selected metrics on the chart as red columns.

Date Format

Select the date and time format for your chart.

Layout

(Pie Chart only) Select how to display values in your chart. By default, the total value for all metrics displays in the middle of the pie chart. Also, when you select a drill-down IP address in your pie chart, each drill-down value is represented by a slice. By default, a slice that represents all other remaining values is included in the pie chart, illustrating the drill-down values metrics in proportion to the total value. You also can display percents instead of accounts in the pie chart.

Status Indication

By default, the color for the metric display indicates the alert status. The colors red, orange, and yellow mean that an alert has been triggered, which is associated the metric. Each color indicates the severity of the alert, which you can configure. For more information, see the Alerts section.

The following tables describes the available chart options for each chart type:

Chart type	Available Options
Area	• Units • Suffix Notation

Chart type	Available Options
	• Legend
Bar	• Units (log scale is unavailable) • Suffix Notation • Labels
Column	• Units • Suffix Notation • Legend
Candlestick	• Units • Suffix Notation
Heatmap	• Units (log scale is unavailable)
Line	• Units • Suffix Notation • Legend
Line & Column	• Units • Suffix Notation • Error Indication • Legend
List	• Units • Suffix Notation • Labels • Date Format
Pie	• Units • Suffix Notation • Layout • Labels • Legend
Single value	• Units • Suffix Notation • Status Indication • Labels • Date Format
Status	• Units • Suffix Notation

For more information, see the Configure a Chart and Chart Types sections.

Configure a text box widget in markdown

The Text Box widget enables you to type and display custom text in a dashboard region. It is a helpful tool for adding notes about a chart or data in a dashboard.

The text widget supports the Markdown syntax, which enables you to format text and add metric variables that display updated metric data dynamically. Markdown is a simple formatting syntax that converts plain text into HTML with non-alphabetic characters, such as “#” or “*”. A new Text Box widget contains example text that is already formatted in Markdown.

1. Open the Widget Configuration dialog.
   - Click anywhere within the text box widget when Edit Layout is selected from the dashboard command menu.
   - Click the text widget command menu and selecting Edit when Edit Layout is not selected.
2. In the Widget Configuration dialog, Edit text in the left Editor pane.
   The HTML output text dynamically displays in the right Preview pane.
3. Click **Save**.

Formatting text in Markdown syntax
The following table shows common Markdown formats that are supported in the Text Box widget.

Additional Markdown format examples are provided in the GitHub Guides: Mastering Markdown. However, not all Markdown syntax formatting options are supported in the text widget.

Note: Adding emojis in Markdown syntax is unsupported. However, copying and pasting a Unicode block emoji is supported in the text widget. For more information, see Unicode Emoji Chart website.

Format	Description	Example
Headings	Place a number sign (#) before your text to format headings. The level of heading is determined by the amount of number signs.	####Example H4 heading
Unordered lists	Place a single asterisk (*) before your text to format bulleted lists.	* Example 1 * Example 2
Ordered lists	Place a single number and period (1.) before your text to format numbered lists.	1. Example 1 2. Example 2
Bold	Place double asterisks before and after your text to format bold.	**bold text**
Italics	Place an underscore before and after your text to format italics.	_italicized text_
Hyperlinks	Place link text in brackets before the URL in parentheses. Or type your URL.	[Visit our home page](http://www.extrahop.com) http://www.extrahop.com
Blockquotes	Place a right angle bracket and a space before your text to format a blockquote.	On the ExtraHop website: > Access the live demo and review case studies.

Adding images in markdown syntax
You can add images by linking to them. Make sure your image is hosted on a network that is accessible to the Discover appliance.

Links to images must be specified in the following format:

```
![<alt_text>](<file_path>)
```

Where <alt_text> is the alternative text and <file_path> is the path of the image. For example:

```
![Graph](/images/graph_1.jpg)
```

Note: You also can add images by encoding them to Base64. For more information, see the following post on the ExtraHop Customer forum, "Putting Images in Text Boxes."

Adding metrics in markdown syntax
You can add metric variables to a text widget by writing metric queries in Markdown.

The Markdown format for writing metric queries is:

```
%%metric:<definition>%%
```

Where *<definition>* is replaced with a JSON-defined structure that is based on the ExtraHop REST API query structure.

Note: The following metric queries are unsupported in the Text Box widget:

- Time-series queries
- Mean calculations
- Multiple object_ids
- Multiple metric_spec
- Multiple percentiles

A metric query must contain the following parameters:

- `object_type`
- `object_ids`
- `metric_category`
- `metric_spec`

To retrieve the object_type, metric_spec, and metric_category values for a metric name:

1. Click Settings
2. Click Metric Catalog
3. Type the metric name in the search field
4. Select the metric, and look for the values in the REST API Parameters section.

For more information, see the Metric Catalog section.

You can retrieve object_ids from the URL that you are browsing.

Object Type	URL Parameter
Application	applicationOID=
Network	networkOID=
Group	deviceGroupOID=
Device	deviceOID=

Metric variables examples
The following examples show you how to write top-level metric queries for application, device, and network objects, and detail metric queries.

Application queries

To specify the All Activity object, the `object_ids` is "`0`".

This example query shows how you can retrieve HTTP metrics from the All Activity object, and displays the following output: "`Getting [value] HTTP requests and [value] HTTP responses from All Activity.`"

```
Getting
%%metric:{
"object_type": "application",
"object_ids": [0],
"metric_category": "http",
"metric_specs": [{"name":"req"}]
}%%HTTP requests and
%%metric:{
"object_type": "application",
"object_ids": [0],
"metric_category": "http",
"metric_specs": [{"name":"rsp"}]
}%%
HTTP responses from All Activity.
```

Device queries

You must specify either a client ("`_client`") or server ("`_server`") in the `metric_category`. To retrieve metrics for a specific device, specify the device object ID number in `object_ids`. To retrieve the device object ID (`deviceOid`), search for the device object in the ExtraHop global search. Select the device from your search results. The "`deviceOid=`" value will be embedded in the URL query string.

This example query shows how to retrieve metrics from a device client object, and displays the following output: "`Getting [value] CLIENT DNS response errors from a specific device.`"

```
Getting
%%metric:{"object_type": "device",
"object_ids": [8],
"metric_category": "dns_client",
"metric_specs": [{"name":"rsp_error"}]
}%%
CLIENT DNS response errors from a specific device.
```

This example query shows how to retrieve metrics from a device server object, and displays the following output: "`Getting [value] SERVER DNS response errors from a specific device.`"

```
Getting
%%metric:{
"object_type": "device",
"object_ids": [156],
"metric_category": "dns_server",
"metric_specs": [{"name":"rsp_error"}]
}%%
SERVER DNS response errors from a specific device.
```

Network queries

To specify All Networks, the `object_type` is "`capture`" and the `object_ids` is "`0`." To specify a specific VLAN, the `object_type` is "`vlan`" and the `object_ids` is the VLAN number.

This example query shows how to retrieve metrics for all networks, and displays the following output: "`Getting [value] broadcast packets from all networks.`"

```
Getting
%%metric:{
"object_type": "capture",
```

```
"object_ids": [0],
"metric_category": "net","metric_specs":
 [{"name":"frame_cast_broadcast_pkts"}]
}%%
broadcast packets from all networks.
```

This example query shows how to retrieve metrics for a specific VLAN and displays the following output: "`Getting [value] broadcast packets from VLAN 3.`"

```
Getting
%%metric:{
"object_type": "vlan",
"object_ids": [3],
"metric_category": "net",
"metric_specs": [{"name":"frame_cast_broadcast_pkts"}]
}%%
broadcast packets from VLAN 3.
```

Group queries

To specify a group, the `object_type` is "`activity_group`" or "`custom_group`." You must specify either a client ("`_client`") or server ("`_server`") in the `metric_category`. The `object_ids` for the specific group must be retrieved from the REST API Explorer.

This example query shows how to retrieve metrics for all networks, and displays the following output: "`Getting [value] HTTP responses from the HTTP Client Activity Group.`"

```
Getting
%%metric:{
"object_type": "activity_group",
"object_ids": [17],
"metric_category": "http_client",
"metric_specs": [{"name":"req"}]
}%%
HTTP responses from the HTTP Client Activity Group.
```

Detail metric queries

If you want to retrieve detail metrics, your metric query should contain additional key parameters, such as key1 and key2:

- object_type
- object_ids
- metric_category
- metric_spec
  - name
  - key1
  - key2

Note: The key parameters act as a filter for displaying detail metric results.

For built-in detail metrics, you can retrieve detail metric parameters from the Metric Catalog. For example, type HTTP Responses by URI, and then look at the parameter values in the REST API Parameters section. You must supply the `object_ids`.

This example shows how to retrieve HTTP requests by URI for the All Activity application (`object_ids` is "`0`"):

```
%%metric:{
```

```
"object_type": "application",
"object_ids": [0],
"metric_category": "http_uri_detail",
"metric_specs": [{"name":"req"}]
}%%
```

This example query shows you how to retrieve HTTP requests by URIs that contain a key value for "pagead2" for the All Activity application (object_ids is "0"):

```
%%metric:{
"metric_category": "http_uri_detail",
"object_type": "application",
"object_ids": [0],
"metric_specs": [
{
"name": "req",
"key1": "/pagead2/"
}
]
}%%
```

This example query shows how to retrieve count metrics for all networks and displays the following output: "Getting [value] detail ICA metrics on all networks."

```
Getting
%%metric:{
"object_type": "capture",
"object_ids": [0],
"metric_category": "custom_detail",
"metric_specs": [{
"name":"custom_count",
"key1":"network-app-byte-detail-ICA"
}],
}%%
detail ICA metrics on all networks.
```

This example query shows how to retrieve a custom dataset statistic with topn keys and percentiles, and displays the following output: "The fifth percentile is: [value]."

```
The fifth percentile is:
%%metric:{
"object_type": "vlan",
"object_ids": [1],
"metric_category": "custom_detail",
"metric_specs": [{
"name": "custom_dset",
"key1": "myCustomDatasetDetail",
"key2": "/10.10.7/",
"calc_type": "percentiles",
"percentiles": [5]
}],
}%%
.
```

Note: Sampleset metrics are unsupported in the text box widget. For example, adding the "calc_type": "mean" parameter to your text box query is unsupported.

Share a dashboard

You can share custom dashboards with other ExtraHop users and decide whether to give them view or edit access.

1. Click **Dashboards**.
2. In the left pane, under My Dashboards, click the name of a dashboard.
3. Click the command menu button in the Navigation bar and select **Share**.
4. Specify which users can view the dashboard.
   - To allow specific users to view the dashboard, click **Only specified users can view or edit**. In the Specify users: area, type the name of a user and select the author name from the dropdown list. Select **Can view** and click the green plus icon. Repeat this process for additional users.
   - To allow all users to view the dashboard, click **All users can view; only specified users can edit**.
5. Specify which users can edit the dashboard:
   a) In the Specify users: area, type the name of a user.
   b) Select the author name from the dropdown list.
   c) Select **Can edit**.
   d) Click the green plus icon.
   e) Repeat the process for additional users.

   Note: Users that can edit a dashboard can also view the dashboard.
6. Click **Save**.

Remove view or edit access to a dashboard

1. Click **Dashboards**.
2. In the left pane, under My Dashboards, click the name of a dashboard.
3. Click the command menu button in the Navigation bar and select **Share**.
4. Click the red delete icon next to the name of the user you would like to remove access for.
5. Click **Save**.

View a dashboard

You can opt to show hover-over descriptions for available protocols in dashboards. You can also select between two options to view dashboards: in presentation mode or as a widget slideshow. Statuses and alerts can be viewed in dashboards through widgets.

Click **Dashboards**.

- To show or hide descriptions, click the command menu button in the upper right corner of the page, and select **Show Descriptions** to display highlighting for available descriptions. Hover your mouse over the highlighted box to display the description.

  You can also view descriptions in charts that display traffic from individual ports. Descriptions are provided for protocols that the Discover appliance parses.
- To view a dashboard in presentation mode, click the command menu button in the upper right corner of the page, and select **Presentation Mode** to enter a full-screen display of the metrics on the currently selected dashboard. Click the **Exit Presentation Mode** button to return to the previous display.

Tip: You can open a dashboard in presentation mode directly by appending `/presentation` to the URL. For example:

```
https://<extrahop_ip>/extrahop/#/Dashboard/437/presentation
```

- To view a dashboard as a widget slideshow, click the command menu button in the upper right corner of the page, select **Widget Slideshow**, and select a time increment to view a slideshow of widgets within the current region. Click the **X** in the upper right corner of the screen to return to the previous display.

Status widgets, or service availability widgets, are based on alerts. The service status is displayed in a bar graph with red, orange, yellow, or green bars based on the severity and type of configured alerts.

A user-defined, detailed alert can be associated with a device, group, or application widget. When you configure an alert for a specific metric, any alert of that metric type will appear on the widget. For example, if you configure an alert to fire on HTTP responses, HTTP errors, and HTTP response times, all three metrics appear on the widget. When an alert fires, the bar on the widget associated with that alert is colored based on the severity level set in the alert. If multiple alerts fire on the same widget, the color of the bar reflects the most severe alert.

To display a list of all alerts related to the widget, click **Show Related Alerts**.

Organize dashboards

1. Create folders for dashboards.
   a) At the bottom of the left pane, click the configuration button to the right of New Dashboard.
   b) Select **New Folder**.
   c) Type a name for the folder and click **Save**.
2. Add dashboards to a folder.
   a) At the bottom of the left pane, click the command menu to the right of **New Dashboard**.
   b) Click **Edit Dock**.

      In Edit mode, you can organize, edit, and create new dashboards.
   c) Drag-and-drop any of the dashboards you created into a folder.

      If dashboards are sorted in ascending or descending order, the drag-and-drop functionality is disabled. To enable this functionality again, click the sort icon in the upper right header of the dashboard dock until the custom sort icon displays.
   d) Click the right-most button in the panel to save and exit edit mode.

Export dashboard data

You can export data from any chart or table in your dashboard to a CSV or Excel file. You cannot export content from a text box widget.

1. Right-click the chart or table that you want to export.
2. Select **Export to CSV** or **Export to Excel**.

Your file will be downloaded to your local computer.

Delete a dashboard

1. Click the command menu in the upper right corner of the page, and select **Delete**.
2. Click **Delete Dashboard** in the Are you sure? dialog box.

Protocols

You can view metrics about protocols through the following views:

Devices

Displays device metrics to troubleshoot network issues at the device level.

Groups

Displays a select group of devices, filtering out the devices that are not likely to be related to the traffic being examined. Group metrics are aggregated and viewable by all ExtraHop users on the network.

Applications

Displays applications. Some applications use multiple devices, and some devices host multiple applications. The Discover appliance provides a set of default applications based on all traffic. You can modify the default application template to suit the needs of your organization, and you can add your own applications.

Networks

Displays network capture attributes, network alerts, and network traffic details. Networks is the entry point into the network capture. The metrics that are collected provide a summary of all network activity retrieved in the capture.

AAA

ExtraHop appliances collect metrics about Authentication, Authorization, and Accounting (AAA) activity.

AAA applications page

Application toolbar

The AAA application toolbar includes the following controls:

Errors

The chart shows the number of AAA errors. Mouse over the points to view a summary of a specific time or date. The table lists the AAA error messages and number of occurrences.

Clients

The chart shows processing time for all clients. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client as well as total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows processing time for all servers. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the `Request Bytes` counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with AAA requests.

Response L2 Bytes

The number of L2 bytes associated with AAA responses.

Request Packets

The number of packets associated with AAA requests.

Response Packets

The number of packets associated with AAA responses.

Request RTOs

The number of retransmission timeouts caused by congestion when clients were sending AAA requests. A retransmission timeout is a one-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

The number of retransmission timeouts caused by congestion when servers were sending AAA responses. A retransmission timeout is a one-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

The number of zero window advertisements sent by AAA clients. A device advertises a zero window when it cannot process incoming data as quickly as it is arriving.

Response Zero Window

The number of zero window advertisements sent by servers while receiving AAA requests. A device advertises a zero window when it cannot process incoming data as quickly as it is arriving.

AAA Metrics

Contains the following metrics:

Requests

The number of AAA requests.

Responses

The number of AAA responses.

Errors

The number of AAA errors for the selected time interval.

Aborts

The number of aborted AAA sessions.

RADIUS Requests

The number of RADIUS requests.

Diameter Requests

The number of Diameter requests.

Methods

Displays the selected method types for the AAA client or server.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

AAA devices page

AAA Device Toolbar

The AAA device toolbar includes the following controls:

AAA Metric Type

Display metrics for devices acting as an AAA client or AAA server.

Errors

Click the **Errors** button to display the list of error messages sent to or received by the current device over the time interval. Errors are formatted as follows: Results-Code-Description:Session-Id:Error-Reporting-Host:Subscription-ID-Data.

- Session-Id frequently contains multiple semicolon-separated records.
- Error-Reporting-Host is not always present.

Records

Displays results for records that match the selected metric source and protocol.

AAA Client

If you select **Client** for the AAA Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of total requests that the device sent when acting as an AAA client.

Responses

Number of responses that the device received when acting as an AAA client.

Errors

Number of AAA errors for the selected time interval.

Aborts

Number of aborted sessions that occurred when the device is acting as an AAA client.

AAA Server

If you select **Server** for the AAA Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of total requests that the device received when acting as an AAA server.

Responses

Number of responses that the device sent when acting as an AAA server.

Errors

Number of AAA errors for the selected time interval.

Aborts

Number of aborted sessions that occurred when the device is acting as an AAA server.

Messages

Selected message types for the AAA server.

Status Codes

The AAA status codes for the selected time interval.

Processing Time Distribution

Displays a histogram of times it took the server to process requests. Move the mouse pointer over each bar to display the time range it represents and the number of requests in this bin.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

AAA groups page

AAA Groups Toolbar

The AAA groups toolbar includes the following controls:

Metric Type

Click the **Metric Type** drop-down list and select either **Client** or **Server** to display metrics for devices in the current group acting as an AAA client or AAA server, respectively.

Errors

Click the **Errors** button to display the list of error messages sent to or received by the current member over the time interval. Errors are formatted as follows: Results-Code-Description:Session-Id:Error-Reporting-Host:Subscription-ID-Data.

Session-Id frequently contains multiple semicolon-separated records. Error-Reporting-Host is not always present.

Records

Displays results for records that match the selected metric source and protocol.

AAA Client

If you select **Client** for the AAA Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of AAA requests for the selected time interval.

Responses

Number of AAA responses for the selected time interval.

Errors

Number of AAA errors for the selected time interval.

Aborts

Number of AAA aborted requests for the selected time interval.

Diameter Requests

Number of Diameter requests for the selected time interval.

Radius Requests

Number of RADIUS requests for the selected time interval.

AAA Server

If you select **Server** for the AAA Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of AAA requests for the selected time interval.

Responses

Number of AAA responses for the selected time interval.

Errors

Number of AAA errors for the selected time interval.

Aborts

Number of AAA aborted requests for the selected time interval.

Diameter Requests

Number of Diameter requests for the selected time interval.

Radius Requests

Number of RADIUS requests for the selected time interval.

Messages

Selected message types for the AAA server.

CIFS

ExtraHop appliances collect metrics about Common Internet File System (CIFS) activity.

CIFS devices page

CIFS Devices Toolbar

The CIFS device page toolbar includes the following controls:

CIFS Metric Type

Displays metrics for the current device acting as a CIFS client or CIFS server.

Errors

Displays the list of error messages sent to or received by the current device over the selected time interval.

Warnings

Displays the list of warning messages sent to or received by the current device over the selected time interval.

Methods

Displays the list of methods and associated bytes sent and received by the current device for the selected time interval. Methods are broken out by key parameters, such as the accessed file name and file access time.

Users

Displays the list of users accessing the file server and associated bytes sent and received for the selected time interval.

Files

Displays the list of files accessed and associated bytes sent and received for the selected time interval. The access time indicates the time to access a file on a CIFS partition and is measured by timing the first READ or WRITE on every flow.

Records

Displays results for records that match the selected metric source and protocol.

Where file name detail is presented, the Discover appliance displays both the file path and mount point, if available. The prefix '...' indicates that either the mount point or part of the path is not available. This may occur in instances when the capture process was restarted after the "mount" or a "cd" command was issued, or when the commands were lost due to desyncs.

Click the counters next to individual CIFS metrics to show the IP Address CIFS Metrics details for CIFS peer devices. For CIFS servers, the peer devices are CIFS clients. For CIFS clients, the peer devices are CIFS servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

CIFS Server

Displays additional IP address details.

Responses

Specifies the number of responses that the device sent when acting as a CIFS server.

Errors

Specifies the number of errors sent by the CIFS server.

Warnings

Displays the list of warning messages sent to or received by the CIFS server over the selected time interval.

Reads

Specifies the number of read operation requests that the device received when acting as a CIFS server.

Writes

Specifies the number of write operation requests that the device received when acting as a CIFS server.

Locks

Specifies the number of lock operation requests that the device received when acting as a CIFS server.

FSInfo

Specifies the number of file system metadata queries that the device received when acting as a CIFS server.

CIFS Client

Displays additional IP address details.

Responses

Specifies the number of responses that the device received when acting as a CIFS client.

Errors

Specifies the number of errors sent by the CIFS client.

CIFS groups page

CIFS Groups Toolbar

The CIFS groups toolbar includes the following controls:

CIFS Metric Type

Displays metrics for devices in the current group acting as a CIFS client or server, respectively.

Errors

Displays the list of error messages sent to or received by devices in the current group over the selected time interval.

Warnings

Displays the list of warning messages sent to or received by devices in the current group over the selected time interval.

Methods

Displays the list of methods and associated bytes sent and received by devices in the current group during the selected time interval. Methods are broken out by key parameters, such as the accessed file name.

Users

Displays the list of users accessing the file server and associated bytes sent and received for the selected time interval.

Files

Displays the list of files accessed and associated bytes sent and received for the selected time interval. Access Time indicates the time it took for the server to access a file on disk.

Records

Displays results for records that match the selected metric source and protocol.

Where file name detail is presented, the Discover appliance displays both the file path and mount point, if available. The prefix '...' indicates that either the mount point or part of the path is not available. This may occur in instances when the capture process was restarted after the "mount" or a "cd" command was issued, or when the commands were lost due to desyncs.

CIFS Server

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Responses

Specifies the number of responses sent by the CIFS server.

Errors

Specifies the number of errors sent by the CIFS server.

Warnings

Displays the list of warning messages sent to or received by devices in the CIFS server over the selected time interval.

Reads

Specifies the number of read operations requested from the CIFS server.

Writes

Specifies the number of write operations requested from the CIFS server.

Locks

Specifies the number of lock operations requested from the CIFS server.

CIFS Client

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Responses

Specifies the number of responses received by the CIFS client.

Errors

Specifies the number of errors sent by the CIFS client.

Warnings

Displays the list of warning messages sent to or received by the CIFS client over the selected time interval.

Reads

Specifies the number of read operations requested by the CIFS client.

Writes

Specifies the number of write operations requested by the CIFS client.

Locks

Specifies the number of lock operations requested by the CIFS client.

Methods

Displays the CIFS methods for the selected time interval.

Click the counter next to the method to break it down by group members in the table.

Database

ExtraHop appliances collect metrics about database activity.

Database applications page

Database Application Toolbar

The Database application toolbar includes the following controls:

Errors

The chart shows the total count for DB errors. Mouse over the points to view a summary of a specific time or date. The table lists DB error messages and the number of occurrences.

Methods

The chart shows responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists methods, number of responses, total time, and processing time (ms) associated with each method. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Users

The chart shows the number of responses and errors for all users. Mouse over the chart to view a summary of a specific time or date. The table displays the list of users, and the number of responses and errors associated with each user.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the `Request Bytes` counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with database requests.

Response L2 Bytes

The number of L2 bytes associated with database responses.

Request Packets

The number of packets associated with database requests.

Response Packets

The number of packets associated with database responses.

Request RTOs

The number of retransmission timeouts caused by congestion when clients were sending database requests. A retransmission timeout is a one-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

The number of retransmission timeouts caused by congestion when servers were sending database responses. A retransmission timeout is a one-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

The number of zero window advertisements sent by database clients. A device advertises a zero window when it cannot process incoming data as quickly as it is arriving.

Response Zero Window

The number of zero window advertisements sent by servers while receiving database requests. A device advertises a zero window when it cannot process incoming data as quickly as it is arriving.

DB Metrics

Contains the following metrics:

Requests

The number of database requests.

Responses

The number of database responses.

Response Errors

The number of database response errors.

Transaction Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

Database devices page

Database Devices Toolbar

The Database device toolbar includes the following controls:

Database Metric Type

Displays statistics for the current device acting as a database client or database server.

Errors

Displays the list of error messages sent to or received by the current device over the time interval.

Methods

Displays the list of names and the associated number of responses and errors.

Users

Displays the list of users accessing the database server and associated bytes sent and received for the selected time interval.

Clients or Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

Click the counters next to individual database metrics to show the IP Address Database Metrics for database peer devices. For database servers, the peer devices are database clients. For database clients, the peer devices are database servers.

Device Details

Click the counters next to individual database metrics to show the IP Address Database Metrics for database peer devices. For database servers, the peer devices are database clients. For database clients, the peer devices are database servers.

By IP

Displays database metrics by IP address.

By Database

Displays database metrics by database. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

Database Client

If you select Client for the Database Metric Type, the Discover appliance displays the following metrics:

Responses

Specifies the number of responses that the device received when acting as a database client. Click to display the list of servers from which responses were sent.

Errors

Specifies the number of database protocol errors for the selected time interval. Click to display the list of servers for which there were errors.

Requests Aborted

Specifies the number of requests that the device began to send but did not send completely when acting as a database client.

Responses Aborted

Specifies the number of responses that the device began to receive but did not receive completely when acting as a database client.

Database Server

If you select Server for the Database Metric Type, the Discover appliance displays the following metrics:

Responses

Specifies the number of responses that the device sent when acting as a database server. Click to display the list of clients to which responses were sent.

Errors

Specifies the number of database protocol errors for the selected time interval. Click to display the list of clients for which there were errors.

Requests Aborted

Specifies the number of requests that the device began to receive but did not receive completely when acting as a database server.

Responses Aborted

Specifies the number of responses that the device began to send but did not send completely when acting as a database server.

Methods

Displays the database methods for the selected time interval. Methods will vary for each specific device.

Transaction Metrics

Displays the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th

and 75th percentile values. Move the mouse pointer over each component to display a five-number statistical summary.

ReqXfer

The request transfer time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

The server processing time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

The response transfer time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

Request Size

Displays the range of request sizes for all transactions associated with the current device. Mouse over the chart to see the five-number summary. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device, database, or IP address.

Response Size

Displays the range of response sizes for all transactions associated with the current device. Mouse over the chart to see the five-number summary. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean response size for each peer device, database, or IP address.

Transactions Per Second

Displays the number of database protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red data points to list the peer devices associated with the errors at this point in time. Click and drag across the chart to select a particular region. Select a database from the **Databases** drop-down list and then click the red data points to display results associated with that database only. For detailed error information, click **Errors**.

Response Time Breakdown

Displays the area chart containing median request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Database devices timing page

The timing charts draw data from the **Time Selector** drop-down list on the navigation toolbar. The events observed during this interval are used to fill the bins of a histogram that displays a distribution of timing data. Timing charts use a logarithmic horizontal axis that simultaneously displays events that took milliseconds and those that took seconds.

Request Transfer Time

Displays a histogram of times it took to transfer requests from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Processing Time

Displays a histogram of times it took the server to process requests. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Response Transfer Time

Displays a histogram of times it took to transfer the response from the server to the client. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Database groups page

Database Groups Toolbar

The Database groups toolbar includes the following controls:

Database Metric Type

Displays metrics for members in the current group acting as a database client or server, respectively.

Errors

Displays the list of error messages sent to or received by members in the current group over the time interval.

Methods

Displays the list of names and the associated processing times for the stored procedures executed within the databases belonging to the current group during the selected time interval.

Users

Displays the list of users accessing the database servers in this group and associated bytes sent and received for the selected time interval.

Database Client

If you select Client for the Database Metric Type, the Discover appliance displays the following metrics. Click the counter to break down the responses by group members in the table at the bottom of the page.

Responses

Specifies the number of database protocol responses received by all members of the current group during the selected time interval.

Errors

Specifies the number of database protocol errors received by all members of the current group during the selected time interval.

Requests Aborted

Specifies the number of requests that members of the group began to send but did not send completely when acting as a database client.

Responses Aborted

Specifies the number of responses that members of the group began to receive but did not receive completely when acting as a database client.

Database Server

If you select Server for the Database Metric Type, the Discover appliance displays the following metrics. Click the counter to break down the responses by group members in the table at the bottom of the page.

Responses

Specifies the number of database protocol responses sent by all members of the current group during the selected time interval.

Errors

Specifies the number of database protocol errors sent by all members of the current group during the selected time interval.

Requests Aborted

Specifies the number of requests that members of the group began to receive but did not receive completely when acting as a database server.

Responses Aborted

Specifies the number of responses that members of the group began to send but did not send completely when acting as a database server.

Methods

Displays the database methods for the selected time interval.

Database groups all methods page

Database Groups Toolbar

The Database groups toolbar includes the following controls:

Database Metric Type

Displays metrics for members in the current group acting as a database client or server, respectively.

Records

Displays results for records that match the selected metric source and protocol.

Methods

This section displays the database methods for the selected time interval. Click to display additional per-client or per-server details.

Database Client

This table lists the peer members associated with the database client.

Database Server

This table lists the peer members associated with the database server.

Database groups processing time page

Database Groups Toolbar

The Database groups toolbar includes the following controls:

Database Metric Type

Displays metrics for members in the current group acting as a database client or server, respectively.

Records

Displays results for records that match the selected metric source and protocol.

Server Processing Time

Shows median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

DHCP

ExtraHop appliances collect metrics about Dynamic Host Configuration Protocol (DHCP) activity.

DHCP applications page

DHCP Applications Toolbar

The DHCP applications toolbar includes the following controls:

Errors

Displays a chart that shows the number of DHCP errors.

Clients

Displays chart and table information about DHCP client activity. The chart shows the total number of client responses compared to processing time.

The table lists client IP addresses, the host and device associated with each client, the number of requests by each client, and total processing time.

Servers

Displays chart and table information about DHCP server activity. The chart shows the total number of server responses compared to processing time. The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with DHCP requests.

Response L2 Bytes

The number of L2 bytes associated with DHCP responses.

Request Packets

The number of packets associated with DHCP requests.

Response Packets

The number of packets associated with DHCP responses.

DHCP Metrics

Contains the following metrics:

Requests

The number of DHCP requests.

Responses

The number of DHCP responses.

Errors

Displays the number of DHCP errors.

Requests by Message Type

Displays the number of DHCP requests broken out by the message type.

Responses by Message Type

Displays the number of DHCP requests broken out by the message type.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Processing Time

Displays the median processing time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing time metrics.

DHCP devices page

DHCP Devices Toolbar

The DHCP devices toolbar includes the following controls:

DHCP Metric Type

From the drop-down menu, select the type of metrics for the current device.

Errors

Displays the list of error messages sent or received by the current device over the selected time interval.

Clients or Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

Records

Displays results for records that match the selected metric source and protocol.

DHCP Server

Requests

Displays the number of requests that the device received.

Responses

Displays the number of responses that the device sent.

Response Errors

Displays the number of response errors.

Requests by Message Type

Displays the number of requests that the device received for the message type.

Responses by Message Type

Displays the number of requests that the device received for the message type.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Server Processing Time

Displays the median server processing time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing time metrics.

Processing Time Distribution

Displays a histogram of the server time taken to process requests.

Requests by Record Type

Displays the categorization of all request types sent or received by the current device.

Responses by Record Type

Displays the categorization of all response types sent or received by the current device.

DHCP groups page

DHCP Groups Toolbar

The DHCP groups toolbar includes the following controls:

DHCP Metric Type

Displays metrics for members in the current group acting as a DHCP client or server, respectively.

Errors

Displays the list of error messages sent to or received by members in the current group over the time interval.

DHCP Client

If you select Client for the DHCP Metric Type, the Discover appliance displays the following metrics:

Requests

Specifies the number of requests that the device sent when acting as a DHCP client. Click the counter to display the list of servers to which requests were sent.

Responses

Specifies the number of responses that the device received when acting as a DHCP client. Click the counter to display the list of servers from which the responses were received.

Response Errors

Specifies the number of response errors for the selected time interval when acting as a DHCP client. Click the counter to display the list of servers associated with the errors.

DHCP Server

If you select Server for the DHCP Metric Type, the Discover appliance displays the following metrics:

Requests

Specifies the number of requests that the device sent when acting as a DHCP server. Click the counter to display the list of clients from which requests were received.

Responses

Specifies the number of responses that the device received when acting as a DHCP server. Click the counter to display the list of clients to which the responses were sent.

Response Errors

Specifies the number of response errors for the selected time interval when acting as a DHCP server. Click the counter to display the list of clients associated with the errors.

Requests by Message Type

Displays the number of DHCP requests broken out by the message type.

Responses by Message Type

Displays the number of DHCP requests broken out by the message type.

DHCP groups processing time page

Server Processing Time

Shows median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

DNS

ExtraHop appliances collect metrics about Domain Name System (DNS) activity.

DNS applications page

DNS Applications Toolbar

The DNS applications toolbar includes the following controls:

Errors

Displays a chart that shows the number of errors.

Host Queries

The chart shows the total number of host queries compared to processing time during the selected time interval. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists DNS hosts, number of host queries, and the processing time.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word `via`. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Host Query

Displays application metrics by host query.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the **Request Bytes** counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

DNS Metrics

Contains the following metrics:

Requests

The number of requests received.

Request Timeouts

The number of request timeouts. A request timeout occurs when there is a repeated request without a response to the first request.

Truncated Requests

The number of requests that were sent but were truncated in transit. A truncated request is indicated by the truncated bit in the message and occurs when the message is larger than the underlying transmission channel allows.

Responses

The number of responses received.

Response Errors

The number of response errors.

Truncated Responses

The number of responses that were sent but were truncated in transit. A truncated request is indicated by the truncated bit in the message and occurs when the message is larger than the underlying transmission channel allows.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Requests by Opcode

Displays all request opcode types sent or received by the current application. For each field, click to display the devices to or from which these requests were sent or received.

Query

Number of DNS QUERY Opcodes sent or received by the current application. DNS Queries are the most-frequently encountered DNS Opcode type.

Responses by Response Code

Displays all response codes broken down by request opcode and request record type sent (if server) or received (if client) by the current device. The format of the entry is `ERROR/REQUEST_OPCODE:REQUEST_RECORD`. For each field, click to display the devices to or from which these requests were sent or received.

The response code bar categories include the following:

NOERROR

Successful transaction; no error.

FORMERROR

Format Error.

SERVFAIL

DNS Server Failed.

NXDOMAIN

No such domain.

NOTIMPL

No handler implemented for this query type.

REFUSED

Query administratively refused.

UPDATEERR

Error in handling UPDATE request.

TSIGERR

Error in handling TSIG request.

OTHER

All other response code types.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Processing Time

Displays the mean processing time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing time metrics. Click and drag across the chart to select a particular region.

Click the graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

DNS devices page

DNS Devices Toolbar

The DNS devices toolbar includes the following controls:

DNS Metric Type

Displays metrics for the current device acting as a DNS client or DNS server.

Errors

Displays the list of DNS queries made to or from this device, sorted by **Host Query** frequency. Click the Query Errors header to sort the list by the number of DNS errors encountered.

Servers

When acting as a DNS client, displays a chart showing the total number of responses compared to processing time during the selected time interval.

Clients

When acting as a DNS client, displays a chart showing the total number of responses compared to processing time during the selected time interval.

Records

Displays results for records that match the selected metric source and protocol.

DNS Client

If you select **Client** for the DNS Metric Type, the Discover appliance displays the following metrics. For each field, click to display the devices to which these requests were made.

Requests

Specifies the number of requests that the device sent when acting as a DNS client.

Request Timeouts

Specifies the number of request timeouts when the device is acting as a DNS client. A request timeout occurs when there is a repeated request without a response to the first request. A high number here may indicate server unresponsiveness or a client misconfiguration.

Truncated Requests

Specifies the number of requests that were sent, but were truncated in transit, when the device is acting as a DNS client. A truncated request is indicated by the truncated bit in the message and occurs when the message is larger than the underlying transmission channel allows.

Responses

Specifies the number of responses that the device received when acting as a DNS client.

Response Errors

Specifies the number of responses received with a code other than NOERROR, when the device is acting as a DNS client.

Truncated Responses

Specifies the number of truncated responses that the device received when acting as a DNS client. A truncated response is indicated by the truncated bit in the message and occurs when the message is larger than the underlying transmission channel allows.

DNS Server

If you select **Server** for the DNS Metric Type, the Discover appliance displays the following metrics. For each field, click to display the devices from which these requests were received.

Requests

Specifies the number of requests that the device received when acting as a DNS server.

Request Timeouts

Specifies the number of request timeouts when the device is acting as a DNS server. A request timeout occurs when there is a repeated request without a response to the first request. A high number here might indicate a problem with this DNS server.

Truncated Requests

Specifies the number of requests that were received, but were truncated in transit, when the device is acting as a DNS server. A truncated request is indicated by the truncated bit in the message and occurs when the message is larger than the underlying transmission channel allows.

Responses

Specifies the number of responses that the device sent when acting as a DNS server.

Response Errors

When the device is acting as a DNS server, specifies the number of responses sent with a code other than NOERROR.

Truncated Responses

Specifies the number of responses sent, but later truncated, when the device is acting as a DNS server. A truncated response is indicated by the truncated bit in the message and occurs when the message is larger than the underlying transmission channel allows.

Requests by Opcode

Displays all request opcode types sent or received by the current device. For each field, click to display the devices to or from which these requests were sent or received.

Query

Specifies the number of DNS QUERY Opcodes sent or received by the current device. DNS Queries are the most-frequently encountered DNS Opcode type.

Notify

Specifies the number of DNS NOTIFY Opcodes sent or received by the current device. DNS Notify is used as a synchronization method between DNS servers.

Update

Specifies the number of DNS UPDATE Opcodes sent or received by the current device. DNS Update is used as a synchronization method between DNS servers.

Other

Specifies the number of other miscellaneous DNS Opcodes sent or received by the current device.

Responses by Response Code

Displays all response codes broken down by request opcode and request record type sent (if server) or received (if client) by the current device. The format of the entry is `ERROR/REQUEST_OPCODE:REQUEST_RECORD`. For each field, click to display the devices to or from which these requests were sent or received.

The response code bar categories include:

NOERROR

Successful transaction; no error.

FORMERROR

Format Error.

SERVFAIL

DNS Server Failed.

NXDOMAIN

No such domain.

NOTIMPL

No handler implemented for this query type.

REFUSED

Query administratively refused.

UPDATEERR

Error in handling UPDATE request.

TSIGERR

Error in handling TSIG request.

OTHER

All other response code types.

Transactions Per Second

Displays the number of DNS transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red data points to list the peer devices associated with the errors at this point in time. Click and drag across the chart to select a particular region.

Server Processing Time

Displays the median server processing time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing time metrics. Click and drag across the chart to select a particular region.

Processing Time Distribution

Displays a histogram of times it took the server to process requests. Move the mouse pointer over each bar to display the time range it represents and the number of requests in this bin.

Requests by Record Type

Shows the categorization of all request types sent or received by the current device. Click a bar to display the device to which (if client) or from which (if server) the query was sent.

The request query bar categories displayed include:

- A. Address
- NS. Name Server
- CNAME. Canonical Name
- SOA. Start Of Authority
- PTR. Pointer Record

- MX. Mail Exchanger
- TXT. Text
- AAAA. IPv6 Address
- SRV. Service
- TSIG. Secured Signed Request class
- IXFR. Incremental Zone Transfer
- AXFR. Zone Transfer
- ANY. Any available
- Other. All other categories

Responses by Record Type

Shows the categorization of all response types sent or received by the current device. Click a bar to display the device from which (if client) or to which (if server) the response was sent.

The request query bar categories displayed include:

- A. Address
- NS. Name Server
- CNAME. Canonical Name
- SOA. Start Of Authority
- PTR. Pointer Record
- MX. Mail Exchanger
- TXT. Text
- AAAA. IPv6 Address
- SRV. Service
- TSIG. Secured Signed Request class
- IXFR. Incremental Zone Transfer
- AXFR. Zone Transfer
- Other. All other categories

It is possible for multiple answers to be sent in response to a single query.

DNS groups page

DNS Groups Toolbar

The DNS groups toolbar includes the following controls:

DNS Metric Type

Displays metrics for members in the current group acting as a DNS client or DNS server, respectively.

Errors

Displays the number of query errors by host.

Host Queries

Displays the list of DNS queries made to or from any member in the current group. The list is sorted by Host Query frequency. Click the Query Errorsheader to sort the list by the number of DNS errors encountered.

Records

Displays results for records that match the selected metric source and protocol.

DNS Client

If you select **Client** for the DNS Metric Type, the Discover appliance displays the following metrics. Click the metric to break down DNS requests by group members in the table at the bottom of the page.

Requests

Specifies the number of DNS requests made by all members of the group.

Request Timeouts

Specifies the number of DNS requests made by any member of the group to which no response was received.

Truncated Requests

Specifies the number of malformed, truncated DNS requests sent by any member of the group.

Responses

Specifies the number of DNS responses received by all members of the group.

Response Errors

Specifies the number of DNS response errors received by all members of the group.

Truncated Responses

Specifies the number of malformed, truncated DNS responses received by all members of the group.

DNS Server

If you select **Server** for the DNS Metric Type, the Discover appliance displays the following metrics. Click the metric to break down DNS requests by group members in the table at the bottom of the page.

Requests

Specifies the number of DNS requests received by all members of the group.

Request Timteouts

Specifies the number of DNS requests received by any member of the group to which no response was sent.

Truncated Requests

Specifies the number of malformed, truncated DNS requests received by all members of the group.

Responses

Specifies the number of DNS responses sent by all members of the group.

Response Errors

Specifies the number of DNS response errors sent by all members of the group.

Truncated Responses

Specifies the number of malformed, truncated DNS responses sent by all members of the group.

Requests by Opcode

Shows the breakdown of all opcodes sent (if server) or received (if client) by members in the selected group. For each opcode, click to break down by group members in the table at the bottom of the page.

Query

Specifies the number of DNS QUERY Opcodes sent or received by all members of the group. DNS Queries are the most-frequently encountered DNS Opcode type.

Notify

Specifies the number of DNS NOTIFY Opcodes sent or received by all members of the group. DNS Notify is used as a synchronization method between DNS servers.

Update

Specifies the number of DNS UPDATE Opcodes sent or received by all members of the group. DNS Update is used as a synchronization method between DNS servers.

Other

Specifies the number of other miscellaneous DNS Opcodes sent or received by all members of the group.

Requests by Record Type

Shows the breakdown of all request types sent or received by members in the selected group. For each query type, click to break down by group members in the table at the bottom of the page.

The request query bar categories displayed include:

- A. Address
- NS. Name Server
- CNAME. Canonical Name
- SOA. Start Of Authority
- PTR. Pointer Record
- MX. Mail Exchanger
- TXT. Text
- AAAA. IPv6 Address
- SRV. Service
- TSIG. Secured Signed Request class
- IXFR. Incremental Zone Transfer
- AXFR. Zone Transfer
- ANY. Any available
- Other. All other categories

Responses by Record Type

Shows the breakdown of all record types sent (if server) or received (if client) by members in the selected group. For each query type, click to break down by group members in the table at the bottom of the page.

The request query bar categories displayed include:

- A. Address
- NS. Name Server
- CNAME. Canonical Name
- SOA. Start Of Authority
- PTR. Pointer Record
- MX. Mail Exchanger
- TXT. Text
- AAAA. IPv6 Address
- SRV. Service
- TSIG. Secured Signed Request class
- IXFR. Incremental Zone Transfer
- AXFR. Zone Transfer
- ANY. Any available
- Other. All other categories

Responses by Response Code

Shows the categorization of all response codes broken down by request opcode and request record type sent (if server) or received (if client) by members in the selected group. The format of the entry is `ERROR/REQUEST_OPCODE:REQUEST_RECORD`.

The response code bar categories include:

- NOERROR. Successful transaction; no error.
- FORMERROR. Format Error.
- SERVFAIL. DNS Server Failed.

- NXDOMAIN. No such domain.
- NOTIMPL. No handler implemented for this query type.
- REFUSED. Query administratively refused.
- UPDATEERR. Error in handling UPDATE request.
- TSIGERR. Error in handling TSIG request.
- OTHER. All other response code types.

Click the counter next to the response code to break it down by group members in the table.

DNS groups processing time page

Server Processing Time

Shows median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

FIX

ExtraHop appliances collect metrics about Financial Information Exchange (FIX) activity.

FIX applications page

FIX Applications Toolbar

The FIX applications toolbar includes the following controls:

Errors

The chart shows the number of FIX errors. Mouse over the points to view a summary of a specific time or date. The table lists FIX error messages and the number of times each occurred.

Senders

The chart shows showing the number of FIX senders. Mouse over the points to view a summary of a specific time or date. The table lists senders and the count associated with each sender.

Targets

The chart shows the number of FIX targets. Mouse over the points to view a summary of a specific time or date. The table lists targets and the count associated with each target.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each

server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Target

Displays application metrics by target.

By Sender

Displays application metrics by sender.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with DNS requests.

Response L2 Bytes

The number of L2 bytes associated with DNS responses.

Request Packets

The number of packets associated with DNS requests.

Response Packets

The number of packets associated with DNS responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

FIX Metrics

Contains the following metrics:

Requests

Specifies the number of requests for the application.

Responses

Specifies the number of responses for the application.

Response Errors

Specifies the number of responses by error for the application.

Methods

Methods exchanged by device over the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads

to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

FIX devices page

FIX Devices Toolbar

The FIX device toolbar includes the following controls:

FIX Metric Type

Displays metrics for devices acting as a FIX client or FIX server.

Errors

Click the **Errors** button to display the list of FIX session-level reject reasons (error messages) sent to or received by the current device over the selected time interval. These metrics do not include the processing of order and trade errors.

Senders

Click the **Senders** button to display a list of institutions sending the FIX message as it appears in the SenderCompID field.

Targets

Click the **Targets** button to display a list of institutions receiving the FIX message as it appears in the TargetCompID field.

Records

Displays results for records that match the selected metric source and protocol.

FIX Details specifies the type of additional FIX information displayed. Moving the cursor over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays FIX metrics by IP addresses.

By Sender

Displays FIX metrics by sender.

By Target

Displays FIX metrics by target.

For example, FIX Responses is a top-level metric showing how many responses were received by the FIX server during the selected time frame. Selecting **By IP** in the drop-down list while moving the cursor over the FIX Responses counter shows which IP addresses originated these responses. Selecting **By IP** from the drop-down list while moving the cursor over the FIX Responses counter shows the IP addresses of the responses.

FIX Metrics by IP Address

Click **By IP** in the drop-down list to display the following information in the details table.

IP Address

Represents the FIX server's IP address.

Host

Represents the DNS host name of the FIX server determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding FIX server device.

<Metric value>

Displays the value for the selected metric.

FIX Metrics by Sender

Click **By Sender** in the drop-down list to display the following information in the details table.

Sender

Displays a list of senders.

<Metric value>

Displays the value for the selected metric.

FIX Metrics by Target

Click By Target in the drop-down list to display the following information in the details table.

Target

Displays a list of targets.

<Metric value>

Displays the value for the selected metric.

FIX Client

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Requests

The number of requests received.

Responses

The number of responses received.

Errors

Number of errors sent.

POS Duplicate

Number of POS duplicates recieved.

POS Resend

Number of POS resend received.

FIX Servers

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Requests

The number of requests received.

Responses

The number of responses received.

Errors

Number of errors sent.

POS Duplicate

Number of POS duplicates recieved.

POS Resend

Number of POS resend received.

Methods

Methods exchanged by device over the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Versions

FIX versions used over the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

FIX groups page

FIX Groups Toolbar

The FIX groups toolbar includes the following controls:

FIX Metric Type

Displays metrics for groups acting as a FIX client or FIX server.

Errors

Click the **Errors** button to display the list of FIX session-level reject reasons (error messages) sent to or received by the current group over the selected time interval. These metrics do not include the processing of order and trade errors.

Senders

Click the **Senders** button to display a list of institutions sending the FIX message as it appears in the SenderCompID field.

Targets

Click the **Targets** button to display a list of institutions receiving the FIX message as it appears in the TargetCompID field.

Records

Displays results for records that match the selected metric source and protocol.

FIX Client

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Requests

The number of requests received.

Responses

The number of responses received.

Errors

Number of errors sent.

POS Duplicate

Number of POS duplicates recieved.

POS Resend

Number of POS resend received.

FIX Servers

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Requests

The number of requests received.

Responses

The number of responses received.

Errors

Number of errors sent.

POS Duplicate

Number of POS duplicates recieved.

POS Resend

Number of POS resend received.

Methods

Methods exchanged by device over the selected time interval. Click the counter to display additional per-client or per-server IP address details.

FTP

ExtraHop appliances collect metrics about File Transfer Protocol (FTP) activity.

FTP applications page

FTP Applications Toolbar

The FTP application toolbar includes the following controls:

Errors

The chart shows the number FTP errors. Mouse over the points to view a summary of a specific time or date. The table lists FTP error messages and the number of times each occurred.

Warnings

The chart shows the FTP warnings (4xx error messages) transferred. The table lists the FTP warning messages and the number of times each occurred.

Users

The chart shows the number of responses and errors for all users. Mouse over the chart to view a summary of a specific time or date. The table lists users and the number of responses and errors associated with each user.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Users

Displays application metrics by user.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

FTP Metrics

Contains the following metrics:

Requests

The number of requests received.

Responses

The number of responses received.

Response Warnings

The number of responses with an FTP status code of 4xx.

Response Errors

The number of FTP response errors.

Methods

Displays the FTP commands for the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Examples of FTP commands:

CWD

Allows the user to work with a different directory or dataset for file storage or retrieval without altering his log on or accounting information.

DELE

Causes the file specified in the path name to be deleted at the server site.

EPSV

Puts connection into extended passive mode.

LIST

Gets information for a specific working directory, if explicitly specified, or the current one if none is specified.

MDTM

Gets last-modified time of a file.

MLSD

Gets the contents of a directory.

PASS

Is a Telnet string specifying the user's password. This command must be immediately preceded by the user name command.

PASV

Requests the server-DTP to "listen" on a data port (which is not its default data port) and to wait for a connection rather than initiate one on receipt of a transfer command.

PORT

Is a HOST-PORT specification for the data port to be used in data connection.

PWD

Causes the name of the current working directory to be returned in the reply.

QUIT

Terminates a USER, and if file transfer is not in progress, the server closes the control connection. If file transfer is in progress, the connection will remain open for the result response, and the server will then close it.

RETR

Causes the server-DTP to transfer a copy of the file, specified in the path name, to the server.

SIZE

Gets the size of a file.

STOR

Causes the server-DTP to accept the data transferred via the data connection, and to store the data as a file at the server site.

SYST

Used to find out the type of operating system at the server.

TYPE

Puts the transfer mode into ASCII or Binary mode.

Status Codes

Displays the FTP reply codes for the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Examples of FTP reply codes:

1xx

Positive Preliminary reply

2xx

Positive Completion reply

3xx

Positive Intermediate reply

4xx

Transient Negative Completion reply

5xx

Permanent Negative Completion reply

6xx

Protected reply

Examples of specific reply codes:

200

OK

221

Service closing control connection

225

Data connection open

226

Closing data connection

227

Entering passive mode

230

User logged in – proceed

250

Requested file action okay

500

Syntax error, command unrecognized. This may include errors such as command line too long.

501

Syntax error in parameters or arguments

502

Command not implemented

503

Bad sequence of commands

504

Command not implemented for that parameter

530

Not logged in

550

Requested action not taken – file not available

553

Requested action not taken – filename not allowed

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

FTP devices page

Note: Where file name detail is presented, the Discover appliance displays both the file path and mount point, if available. The prefix '...' indicates that either the mount point or part of the path is not available. This may occur in instances when the capture process was restarted after the "mount" or a "cd" command was issued, or when the commands were lost due to desyncs.

FTP Devices Toolbar

The FTP metrics toolbar includes the following controls:

FTP Metric Type

Displays metrics for the current device acting as an FTP client or server.

Errors

Displays the list of 5xx error messages sent to or received by the current device over the selected time interval.

Warnings

Displays the list of 4xx error messages sent to or received by the current device over the selected time interval.

Files

Displays the list of files accessed, associated bytes sent and received, and associated errors for the selected time interval.

Clients or Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

Records

Displays results for records that match the selected metric source and protocol.

FTP Details

Specifies the type of additional FTP information displayed. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays FTP metrics by IP addresses.

By User

Displays FTP metrics by user name.For example, FTP Requests is a top-level metric showing how many requests were received by the FTP server during the selected time frame. Selecting **By IP** in the drop-down list while mousing over the FTP Requests counter shows which IP addresses originated these requests. Selecting **By User** in the drop-down list while mousing over the FTP Requests counter shows which FTP user names originated these requests.

IP Address FTP Metrics

Click **By IP** in the drop-down list to display the following information in the details table.

IP Address

Represents the HTTP server's IP address.

Host

Represents the DNS host name of the FTP server determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding FTP server device. For local FTP servers, the link leads to the FTP server device. For remote FTP servers, the link leads to the gateway device through which the requests were routed.

<Metric Value>

Displays the value of the selected metric

FTP Metrics by User

Click **By User** in the drop-down list to display the following information in the details table.

Users

Represents FTP user names that originated these requests.

<Metric Value>

Displays the value of the selected metric.

IP Address FTP Metrics

When you click the counters next to individual FTP metrics, the IP Address FTP Metrics table shows details about FTP peer devices. For FTP servers, the peer devices are FTP clients. For FTP clients, the peer devices are FTP servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

FTP Server

Displays additional IP address details.

Requests

Specifies the total number of FTP requests received on the command connection when the device is acting as an FTP server.

Responses

Specifies the number of responses that the device sent when acting as an FTP server.

Errors

Specifies the number of errors sent by the FTP server.

FTP Client

Displays additional IP address details.

Requests

Specifies the total number of FTP requests sent on the command connection when the device is acting as an FTP client.

Responses

Specifies the number of responses that the device received when acting as an FTP client.

Errors

Specifies the number of errors received by the FTP client.

Data Channel

Displays additional IP address details.

Requests

Specifies the number of data channel requests sent or received by the current device.

Connects

Specifies the number of responses sent or received by the current device.

Methods

Displays the FTP commands for the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Examples of FTP commands:

CWD

Allows the user to work with a different directory or dataset for file storage or retrieval without altering his log on or accounting information.

DELE

Causes the file specified in the path name to be deleted at the server site.

EPSV

Puts connection into extended passive mode.

LIST

Gets information for a specific working directory, if explicitly specified, or the current one if none is specified.

MDTM

Gets last-modified time of a file.

MLSD

Gets the contents of a directory.

PASS

Is a Telnet string specifying the user's password. This command must be immediately preceded by the user name command.

PASV

Requests the server-DTP to "listen" on a data port (which is not its default data port) and to wait for a connection rather than initiate one on receipt of a transfer command.

PORT

Is a HOST-PORT specification for the data port to be used in data connection.

PWD

Causes the name of the current working directory to be returned in the reply.

QUIT

Terminates a USER, and if file transfer is not in progress, the server closes the control connection. If file transfer is in progress, the connection will remain open for the result response, and the server will then close it.

RETR

Causes the server-DTP to transfer a copy of the file, specified in the path name, to the server.

SIZE

Gets the size of a file.

STOR

Causes the server-DTP to accept the data transferred via the data connection, and to store the data as a file at the server site.

SYST

Used to find out the type of operating system at the server.

TYPE

Puts the transfer mode into ASCII or Binary mode.

Status Codes

Displays the FTP reply codes for the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Examples of FTP reply codes:

1xx

Positive Preliminary reply

2xx

Positive Completion reply

3xx

Positive Intermediate reply

4xx

Transient Negative Completion reply

5xx

Permanent Negative Completion reply

6xx

Protected reply

Examples of specific reply codes:

200

OK

221

Service closing control connection

225

Data connection open

226

Closing data connection

227

Entering passive mode

230

User logged in – proceed

250

Requested file action okay

500

Syntax error, command unrecognized. This may include errors such as command line too long.

501

Syntax error in parameters or arguments

502

Command not implemented

503

Bad sequence of commands

504

Command not implemented for that parameter

530

Not logged in

550

Requested action not taken – file not available

553

Requested action not taken – filename not allowed

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Request Size

Displays the range of request sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Response Size

Displays the range of response sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Read and Write Bytes

Displays the area chart containing the breakdown of bytes by reads and writes over time. Click and drag across the chart to select a particular region.

FTP devices timing page

The timing charts draw data from the Time Selector drop-down list on the navigation toolbar. The events observed during this interval are used to fill the bins of a histogram that displays a distribution of timing data. Timing charts use a logarithmic horizontal axis that simultaneously displays events that took milliseconds and those that took seconds.

Request Transfer Time

Displays a histogram of times it took to transfer requests from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Processing Time

Displays a histogram of times it took the server to process requests. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Response Transfer Time

Displays a histogram of times it took to transfer the response from the server to the client. Mouse over each bar to display the time range it represents and the number of requests in this bin.

FTP groups page

Note: Where file name detail is presented, the Discover appliance displays both the file path and mount point, if available. The prefix '...' indicates that either the mount point or part of the path is not available. This may occur in instances when the capture process was restarted after the "mount" or a "cd" command was issued, or when the commands were lost due to desyncs.

FTP Groups Toolbar

The FTP metrics toolbar includes the following controls:

FTP Metric Type

Display metrics for the current device acting as an FTP client or server, respectively.

Errors

Displays the list of 5xx error messages sent to or received by the current device over the selected time interval.

Warnings

Displays the list of 4xx error messages sent to or received by the current device over the selected time interval.

Files

Displays the list of files accessed, associated bytes sent and received, and associated errors for the selected time interval.

Records

Displays results for records that match the selected metric source and protocol.

FTP Client

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of data requests sent by the FTP client.

Responses

Specifies the number of responses received by the FTP client.

Errors

Specifies the number of errors received by the FTP client.

Warnings

Specifies the number of warnings received by the FTP client.

FTP Server

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of data requests received by the FTP server.

Responses

Specifies the number of responses sent by the FTP server.

Errors

Specifies the number of errors sent by the FTP server.

Warnings

Specifies the number of warnings received by the FTP server.

Data Channel

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of data channel requests sent or received by the current device.

Connects

Specifies the number of responses sent or received by the current device.

Methods

Displays the FTP methods for the selected time interval. Commands include RETR (get), STOR (put), and more. Click the counter next to each method to break it down by group members in the table at the bottom of the page.

Status Codes

Displays the FTP status codes for the selected time interval. Click the counter next to each status code to break it down by group members in the table.

FTP groups processing time page

The **Server Processing Time** bar graph shows median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

HTTP-AMF

ExtraHop appliances collect metrics about Hypertext Transfer Protocol (HTTP) Action Message Format (AMF) activity.

HTTP-AMF devices page

HTTP-AMF Devices Page

The HTTP-AMF device toolbar includes the following controls:

HTTP-AMF Metric Type

Displays metrics for the current device acting as an HTTP-AMF client or HTTP-AMF server.

Clients or Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

HTTP-AMF Details specifies the type of additional HTTP-AMF information displayed. Moving the mouse pointer over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays HTTP-AMF metrics by IP addresses.

By Target URI

Displays HTTP-AMF metrics by Target URI.

For example, HTTP-AMF Requests is a top-level metric showing how many requests were received by the HTTP server during the selected time frame. Selecting **By IP** in the drop-down list while moving the mouse pointer over the Requests counter shows which IP addresses originated these requests. Selecting **By Target URI** from the drop-down list while moving the mouse pointer over the HTTP-AMF Requests counter shows which URIs were accessed by the requestors.

IP Address HTTP-AMF Metrics

Click **By IP** in the drop-down list to display the following information in the details table.

IP Address

Represents the HTTP-AMF server's IP address.

Host

Represents the DNS host name of the HTTP-AMF server determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding HTTP-AMF server device. For local HTTP-AMF servers, the link leads to the HTTP server device. For remote HTTP-AMF servers, the link leads to the gateway device through which the requests were routed.

<Metric value>

Displays the value for the selected metric.

Processing Time

Represents the time in milliseconds it took for HTTP servers to process requests for the currently selected HTTP client. Timing information is expressed as a confidence interval around the mean value bounded by one standard deviation. This metric is available for successful HTTP Responses only.

HTTP-AMF Metrics by Target URI

Click **By Target URI** in the drop-down list to display the following information in the details table.

Target URI

Represents the full HTTP target URI.

<Metric value>

Displays the value for the selected metric.

Processing Time

Represents the time in milliseconds it took to process URIs requested by the currently selected HTTP client. Timing information is expressed as a confidence interval around the mean value bounded by one standard deviation. This metric is available for successful HTTP Responses only.

HTTP-AMF Client

If you select **Client** for the **HTTP-AMF Metric Type**, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of requests that the device sent when acting as an HTTP-AMF client.

Responses

Number of responses that the device received when acting as an HTTP-AMF client.

Errors

Number of HTTP-AMF errors for the selected time interval.

Requests w/o Length

Number of requests that had no length, that the device received when acting as an HTTP-AMF client.

Responses w/o Length

Number of responses that had no length, that the device sent when acting as an HTTP-AMF client.

HTTP-AMF Server

If you select **Server** for the HTTP-AMF Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of requests that the device received when acting as an HTTP-AMF server.

Responses

Number of responses that the device sent when acting as an HTTP-AMF server.

Errors

Number of HTTP-AMF errors for the selected time interval.

Requests w/o Length

Number of requests that had no length, that the device received when acting as an HTTP-AMFs server.

Responses w/o Length

Number of responses that had no length, that the device sent when acting as an HTTP-AMF server.

HTTP-AMF groups page

HTTP-AMF Client

If you select Client for the HTTP-AMF Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of HTTP-AMF requests for the selected time interval.

Responses

Number of HTTP-AMF responses for the selected time interval.

Errors

Number of HTTP-AMF errors for the selected time interval.

Requests w/o Length

Number of HTTP-AMF requests without length.

Responses w/o Length

Number of HTTP-AMF responses without length.

HTTP-AMF Server

If you select Server for the HTTP-AMF Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of HTTP-AMF requests for the selected time interval.

Responses

Number of HTTP-AMF responses for the selected time interval.

Errors

Number of HTTP-AMF errors for the selected time interval.

Requests w/o Length

Number of HTTP-AMF requests without length.

Responses w/o Length

Number of HTTP-AMF responses without length.

IBMMQ

ExtraHop appliances collect metrics about IBM message queue (IBMMQ) activity.

IBMMQ applications page

IBMMQ Applications Toobar

The IBMMQ application toolbar includes the following controls:

Errors

The chart shows the number of IBMMQ errors. Mouse over the chart to view a summary of a specific time or date. The table lists IBMMQ error messages and the number of times each occurred.

Warnings

The chart shows the IBMMQ warnings (4xx error messages) transferred. The table lists IBMMQ warning messages and the number of times each occurred.

PUT/GET Radio

The chart shows the total PUT and GET counts for all server IPs. Mouse-over the chart to view a summary of a specific time or date. The table lists server IP addresses, the host and device associated with each server, and PUT and GET count for each server.

Clients

The chart shows round-trip time for all clients. Mouse over the points to view a five-number summary of round-trip time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, and round-trip time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows round-trip time for all servers. Mouse over the points to view a five-number summary of round-trip time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, and round-trip time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Queue

Displays application metrics by queue name.

By Channel

Displays application metrics by channel.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the Request Bytes counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

IBMMQ Metrics

Contains the following metrics:

Requests

The number of IBMMQ requests.

Responses

The number of IBMMQ responses.

Client Messages

The number of IBMMQ client messages sent or received.

Server Messages

The number of IBMMQ server messages transferred.

Errors

Number of IBMMQ errors for the selected time interval.

Warnings

Number of IBMMQ warnings for the selected time interval.

Server to Server

The number of IBMMQ server-to-server message types transferred.

Client to Server

The number of IBMMQ client-to-server message types transferred.

Methods

Displays the IBMMQ methods for the selected time interval.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

MQGET and MQPUT

Displays the GET and PUT count for the current device over the selected time interval.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

Note: When the system detects only server-to-server traffic, the metrics that are gathered for client-to-server transactions only are zero or blank.

IBMMQ devices page

IBMMQ Devices Toolbar

The IBMMQ device toolbar includes the following controls:

IBMMQ Metric Type

Displays statistics for the current device acting as a IBMMQ client or server.

Errors

Displays the list of 5xx error messages sent to or received by the current device over the selected time interval.

Warnings

Displays the list of 4xx error messages sent to or received by the current device over the selected time interval.

PUT/GET Ratio

Displays the PUT and GET counts for each IBMMQ device.

IBMMQ details specify the type of additional IBMMQ information displayed. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays IBMMQ metrics by IP addresses.

By Channel

Displays IBMMQ metrics by channel.

By Queue

Displays IBMMQ metrics by queue name.

For example, IBMMQ Requests is a top-level metric showing how many requests were received by the IBMMQ server during the selected time frame. Selecting **By IP** in the drop-down list while mousing over the IBMMQ Requests counter shows which IP addresses originated these requests.

IP Address IBMMQ Metrics

Move the mouse pointer over the counter, and click **By IP** in the drop-down list to display the following information in the details table.

IP Address

Represents the IBMMQ server's IP address.

Host

Represents the DNS hostname of the IBMMQ server determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding IBMMQ server device. For local IBMMQ servers, the link leads to the IBMMQ server device. For remote IBMMQ servers, the link leads to the gateway device through which the requests were routed.

Counter Name

Identifies the metric name and count by device associated with the counter that was clicked to open this table.

Processing Time

Represents the time in milliseconds it took for IBMMQ servers to process requests for the currently selected IBMMQ client. Timing information is expressed as a confidence interval around the mean value bounded by one standard deviation. This metric is available for successful IBMMQ Responses only.

IBMMQ

Move the mouse pointer over the counter, and click **By Channel** in the drop-down list to display the following information in the details table.

IBMMQ

Represents the channel on which the IBM MQ communication is occurring.

Counter Name

Identifies the metric name and count by device associated with the counter that was clicked to open this table.

IBMMQ Metrics by Queue

Move the mouse pointer over the counter, and click **By Queue** in the drop-down list to display the following information in the details table.

IBMMQ Queue

Represents the queue name on which the IBM MQ communication is occurring.

Counter Name

Identifies the metric name and count by device associated with the counter that was clicked to open this table.

IBMMQ Client

If you select **Client** for the Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of requests that the device sent when acting as an IBM MQ client.

Responses

Number of responses that the device received when acting as an IBM MQ client.

Client Messages

Number of client messages that the device sent or received when acting as an IBM MQ client.

Server Messages

Number of server messages that the device sent or received when acting as an IBM MQ client.

Errors

When the device is acting as an IBM MQ client, the number of responses indicating an error, broken down by specific error.

Warnings

When the device is acting as an IBM MQ client, the number of responses received, broken down by IBM MQ warning message.

PCF Errors

When the device is acting as an IBM MQ client, the number of PCF error responses, broken down by specific error. Programmable command formats (PCFs) provide a way to manipulate queue manager objects, such as queues, namelists, and channels.

PCF Warnings

When the device is acting as an IBM MQ client, the number of responses received indicating a PCF warning, broken down by specific warning message. Programmable command formats (PCFs) provide a way to manipulate queue manager objects, such as queues, namelists, and channels.

IBMMQ Server

If you select **Server** for the Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table below.

Requests

Number of requests that the device received when acting as an IBM MQ server.

Responses

Number of responses that the device sent when acting as an IBM MQ server.

Client Messages

Number of client messages that the device sent or received while acting as an IBM MQ server.

Server Messages

Number of server messages that the device sent or received when acting as an IBM MQ server.

Errors

When the device is acting as an IBM MQ server, the number of responses indicating an error, broken down by specific error.

Warnings

Number of IBMMQ warnings for the selected time interval.

PCF Errors

Number of IBMMQ PCF errors sent or received within the selected time interval.

PCF Warnings

When the device is acting as an IBM MQ server, the number of responses sent indicating a PCF warning, broken down by specific warning message. Programmable command formats (PCFs) provide a way to manipulate queue manager objects, such as queues, namelists, and channels.

Methods

Displays the IBMMQ methods for the selected time interval.

Message Formats

Displays the IBMMQ message formats for the selected time interval.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Request Size

Displays the range of request sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Response Size

Displays the range of response sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

MQGET/MQPUT

Displays the GET and PUT count for the current device over the selected time interval. (Client-to-server transactions only.)

Note: When the system detects only server-to-server traffic, the metrics that are gathered for client-to-server transactions only are zero or blank.

IBMMQ devices PCF details page

Click the **PCF Details** node to display information specific to the administrative PCF channel.

IBMMQ Client

If you select **Client** for the Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of IBMMQ requests sent or received within the selected time interval.

Responses

Number of IBMMQ responses sent or received within the selected time interval.

Errors

Number of IBMMQ errors for the selected time interval.

Warnings

Number of IBMMQ warnings for the selected time interval.

IBMMQ Server

If you select **Server** for the Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table below.

Requests

Number of IBMMQ requests sent or received within the selected time interval.

Responses

Number of IBMMQ responses sent or received within the selected time interval.

Errors

Number of IBMMQ errors for the selected time interval.

Warnings

Number of IBMMQ warnings for the selected time interval.

PCF Methods

Displays the IBMMQ PCF methods for the selected time interval.

PCF Errors

Displays the IBMMQ PCF errors for the selected time interval.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

IBMMQ devices error details page

Click the **Error Details** node to display additional IBMMQ warnings and error details.

IBMMQ groups page

IBMMQ Groups Toolbar

The IBMMQ groups toolbar includes the following controls:

IBMMQ Metric Type

Displays metrics for members in the current group acting as an IBMMQ client or IBMMQ server, respectively.

Errors

Displays the list of 5xx error messages sent to or received by the current member over the selected time interval.

Warnings

Displays the list of 4xx error messages sent to or received by the current member over the selected time interval.

IBMMQ Client

If you select **Client** for the Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of IBMMQ requests sent or received within the selected time interval.

Responses

Number of IBMMQ responses sent or received within the selected time interval.

Client Messages

Number of IBMMQ client messages sent or received within the selected time interval.

Server Messages

Number of IBMMQ server messages sent or received within the selected time interval.

Errors

Number of IBMMQ errors for the selected time interval.

Warnings

Number of IBMMQ warnings for the selected time interval.

PCF Errors

Number of IBMMQ PCF errors sent or received within the selected time interval.

PCF Warnings

Number of IBMMQ PCF requests sent or received within the selected time interval.

IBMMQ Server

If you select **Server** for the Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of IBMMQ requests sent or received within the selected time interval. (Client-to-server transactions only.)

Responses

Number of IBMMQ responses sent or received within the selected time interval. (Client-to-server transactions only.)

Client Messages

Number of IBMMQ client messages sent or received within the selected time interval.

Server Messages

Number of IBMMQ server messages sent or received within the selected time interval.

Errors

Number of IBMMQ errors for the selected time interval.

Warnings

Number of IBMMQ warnings for the selected time interval.

PCF Errors

Number of IBMMQ PCF errors sent or received within the selected time interval.

PCF Warnings

Number of IBMMQ PCF requests sent or received within the selected time interval.

Methods

Displays the IBMMQ methods for the selected time interval.

Message Format

Displays the IBMMQ message formats for the selected time interval.

Note: When the system detects only server-to-server traffic, the metrics that are gathered for client-to-server transactions only are zero or blank.

ICA

ExtraHop appliances collect metrics about Independent Computing Architecture (ICA) activity.

ICA applications page

ICAP Application Toolbar

The ICA application toolbar includes the following controls:

Users

The chart shows the total number of launches compared to load time. Mouse over the points to view a five-number summary of load time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists users, the number of launches by each user, and the login time, load time, network latency, and round-trip time for each user. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Applications

The chart shows the total number of launches compared to load time. Mouse over the points to view a five-number summary of load time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists applications, the number of launches by each application, and the login time, load time, network latency, and round-trip time for each application. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Clients

The chart shows the total number of launches compared to load time. Mouse over the points to view a five-number summary of load time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of launches by each client, and the login time, load time, network latency, and round-trip time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of launches compared to load time. Mouse over the points to view a five-number summary of load time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of launches by each server, and the login time, load time, network latency, and round trip time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Auth Domains

The chart shows the total number of launches compared to load time. Mouse over the points to view a five-number summary of load time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists domains, the number of launches by each domain, and the login time, load time, network latency, and round-trip time for each domain. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By User

Displays application metrics by user.

By Application

Displays application metrics by application. When a Citrix flow is opaque to analysis, whether because of lost segments or RC5 encryption, the reported application name is ICA or CGP.

By Auth Domain

Displays application metrics auth domain.

L2-L4 Metrics

Contains the following metrics:

Client L2 Bytes

The number of L2 bytes transmitted by the Citrix ICA client.

Server L2 Bytes

The number of L2 bytes transmitted by the Citrix ICA server.

Client Packets

The number of packets transmitted by Citrix ICA clients.

Server Packets

The number of packets transmitted by the Citrix ICA server.

Client RTOs

The number of retransmission timeouts caused by congestion when clients were sending Citrix ICA messages. A retransmission timeout is a one-second stall in the TCP connection flow due to excessive retransmissions.

Server RTOs

The number of retransmission timeouts caused by congestion when servers were sending Citrix ICA messages. A retransmission timeout is a one-second stall in the TCP connection flow due to excessive retransmissions.

Client Nagle Delays

The number of connection delays due to a bad interaction between Nagle's Algorithm and delayed ACKs. In some cases, disabling Nagle's Algorithm can mitigate the problem. On the BIG-IP Application Delivery Controller, the Nagle setting in the TCP profile should be disabled and ack_on_push should be enabled.

Server Nagle Delays

The number of connection delays due to a bad interaction between Nagle's Algorithm and delayed ACKs. In some cases, disabling Nagle's Algorithm can mitigate the problem. On the BIG-IP Application Delivery Controller, the Nagle setting in the TCP profile should be disabled and ack_on_push should be enabled.

Client Zero Windows

The number of zero window advertisements sent by clients. A device advertises a zero window when it cannot process incoming data as quickly as it is arriving.

Server Zero Window

The number of zero window advertisements sent by servers. A device advertises a zero window when it cannot process incoming data as quickly as it is arriving.

ICA Metrics

Contains the following metrics:

Client Messages

The number of Citrix ICA client messages transmitted.

Server Messages

The number of Citrix ICA server messages transmitted.

Client CGP Messages

The number of CGP messages sent by the Citrix ICA client. The Client Gateway Protocol (CGP) encapsulates Citrix ICA traffic in support of Session Reliability.

Server CGP Messages

The number of CGP messages sent by the Citrix ICA server. The Client Gateway Protocol (CGP) encapsulates Citrix ICA traffic in support of Session Reliability.

Launches

The number of Citrix ICA sessions that were launched. This count includes encrypted sessions.

Aborts

The number of Citrix ICA sessions that were initiated but closed before a Citrix application finished loading.

Encrypted

The number of Citrix ICA sessions that used an encryption method other than Basic. Certain metrics are not available for these sessions.

Screen Updates Per Second

Displays the number of screen updates per second as a function of time over the selected time interval.

Load Time (ms)

The amount of time from the beginning of the flow until the Discover appliance detects traffic on one of the following virtual channels: Clipboard, Citrix Windows Multimedia Redirection, Citrix Control Virtual Channel, or Zero Latency Font and Keyboard. Subsequent application data launched over the same session is recorded as a launch but does not factor into the load time. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the load time metrics. Click the chart to display a statistical distribution of load time per application for the selected time interval.

Network Latency (ms)

Displays the detected network latency between the ICA client and server as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the network latency metrics. Click the chart to display a statistical distribution of client latency per application for the selected time interval.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Application Launches

Displays the number of ICA launches as a function of time over the selected time interval. The chart is annotated with red data points to indicate aborts. The volume of aborts is denoted by the height of red bars under the chart. Click the red dot to see per-server or per-client details for errors associated with that dot. Click and drag across the chart to select a particular region.

App Client Bytes

Click the chart to display the total bytes per application transmitted within the selected time interval. Click the legend next to the application name to filter the information by application in the Bytes by Virtual Channels table below.

App Server Bytes

Click the chart to display the total bytes per application transmitted within the selected time interval. Click the legend next to the application name to filter the information by application in the Bytes by Virtual Channels table below.

Bytes by Virtual Channel

Displays the breakdown of ICA throughput by virtual channel. If a specific application is selected in the App Client Bytes and App Server Bytes charts above, virtual channel information is displayed specific to the selected application.

Name

Name of the application.

Client Bytes

Represents the client byte count for the currently selected application in the above chart.

Server Bytes

Represents the server byte count for the currently selected application in the above chart.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

ICA devices page

The ICA device toolbar includes the following controls:

ICA Devices Toolbar

The ICA device toolbar includes the following controls:

ICA Metric Type

Displays metrics for devices acting as an ICA client or ICA server.

Users

Click the **Users** button to display the ICA Server or Client: Users information for that device.

All Names

The load time for each user over the selected time interval.

Name

The Citrix user ID.

Load Time (ms)

The amount of time to load the application, including the login time. Load time is measured only for the first application that is loaded. Subsequent application data launched over the same session is recorded as a launch but does not factor into the load time.

Login Time (ms)

The amount of time to log in to the application. Login time is a sub-component of the load time. When the user has gained access through a previous launch, there is no login, so login time for that user is 0.

Network Latency (ms)

Displays the detected network latency between the ICA client and server as a function of time over the selected time interval.

Session Duration (sec)

The duration of each user's session.

Sessions

Click the **Sessions** button to display the ICA Client or Server: Sessions table for the device.

Name

The application name.

Duration (s)

The session duration by application.

Client Types

Click the **Client Types** button to display the ICA Client or Server: Client Types information for the device.

All Names

The number of launches for Citrix receivers over the selected time interval.

Name

The name and version of the Citrix receiver.

Count

Number of launches from that particular version of the receiver.

Auth Domain

Click the **Auth Domain** button to display the ICA Server or Client: Auth. Domain information for that device.

All Names

The load time for each user over the selected time interval.

Name

The device name.

Load Time (ms)

The time from the beginning of the flow until the Discover appliance detects traffic on one of the following virtual channels:

- Clipboard
- Citrix Windows Multimedia Redirection
- Citrix Control Virtual Channel
- Zero Latency Font and Keyboard

Login Time (ms)

The time between the transmission of the Citrix ICA packet that the client sends to the server with its credentials and the Citrix ICA packet that the server sends back to the client with the user name.

Network Latency (ms)

Displays the detected network latency between the ICA client and server as a function of time over the selected time interval.

Session Duration (ms)

The duration of each authentication session.

ICA details specify the type of additional ICA information displayed. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By User

Displays ICA device information by user.

By Application

Displays ICA device information by application. When a Citrix flow is opaque to analysis, whether because of lost segments or RC5 encryption, the reported application name is ICA or CGP.

By IP

Displays ICA device information by IP address.

By Auth Domain

Displays ICA device information by auth domain.

For example, ICA Requests is a top-level metric showing how many requests were received by the ICA server during the selected time frame. Selecting **By IP** in the drop-down list while mousing over the ICA Requests counter shows which IP addresses originated these requests.

Applications

Contains the following metrics:

Launches

Total number of Citrix ICA launch commands within the selected time interval.

Aborts

Total number of Citrix ICA sessions that were initiated but closed before a Citrix application finished loading within the selected time interval.

Encrypted Sessions

Number of encrypted sessions within the selected time interval.

ICA Client or Server

If you select Client or Server for the ICA Metric Type, the Discover appliance displays the following metrics:

Client Messages

Number of ICA client messages sent or received within the selected time interval.

Server Messages

Number of ICA server messages sent or received within the selected time interval.

Client CGP Messages

Number of client CGP messages sent by the client within the selected time interval. The Client Gateway Protocol (CGP) encapsulates ICA traffic.

Server CGP Messages

Number of CGP messages sent by the server within the selected time interval. The Client Gateway Protocol (CGP) encapsulates ICA traffic.

ICA groups page

ICA Groups Toolbar

The ICA groups toolbar includes the following controls:

ICA Metric Type

Click the **Metric Type** drop-down list, and select either **Client** or **Server** to display metrics for members in the current group acting as an ICA client or ICA server, respectively.

Applications

Click the **Applications** button to display the ICA Client or Server: Applications table.

Name

The Citrix user ID.

Launches

Number of Citrix ICA launch commands within the selected time interval.

Aborts

Number of Citrix session aborts within the selected time interval.

Sessions

Click the **Sessions** button to display the ICA Client or Server: Sessions table.

Name

The Citrix user ID.

Duration (sec)

The session duration by application.

Client Types

Click the **Client Types** button to display the ICA Client or Server: Client Types table.

Name

The name and version of the Citrix receiver.

Count

Number of launches from that particular version of the receiver.

Launches

Total number of Citrix ICA launch commands within the selected time interval.

Aborts

Total number of Citrix session aborts within the selected time interval.

ICA Client or Server

If you select Client or Server for the ICA Metric Type, the Discover appliance displays the following metrics:

Client Messages

Number of ICA client messages sent or received within the selected time interval.

Server Messages

Number of ICA server messages sent or received within the selected time interval.

Client CGP Messages

Number of ICA client CGP messages sent or received within the selected time interval.

Server CGP Messages

Number of ICA server CGP messages sent or received within the selected time interval.

iSCSI

ExtraHop appliances collect metrics about Internet Small Computer System Interface (iSCSI) activity.

iSCSI devices page

iSCSI Device Toolbar

The iSCSI device toolbar includes the following controls:

iSCSI Metric Type

Displays metrics for the current device acting as an iSCSI client or iSCSI server.

Errors

Displays the list of error messages broken down by iSCSI initiator sent to or received by the current device over the selected time interval.

OpCodes

Displays the list of iSCSI operation codes broken down by iSCSI initiator sent to or received by the current device over the selected time interval.

Initiators

Displays the list of iSCSI initiators establishing connections to or from the current device over the selected time interval.

IP Address iSCSI Metrics

Click the counters next to individual iSCSI metrics to show the IP Address iSCSI Metrics for iSCSI peer devices. For iSCSI servers, the peer devices are iSCSI clients. For iSCSI clients, the peer devices are iSCSI servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

Target

Displays corresponding iSCSI targets.

iSCSI Server

Click the counter next to each metric to display additional IP address details.

Responses

Specifies the number of responses that the device sent when acting as an iSCSI target.

Errors

Specifies the number of errors sent by the iSCSI server.

Sessions

Specifies the number of iSCSI sessions that the device began when acting as an iSCSI target.

Reads (DataOut)

Specifies the number of read operation requests that the device received when acting as an iSCSI target.

Writes (DataIn)

Specifies the number of write operation requests that the device received when acting as an iSCSI target.

Header Digest

Specifies the number of operations that included optional header digests when the device is acting as an iSCSI target.

Data Digest

Specifies the number of operations that included optional data digests when the device is acting as an iSCSI target.

iSCSI Client

Click the counter next to each metric to display additional IP address details.

Responses

Specifies the number of responses that the device received when acting as an iSCSI initiator.

Errors

Specifies the number of errors sent by the iSCSI client.

Sessions

Specifies the number of iSCSI sessions that the device began when acting as an iSCSI initiator.

Reads (DataOut)

Specifies the number of read operation requests that the device sent when acting as an iSCSI initiator.

Writes (DataIn)

Specifies the number of write operation requests that the device sent when acting as an iSCSI initiator.

Header Digest

Specifies the number of operations that included optional header digests when the device is acting as an iSCSI initiator.

Data Digest

Specifies the number of operations that included optional data digests when the device is acting as an iSCSI initiator.

OpCodes

Displays the list of iSCSI OpCodes sent to or received by the current device over the selected time interval. Click the counter to display additional per-client or per-server IP address details. Click the **OpCodes** button to get OpCodes broken down by iSCSI initiator. OpCodes include:

- Login Request
- Login Response
- Logout Request
- Logout Response
- SCSI Command
- SCSI Response
- Text Request
- Text Response
- SCSI Data-In
- SCSI Data-Out
- SCSI Task Management Response
- SCSI Task Management Function Request
- Ready To Transfer
- Asynchronous Message
- SNACK Request
- Reject
- Last
- NOP-In
- NOP-Out
- Vendor-<hex>

Rejects

Displays the list of reject reasons sent to or received by the current device over the selected time interval. Click the counter to display additional per-client or per-server IP address details. Click the **Errors** button to get errors broken down by iSCSI initiator. Reject reasons include:

- Zero
- Reserved
- Data Digest Error
- SNACK Reject
- Protocol Error
- Command not supported

- Protocol Error
- Immediate Command Reject
- Task in progress
- Invalid Data ACK
- Invalid PDU field
- Long Operation Reject
- Negotiation Reset
- Waiting for Logout

Logins

Displays the iSCSI login errors for the selected time interval. Click the counter to display additional per-client or per-server IP address details. Click the Errors button to get errors broken down by iSCSI initiator.

- Login failures
- Target moved temporarily
- Target moved permanently
- Initiator error
- Authentication failure
- Authorization failure
- Not found
- Target removed
- Unsupported version
- Too many connections
- Missing parameter
- Can't include in session
- Session type not supported
- Session does not exist
- Invalid request during login
- Target errror
- Service unavailable
- Out of resources

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Read and Write Bytes

Displays the area chart containing the breakdown of bytes by reads and writes over time. Click and drag across the chart to select a particular region.

iSCSI groups page

iSCSI Groups Toolbar

The iSCSI groups toolbar includes the following controls:

iSCSI Metric Type

Displays metrics for members in the current group acting as an iSCSI client or server, respectively.

Errors

Displays the list of error messages sent to or received by members in the current group over the selected time interval.

OpCodes

Displays the list of iSCSI operation codes broken down by iSCSI initiator sent to or received by members in the current group over the selected time interval.

Initiators

Displays the list of iSCSI initiators establishing connections to or from members in the current group over the selected time interval.

iSCSI Server

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Responses

Specifies the number of responses sent by the iSCSI server.

Errors

Specifies the number of errors sent by the iSCSI server.

Sessions

Specifies the number of iSCSI sessions received by the iSCSI server.

Reads (DataOut)

Specifies the number of read operations requested from the iSCSI server.

Writes (DataIn)

Specifies the number of write operations requested from the iSCSI server.

Header Digest

Specifies the number of iSCSI operations with optional header digests included.

Data Digest

Specifies the number of iSCSI operations with optional data digests included.

iSCSI Client

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Responses

Specifies the number of responses received by the iSCSI client.

Errors

Specifies the number of errors sent by the iSCSI client.

Sessions

Specifies the number of iSCSI sessions received by the iSCSI server.

Reads (DataOut)

Specifies the number of read operations requested from the iSCSI server.

Writes (DataIn)

Specifies the number of write operations requested from the iSCSI server.

Header Digest

Specifies the number of iSCSI operations with optional header digests included.

Data Digest

Specifies the number of iSCSI operations with optional data digests included.

OpCodes

Displays the list of iSCSI OpCodes sent to or received by members in the current group over the selected time interval. Click the counter next to the metric to break it down by group members in the table at the bottom of the page. Click the **OpCodes** button to get OpCodes broken down by iSCSI initiator. OpCodes include:

- Login Request

- Login Response
- Logout Request
- Logout Response
- SCSI Command
- SCSI Response
- Text Request
- Text Response
- SCSI Data-In
- SCSI Data-Out
- SCSI Task Management Response
- SCSI Task Management Function Request
- Ready To Transfer
- Asynchronous Message
- SNACK Request
- Reject
- Last
- NOP-In
- NOP-Out
- Vendor-<hex>

Rejects

Displays the list of reject reasons sent to or received by the current member over the selected time interval. Click the counter next to the metric to break it down by group members in the table at the bottom of the page. Click the **Errors** button to get errors broken down by iSCSI initiator. Reject reasons include:

- Zero
- Reserved
- Data Digest Error
- SNACK Reject
- Protocol Error
- Command not supported
- Protocol Error
- Immediate Command Reject
- Task in progress
- Invalid Data ACK
- Invalid PDU field
- Long Operation Reject
- Negotiation Reset
- Waiting for Logout

Logins

Displays the iSCSI login errors for the selected time interval. Click the counter next to the metric to break it down by group members in the table at the bottom of the page. Click the **Errors** button to get errors broken down by iSCSI initiator.

- Login failures
- Target moved temporarily
- Target moved permanently
- Initiator error
- Authentication failure
- Authorization failure
- Not found

- Target removed
- Unsupported version
- Too many connections
- Missing parameter
- Can't include in session
- Session type not supported
- Session does not exist
- Invalid request during login
- Target errror
- Service unavailable
- Out of resources

Kerberos

ExtraHop appliances collect metrics about Kerberos activity.

Kerberos applications page

Kerberos Application Toolbar

The Kerberos application toolbar includes the following controls:

Errors

The chart shows the number of Kerberos errors. Mouse over the chart to view a summary of a specific time or date. The table lists Kerberos error messages and the number of times each occurred.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Records

Displays results for records that match the selected metric source and protocol.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

Kerberos Metrics

Contains the following metrics:

Requests

The number of requests received.

Responses

The number of responses received.

Response Errors

The number of response errors.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

Kerberos devices page

Kerberos Device Toolbar

The Kerberos device toolbar includes the following controls:

Kerberos Metric Type

From the drop-down menu, select the type of metrics for the current device.

Errors

Displays the list of error messages sent or received by the current device over the selected time interval.

Clients or Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

FTP Server

Displays additional IP address details.

Requests

The number of requests received.

Responses

The number of responses received.

Errors

Specifies the number of errors sent by the server.

FTP Client

Displays additional IP address details.

Requests

The number of requests received.

Responses

The number of responses received.

Errors

Specifies the number of errors sent by the client.

Requests by Message Type

Displays the number of requests that the device received for the message type.

Responses by Message Type

Displays the number of requests that the device received for the message type.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Kerberos groups page

Kerberos Groups Toolbar

The Kerberos groups toolbar includes the following controls:

Kerberos Metric Type

Displays metrics for members in the current group acting as an Kerberos client or server, respectively.

Errors

Displays the list of error messages sent to or received by members in the current group over the selected time interval.

Kerberos Client

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of data requests sent by the Kerberos client.

Responses

Specifies the number of responses received by the Kerberos client.

Errors

Specifies the number of errors received by the Kerberos client.

Warnings

Specifies the number of warnings received by the Kerberos client.

Kerberos Server

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of data requests received by the Kerberos server.

Responses

Specifies the number of responses sent by the Kerberos server.

Errors

Specifies the number of errors sent by the Kerberos server.

Warnings

Specifies the number of warnings received by the Kerberos server.

Data Channel

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of data channel requests sent or received by the current device.

Connects

Specifies the number of responses sent or received by the current device.

Methods

Displays the Kerberos methods for the selected time interval. Commands include RETR (get), STOR (put), and more. Click the counter next to each method to break it down by group members in the table at the bottom of the page.

Status Codes

Displays the Kerberos status codes for the selected time interval. Click the counter next to each status code to break it down by group members in the table.

L2

ExtraHop appliances collect metrics about L2 activity.

L2 devices page

VLAN Tagged

The number of frames containing VLAN tags observed over the selected time interval. `In` reflects number of VLAN tagged frames received by the device. `Out` reflects number of VLAN tagged frames sent by the device.

Packets

Displays the incoming and outgoing packet rate (packets per second) over the selected time interval. `Current` and `Max` identify the current and maximum packet rates for the given time period, respectively. `Total` identifies the total number of packets for the selected time interval. To view specific statistics for each data point, hover the mouse across the chart to see the packets per second value for each unit on the x-axis of the graph.

Throughput

Displays the incoming and outgoing throughput (bits per second) over the selected time interval. `Current` and `Max` identify the current and maximum throughputs. `Total` identifies the total number of bytes transferred over the selected time interval. To view specific statistics for each data point, move the mouse pointer across the chart to see the throughput in megabits per second for each unit on the x-axis of the graph.

Frame Count by Size

Displays a logarithmic-scale histogram of the distribution of incoming and outgoing Ethernet frame size.

Frame Count by Type

Displays a logarithmic-scale histogram of the distribution of frames by L2 Ethertype (ipv4, ipv6, arp, ipx, mpls, lacp, stp, 802.1X, and other).

Frame Count by Distribution

Displays a logarithmic-scale histogram of the distribution of frames by L2 type (unicast, multicast, and broadcast).

Note: One-second aggregation metrics are available when the specified time interval is six minutes or less. For more information, see the Time Selector section.

L2 devices packets page

The Packets In and Packets Out line charts display the packet rate (in packets per second) for the selected device over the given time interval.

Note: One-second aggregation metrics are available when the specified time interval is six minutes or less. For more information, see the Time Selector section.

L2 devices throughput page

The Throughput In and Throughput Out line charts display the throughput (in bits per second) over the selected time interval.

Note: One-second aggregation metrics are available when the specified time interval is six minutes or less. For more information, see the Time Selector section.

L2 networks page

The L2 network traffic page displays metrics for OSI Layer 2 traffic by packet rate (packets per second) and throughput (in bits per second). It also provides metrics on frame count by L2 Ethertype and by frame size.

Packets

Displays the packet rate (in packets per second) for the selected time interval. On the line chart, `Current` and `Max` identify the current and maximum packet rates for the given time period. `Total` identifies the total number of packets for the selected time interval. The gray bands represent the 5th to 95th percentile of the packet rate historically observed for the specific time of day and day of the week.

Throughput

Displays the throughput (in bits per second) over the selected time interval. In the chart, `Current` and `Max` identify the current and maximum throughputs. `Total` identifies the total number of bytes transferred over the selected time interval. The gray bands represent the 5th to 95th percentile of the throughput historically observed for this time of day and day of the week.

Frame Count by Size

Displays a logarithmic-scale histogram of the distribution of Ethernet frame size. The values on the x-axis (64, 128, 256, 512, 1024, 1513, 1518, and Jumbo) indicate the maximum size of the frame for the category. For example, 256 represents a frame size between 129 and 256 bytes, inclusive.

Frame Count by Type

Displays a logarithmic-scale histogram of the distribution of frames by L2 Ethertype (IPv4, IPv6, ARP, IPX, MPLS, LACP, STP, 802.1X, and other).

Frame Count by Distribution

Displays a logarithmic-scale histogram of the distribution of frames by L2 type (Unicast, Multicast, and Broadcast).

Note: One-second aggregation metrics are available when the specified time interval is six minutes or less. For more information, see the Time Selector section.

L2 networks packets page

The **Packets** line chart displays the packet rate (in packets per second) for the selected time interval.

Note: One-second aggregation metrics are available when the specified time interval is six minutes or less. For more information, see the Time Selector section.

L2 networks throughput page

The **Throughput** line chart displays the throughput (in bits per second) over the selected time interval.

Note: One-second aggregation metrics are available when the specified time interval is six minutes or less. For more information, see the Time Selector section.

L2 networks frame details page

Frame Count by Size

Displays a logarithmic-scale histogram of the distribution of Ethernet frame size. The values on the x-axis (64, 128, 256, 512, 1024, 1513, 1518, and Jumbo) indicate the maximum size of the frame for the category. For example, 256 represents a frame size between 129 and 256 bytes, inclusive.

Frame Count by Type

Displays a logarithmic-scale histogram of the distribution of frames by L2 Ethertype (IPv4, IPv6, ARP, IPX, MPLS, LACP, STP, 802.1X, and other).

Frames

Displays a list of devices and the frame count in and out for a specified frame type. To select a frame type, click a bar in the Frame Count by Size orFrame Count by Type tables.

L2 groups page

VLAN Tagged

The number of frames containing VLAN tags observed over the selected time interval. `In` reflects number of VLAN tagged frames received by the device. `Out` reflects number of VLAN tagged frames sent by the device.

Packets

Displays the incoming and outgoing packet rate (packets per second) over the selected time interval. `Current` and `Max` identify the current and maximum packet rates for the given time period, respectively. Total identifies the total number of packets for the selected time interval. To view specific statistics for each data point, hover the mouse across the chart to see the packets per second value for each unit on the x-axis of the graph.

Throughput

Displays the incoming and outgoing throughput (bits per second) over the selected time interval. `Current` and `Max` identify the current and maximum throughputs. Total identifies the total number of bytes transferred over the selected time interval. To view specific statistics for each data point, move the mouse pointer across the chart to see the throughput in megabits per second for each unit on the x-axis of the graph.

Frame Count by Size

Displays a logarithmic-scale histogram of the distribution of incoming and outgoing Ethernet frame size.

Frame Count by Type

Displays a logarithmic-scale histogram of the distribution of frames by L2 Ethertype (ipv4, ipv6, arp, ipx, mpls, lacp, stp, 802.1X, and other).

Frame Count by Distribution

Displays a logarithmic-scale histogram of the distribution of frames by L2 type (unicast, multicast, and broadcast).

L2 groups packets page

Packets In

Displays how members contribute to the total incoming packet count for the group.

Packets Out

Displays how members contribute to the total incoming packet count for the group.

L2 groups throughput page

Bytes In

Displays how members contribute to the total incoming byte count for the group.

Bytes Out

Displays how members contribute to the total incoming byte count for the group.

L3

ExtraHop appliances collect metrics about L3 activity.

L3 devices device page

Name

The primary name the device uses to communicate on the network. Names are discovered by passively monitoring a variety of naming protocols, including DNS, DHCP, NETBIOS, and Cisco Discovery Protocol. If a device name is not discovered, a NIC manufacturer-based identifier is assigned to the device by looking at the MAC address. If the MAC address range is not registered, or if it belongs to a private MAC address space, the name includes the last six characters of the MAC address (for example, Device 00000c0789b1).

The device-type icon to the left of the device name identifies the activity primarily associated with this device. The device name and type can be edited by clicking on the name and using the edit tools on the Device page.

MAC Address

The MAC address is a unique identifier of the device network interface. For physical devices that have multiple interfaces, one entry per interface is maintained. The vendor icon displays to the left of MAC Address as determined by the MAC OID lookup.

VLAN

The VLAN tag of the device.

IP Address

The Primary IP address the device uses to communicate on the network. By default, Address Resolution Protocol (ARP) traffic is used to determine the mapping from MAC addresses to IP addresses. In the absence of such traffic, IP packet header information is used. If there is no ARP traffic, the IP address 0.0.0.0 is assigned to routing devices, such as gateways, firewalls, and load balancers, to indicate that it handles packets from many sources.

Discovery Time

The time when the device was first discovered. The day of the week, the calendar date, and time are displayed in the following format: Wed Feb 23 09:01.

Description

A user-defined description of the device. To edit the device description, click the device name and use the edit tools on the Device page.

L3 devices page

IP Fragments

Displays the IP fragments in and out for the device or group.

Packet Count by Protocol

Displays the incoming and outgoing packet count for each L3 protocol type. Click a bar in the chart to display the table of devices transmitting or receiving the selected L3 protocol.

Byte Count by Protocol

Displays the incoming and outgoing byte count for each L3 protocol type. Click a bar in the chart to display the table of devices transmitting or receiving the selected L3 protocol. IP types include TCP, UDP, ICMP, SCTP, IPSEC, GRE, ICMP6, VRRP, and OTHER.

Devices and Peer Devices

Displays IP addresses and host names with which the device or group communicates, packet in/out count, and byte in/out count for the currently selected L3 protocol. If no L3 protocol is selected, the packet count and byte count is the sum of all L3 protocol counts for the device or group. Click the device name to navigate to the device.

L3 devices DSCP page

Packets in by DSCP

Displays the number of incoming packets containing DSCP values on the network within the selected time interval. The legend lists the DSCP values with the highest count.

Packets out by DSCP

Displays the number of outgoing packets containing DSCP values on the network within the selected time interval. The legend lists the DSCP values with the highest count.

L3 devices ICMP details page

ICMP Packets In

Displays a list of ICMP response types and associated packet counts received by the current device in the selected time interval.

ICMP Packets Out

Displays a list of ICMP response types and associated packet counts sent by the current device in the selected time interval.

ICMPv6 Packets In

Displays a list of ICMPv6 response types and associated packet counts received by the current device in the selected time interval.

ICMPv6 Packets Out

Displays a list of ICMPv6 response types and associated packet counts sent by the current device in the selected time interval.

- Destination Unreachable:
  - Dest Unreach - Network
  - Dest Unreach - Host
  - Dest Unreach - Protocol
  - Dest Unreach - Port
  - Dest Unreach - Fragmentation Needed
  - Dest Unreach - Source Route
- Time Exceeded:
  - Redirect - Network
  - Redirect - Host
  - Redirect - ToS Network
  - Redirect - ToS Host
- Miscellaneous:
  - Bad Param
  - Source Quench
  - Echo
  - Echo Reply
  - Timestamp
  - Timestamp Reply
  - Info Request
  - Info Reply

- ICMPv6 Destination Unreachable:
  - Dest Unreach - No route
  - Dest Unreach - Prohibited
  - Dest Unreach - Bad scope
  - Dest Unreach - Host
  - Dest Unreach - Port
- ICMPv6 Time Exceeded:
  - Time Exceeded - Transit
  - Time Exceeded - Fragment Reassembly
- ICMPv6 Parameter Problem:
  - Bad Param - Header Error
  - Bad Param - Unknown Next Header
  - Bad Param - Unkown Option
- ICMPv6 Miscellaneous:
  - Packet Too Big
  - Echo
  - Echo Reply
  - MLD Query
  - MLD Report
  - MLD Done
  - ND Router Solicit
  - ND Router Advert
  - ND Neighbor Solicit
  - ND Neighbor Advert
  - ND Redirect
  - Router renumber
  - FQDN Query
  - FQDN Reply
  - MLDv2 Listener Report
  - MLD Mtrace Rsp
  - MLD Mtrace

L3 networks page

The L3 network traffic sub-page displays metrics for OSI Layer 3 traffic by packet count per L3 network protocol and byte count per protocol.

IP Fragments

Displays the number of IP fragments identified in the network capture.

Packet Count by Protocol

Displays the packet count for each L3 protocol type. The values on the x-axis (ICMP6, TCP, UDP, and Other) identify the common L3 protocol types.

Byte Count by Protocol

Displays the byte count for each L3 protocol type. The values on the x-axis (ICMP6, TCP, UDP, and Other) identify the common L3 protocol types.

Devices

Displays the device name, packet in/out count, byte in/out count, and IP fragment in/out count for the currently selected L3 protocol. If no L3 protocol is selected, the packet count and byte count is the sum of all L3 protocol counts for the device. Click the device name to navigate to the device details page.

L3 networks DSCP page

The DSCP sub-page displays the number of packets containing differentiated services code point (DSCP) values.

Packets by DSCP

Displays the number of packets containing DSCP values on the network within the selected time interval. The legend lists the DSCP values with the highest count.

Bytes by DSCP

Displays the number of bytes containing DSCP values on the network within the selected time interval. The legend lists the DSCP values with the highest count.

L3 groups page

IP Fragments

Displays the IP fragments in and out for the device or group.

Packet Count by Protocol

Displays the incoming and outgoing packet count for each L3 protocol type. Click a bar in the chart to display the table of devices transmitting or receiving the selected L3 protocol.

Byte Count by Protocol

Displays the incoming and outgoing byte count for each L3 protocol type. Click a bar in the chart to display the table of devices transmitting or receiving the selected L3 protocol. IP types include TCP, UDP, ICMP, SCTP, IPSEC, GRE, ICMP6, VRRP, and OTHER.

Devices and Peer Devices

Displays IP addresses and host names with which the device or group communicates, packet in/out count, and byte in/out count for the currently selected L3 protocol. If no L3 protocol is selected, the packet count and byte count is the sum of all L3 protocol counts for the device or group. Click the device name to navigate to the device.

L4

ExtraHop appliances collect metrics about L4 activity.

L4 applications page

L4 Application Toolbar

The L4 application toolbar includes the following controls:

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Connections

Displays the TCP connection metrics for the selected time interval.

Connected

Specifies the number of connections initiated by the current device. Click to display the peer devices to which the connections were established and the associated round-trip time.

Closed

Specifies the number of connections closed to or from the current device. Closed connections are explicitly shut down by at least one of the endpoints. Click to display the peer devices for which the connections were closed.

Aborted

Specifies the number of connections aborted by the current device. Aborted connections are reset explicitly by one of the endpoints. In some cases, this indicates that an error occurred. Click to display the peer devices to which the current device aborted the connections.

Expired

Specifies the number of connections to or from the current device no longer tracked due to inactivity. Click to display the peer devices with which the connections were associated.

Established

Number of connections currently open to or from the current application. Click to display the server IP addresses, hosts, and devices with which connections have been established.

Established

Maximum number of established connections observed at any point within the selected time interval.

Request Metrics

Displays the request metrics for the selected time interval.

L2 Bytes

Displays request bytes for the application within the selected time interval.

Packets

Displays request packets for the application within the selected time interval.

RTOs

Displays request RTOs for the application as a function of time within the selected time interval. Request RTOs are transmitted out of the client and into the server.

Nagle Delays

Indicates connection delays due to a bad interaction between Nagle's Algorithm and delayed ACKs. In some cases, disabling Nagle's Algorithm can mitigate the problem. On the BIG-IP Application Delivery Controller, the Nagle setting in the TCP profile should be disabled and ack_on_push should be enabled.

Rcv Wnd Throttles

Number of times the advertised receive window limits the throughput of the connection. In some cases, the read socket buffer size can be increased or receive window scaling can be enabled on the current device to resolve this problem.

Zero Window

Number of zero window advertisements sent by the current device. A zero window indicates that the connection has stalled and the current device is unable to keep up with the rate of data sent. In some cases, the read socket buffer size can be increased on the current device

to resolve this problem. On the BIG-IP Application Delivery Controller, the proxy_buffer_high setting in the TCP profile should be increased.

Response Metrics

Displays the response metrics for the specified time interval.

L2 Bytes

Displays response bytes for the application within the selected time interval.

Packets

Displays response packets for the application within the selected time interval.

RTOs

Displays response RTOs for the application as a function of time within the selected time interval. Response RTOs are transmitted out of the server and into the client.

Nagle Delays

Indicates connection delays due to a bad interaction between Nagle's Algorithm and delayed ACKs. In some cases, disabling Nagle's Algorithm can mitigate the problem. On the BIG-IP Application Delivery Controller, the Nagle setting in the TCP profile should be disabled and ack_on_push should be enabled.

Rcv Wnd Throttle

Number of times the advertised receive window limits the throughput of the connection. In some cases, the read socket buffer size can be increased or receive window scaling can be enabled on the current device to resolve this problem.

Zero Window

Number of zero window advertisements sent by the current device. A zero window indicates that the connection has stalled and the current device is unable to keep up with the rate of data sent. In some cases, the read socket buffer size can be increased on the current device to resolve this problem. On the BIG-IP Application Delivery Controller, the proxy_buffer_high setting in the TCP profile should be increased.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

L4 TCP devices page

TCP Device Toolbar

The TCP device toolbar includes the following controls:

TCP Details

Specifies what type of additional TCP information is displayed when a counter is clicked next to each top-level metric. You can choose between the following options: **By IP** for IP addresses and **By L7 Protocol**. For example, the top-level metric, TCP Closed connections, shows how many connections were closed by the current device during the selected time frame. Selecting **By IP** and clicking on the closed counter will show which IP addresses originated these connections. Selecting **By L7 Protocol** and clicking on the closed counter will show which applications were accessed by the requestor.

Connections

The TCP connection metrics for the specified time interval.

Accepted

Number of inbound connections accepted by the device. Click to display the peer devices from which the connections originated and the associated round-trip time.

Connected

Number of outbound connections initiated by the device. Click to display the peer devices to which the connections were established and the associated round-trip time.

Closed

Number of connections explicitly shut down by the device or its peer. Closed connections are explicitly shut down by at least one of the endpoints. Click to display the peer devices for which the connections were closed.

Aborted

The total number of TCP connections that were forcibly ended between the selected device and another device on the network. Aborted connections might indicate that an error occurred. For more information about the number of aborts for incoming and outgoing connections, click **Details**.

Expired

Number of connections involving the device for which tracking was stopped due to inactivity. Click to display the peer devices with which the connections were associated.

Established

For a given time interval, the number of open connections involving the device at end of the interval. Click to display the peer devices with which connections have been established.

Established Max

Maximum number of established connections observed at any point within the selected time interval.

Desync

Number of times synchronization was lost when processing TCP connections for the device. Large numbers might indicate dropped packets on the monitoring interface, SPAN, or network tap.

TCP Flow Stalls

Number of events in which the device was not responsive.

Connections Chart

Displays the number of accepted, connected, closed, and aborted connections as a function of time over the selected time interval. Click the chart to display a larger version. Date represents the date and time for the currently moused-over point on the graph. Connects, Accepts, Closes, and Abortsrepresent the number of outgoing, incoming, closed, and aborted connections respectively for the currently moused-over point in the graph. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower

quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

Throttling In: Receive Windows and Zero Windows

Represents the incoming receive and zero windows of the current device as a function of time over the selected time interval. Click and drag across the chart to select a particular region.

Throttling Out: Receive Windows and Zero Windows

Represents the outgoing receive and zero windows of the current device as a function of time over the selected time interval. Click and drag across the chart to select a particular region.

L4 TCP devices details page

Specifies what type of additional TCP information is displayed, when a counter is clicked next to each top-level metric. You can choose between the following options: **By IP** for IP addresses and **By L7 Protocol**. For example, TCP Closed connections is a top-level metric showing how many connections were closed by the current device during the selected time frame. Selecting **By IP** and clicking on the closed counter will show which IP addresses originated these connections. Selecting **By L7 Protocol** and clicking on the closed counter will show which applications were accessed by the requestor.

Connections

The TCP connection metrics for the current device.

Accepted

Number of inbound connections accepted by the device. Click to display the peer devices from which the connections originated and the associated round-trip time.

Connected

Number of outbound connections initiated by the device. Click to display the peer devices to which the connections were established and the associated round-trip time.

Closed

Number of connections explicitly shut down by the device or its peer. Closed connections are explicitly shut down by at least one of the endpoints. Click to display the peer devices for which the connections were closed.

Expired

Number of connections involving the device for which tracking was stopped due to inactivity. Click to display the peer devices with which the connections were associated.

Established

For a given time interval, the number of open connections involving the device at end of the interval. Click to display the peer devices with which connections have been established.

Established Max

Maximum number of established connections observed at any point within the selected time interval.

Desync

Number of times synchronization was lost when processing TCP connections for the device. Large numbers might indicate dropped packets on the monitoring interface, port mirror, or network tap.

In

The incoming connection metrics for the current device.

Aborts

Number of connections aborted by the peer of the current device. Aborted connections are reset explicitly by one of the endpoints. In some cases, this indicates that an error occurred. Click to display the peer devices that aborted the connections.

Resets

Number of RSTs received by the current device. TCP resets indicate that a reset packet was sent to forcibly end the TCP connection, and can be used in a variety of situations. Sometimes resets are sent when the receiving device failed to ACK the SYN packet, or it failed to acknowledge another packet sent and retransmitted later in the transaction. Other times, resets may be used to quickly and efficiently end an existing connection to free up resources for more traffic. High volumes of outbound resets should be investigated to determine if they are expected behavior or indicative of a larger issue.

SYNs Received

Number of SYNs received by the current device.

SYNs Unanswered

Number of SYNs received by the device for which there were no corresponding ACKs.

Stray Segments

Number of unexpected TCP packets received by the current device. Stray segments are likely to be recorded when the Discover appliance is first started. Continued large numbers of stray segments could indicate a misconfiguration or deployment problem.

Dropped Segments

Number of episodes in which a segment or a series of segments were lost on the way to the current device and required retransmission. Large values of this counter may indicate network congestion or link reliability problems.

Zero Window

Number of zero window advertisements received by the current device. A zero window indicates that the connection has stalled and the peer device is unable to keep up with the rate of data sent. In some cases, the read socket buffer size can be increased on the peer device to resolve this problem. On the BIG-IP Application Delivery Controller, the proxy_buffer_high setting in the TCP profile should be increased.

Rcv Wnd Throttles

Number of times the advertised receive window of the peer device limits the throughput of the connection. In some cases, the read socket buffer size can be increased or receive window scaling can be enabled on the peer device to resolve this problem.

Snd Wnd Throttles

Number of send window throttles. This indicates that the TCP congestion avoidance on the peer device might be too conservative. In some cases, a different congestion avoidance algorithm can be selected or send window scaling can be enabled on the peer device.

SYNs w/o Timestamps

Number of SYNs without the TCP timestamp option received by the current device.

SYNs w/o SACK

Number of SYNs without the TCP SackOK option received by the current device. This option is necessary to use selective acknowledgments.

RTOs

Number of retransmission timeouts caused by congestion as peers were sending data to the current device. This indicates a relatively long stall in the connection due to packet loss. Enabling selective acknowledgments and fast recovery might reduce such stalls.

Learn more about RTOs on ExtraHop.com

PAWS-Dropped SYNs

Number of PAWS-dropped SYNs. This indicates that a connection failed to initiate because the current device interpreted the SYN as belonging to a previous connection. This problem is often due to network address translation and specifically the timestamp affixed to packets that traverse a network address translation device. PAWS-dropped SYNs may cause a stall in connection setup since the dropped SYN is typically retransmitted after a three-second timer expires. In some cases, increasing the connection linger time on the NAT device or decreasing connection linger time on the current device can mitigate this problem.

Learn more about PAWS-dropped SYNs on ExtraHop.com

Bad Congestion Control

Number of events with bad congestion control, which occurs when the system receives RTOs with in-flight data greater than twice the prior congestion window. This indicates that the peer device is sending too much data, resulting in network congestion and dropped packets.

TCP Flow Stalls

Number of events in which a peer device was not responsive.

Out

The outgoing connection metrics for the current device.

Aborts

Number of connections aborted by the current device. Aborted connections are reset explicitly by one of the endpoints. In some cases, this indicates that an error occurred. Click to display the peer devices to which the current device aborted the connections.

Resets

Number of RSTs sent by the current device. TCP resets indicate that a reset packet was sent to forcibly end the TCP connection, and can be used in a variety of situations. Sometimes resets are sent when the receiving device failed to ACK the SYN packet, or it failed to acknowledge another packet sent and retransmitted later in the transaction. Other times, resets may be used to quickly and efficiently end an existing connection to free up resources for more traffic. High volumes of outbound resets should be investigated to determine if they are expected behavior or indicative of a larger issue.

SYNs Sent

Number of SYNs sent by the current device.

SYNs Unanswered

Number of SYNs sent by the device for which there were no corresponding ACKs.

Dropped Segments

Number of episodes in which a segment or a series of segments were lost on the way to the current device and required retransmission. Large values of this counter may indicate network congestion or link reliability problems.

Tinygrams

Number of tinygrams sent by the current device. This indicates that the TCP payload is being segmented inefficiently, resulting in more packets on the network.

Learn more about tinygrams on the ExtraHop Forum

Nagle Delays

Number of Nagle delays sent by the current device. This indicates connection delays due to a bad interaction between Nagle's Algorithm and delayed ACKs.

Learn more about Nagle delays on the ExtraHop Forum

Zero Window

Number of zero window advertisements sent by the current device. A zero window indicates the connection has stalled because the current device cannot handle the rate of data sent.

Slow Starts

Number of slow starts sent by the current device. This indicates that TCP slow start congestion avoidance has reduced connection throughput. The application on the current device might benefit from connection pooling or persistent connections.

Rcv Wnd Throttle

Number of times the advertised receive window of the current device limits the throughput of the connection. In some cases, the read socket buffer size can be increased or receive window scaling can be enabled on the current device to resolve this problem.

Snd Wnd Throttle

Number of send window throttles. This indicates that the TCP congestion avoidance on the current device might be too conservative. In some cases, a different congestion avoidance algorithm can be selected or send window scaling can be enabled on the current device.

SYNs w/o Timestamps

Number of SYNs without the TCP timestamp option sent by the current device.SYNs

SYNs w/o SACK

Number of SYNs without the TCP SackOK option sent by the current device. This option is necessary to use selective acknowledgments.

RTOs

Number of retransmission timeouts caused by congestion as the current device was sending data to a peer. This indicates a relatively long stall in the connection due to packet loss. Enabling selective acknowledgments and fast recovery might reduce such stalls.

Retransmissions

Number of times data is resent by the current device.

Out of Order

Number of packets sent by the device where the TCP sequence number did not match the sequence number that the Discover appliance was expecting. The reordering may have been introduced at the device itself or by an intermediate device. This can result in reduced connection throughput, increased processing load on the peer device, and additional ACK packets on the network.

Bad Congestion Control

Number of events with bad congestion control, which occurs when the system receives RTOs with in-flight data greater than twice the prior congestion window. This indicates that the current device is sending too much data, resulting in network congestion and dropped packets.

TCP Flow Stalls

Number of events in which the current device was not responsive.

L4 TCP groups page

Connections

The TCP connection metrics for all members in the current group.

Accepted

Number of connections accepted by all members in the current group. Click to break down the number of outgoing connections by each group member in the table at the bottom of the page.

Connected

Number of connections initiated by all members in the current group. Click to break down the number of incoming connections by each group member in the table at the bottom of the page.

Closed

Number of connections closed to or from any member in the current group. Closed connections are explicitly shutdown by at least one of the endpoints. Click to break down the number of closed connections by each group member in the table at the bottom of the page.

Expired

Number of connections to or from any member in the current group no longer tracked due to inactivity. Click to break down the number of expired connections by each group member in the table at the bottom of the page.

Desync

Number of times synchronization was lost when processing TCP connections from or to any member in the current group. Click to break down the number of desyncs by each group member in the table at the bottom of the page.

In

The incoming connection metrics for all members in the current group.

Aborts

Number of connections aborted by the peer of any member in the current group. Click to break down the number of aborts received by each group member in the table at the bottom of the page.

Resets

Number of RSTs received by all members in the current group. Click to break down the number of RSTs received by each group member in the table at the bottom of the page.

TCP resets indicate that a reset packet was sent to forcibly end the TCP connection, and can be used in a variety of situations. Sometimes resets are sent when the receiving member failed to ACK the SYN packet, or it failed to acknowledge another packet sent and retransmitted later in the transaction. Other times, resets may be used to quickly and efficiently end an existing connection to free up resources for more traffic. High volumes of outbound resets should be investigated to determine if they are expected behavior or indicative of a larger issue

SYNs Received

Number of SYNs received by all members in the current group. Click to break down the number of SYNs received by each group member in the table at the bottom of the page.

SYNs Unanswered

Number of SYNs received by all members in the current group for which there were no corresponding ACKs. Click to break down the number of SYNs sent by each group member in the table at the bottom of the page.

Stray Segments

Number of unexpected TCP packets received by all members in the current group. Click to break down the number of stray segments received by each group member in the table at the bottom of the page.

Dropped Segments

Number of episodes in which a segment or a series of segments were lost on the way to the current member and required retransmission. Large values of this counter may indicate network congestion or link reliability problems. Click to break down the number of inbound dropped segments by each group member in the table at the bottom of the page.

Zero Window

Number of zero window advertisements received by all members in the current group. A zero window indicates the connection has stalled because the peer member cannot handle the rate of data sent. Click to break down the number of inbound zero window advertisements by each group member in the table at the bottom of the page.

Rcv Wnd Throttles

Number of times the advertised receive window of the peer member limits the throughput of the connection. Click to break down the number of inbound receive window throttles by each group member in the table at the bottom of the page.

Snd Wnd Throttles

Number of send window throttles. This indicates that the TCP congestion avoidance on the peer member might be too conservative. Click to break down the number of inbound send window throttles by each group member in the table at the bottom of the page.

SYNs w/o Timestamps

Number of SYNs without the TCP timestamp option received by all members of the current group. Click to break down the number of inbound SYNs without timestamps by each group member in the table at the bottom of the page.

SYNs w/o SACK

Number of SYNs without the TCP SackOK option received by all members of the current group. Click to break down the number of inbound SYNs without the TCP SackOK option by each group member in the table at the bottom of the page.

RTOs

Number of retransmission timeouts caused by congestion as peers were sending data to the members of the current group. Click to break down the number of inbound RTOs by each group member in the table at the bottom of the page.

Learn more about RTOs on ExtraHop.com

PAWS-Dropped SYNs

Number of PAWS-dropped SYNs. This indicates that a connection failed to initiate because the current member interpreted the SYN as belonging to a previous connection. Click to break down the number of inbound PAWS-Dropped SYNs by each group member in the table at the bottom of the page.

Learn more about PAWS-dropped SYNs on ExtraHop.com

Bad Congestion Control

Number of events with bad congestion control, which occurs when the system receives RTOs with in-flight data greater than twice the prior congestion window. This indicates that the peer member is sending too much data, resulting in network congestion and dropped packets. Click to break down the number of bad congestion control events by each group member in the table at the bottom of the page.

TCP Flow Stalls

Number of events in which the group was not responsive. Click to break down the number of non-responsive events by each group member in the table at the bottom of the page.

Out

The outgoing connection metrics for all members in the current group.

Aborts

Number of connections aborted by any member in the current group. Aborted connections are reset explicitly by one of the endpoints. In some cases, this indicates that an error occurred. Click to display the number of aborts each group member initiated in the table at the bottom of the page.

Resets

Number of RSTs sent by all members in the current group. Click to break down the number of RSTs sent by each group member in the table at the bottom of the page.

TCP resets indicate that a reset packet was sent to forcibly end the TCP connection, and can be used in a variety of situations. Sometimes resets are sent when the receiving member failed to ACK the SYN packet, or it failed to acknowledge another packet sent and retransmitted later in the transaction. Other times, resets may be used to quickly and efficiently end an existing connection to free up resources for more traffic. High volumes of outbound resets should be investigated to determine if they are expected behavior or indicative of a larger issue.

SYNs Sent

Number of SYNs sent by all members in the current group. Click to break down the number of SYNs sent by each group member in the table at the bottom of the page.

SYNs Unanswered

Number of SYNs sent by all members in the current group for which there were no corresponding ACKs. Click to break down the number of SYNs received by each group member in the table at the bottom of the page.

Dropped Segments

Number of episodes in which a segment or a series of segments were lost on the way to the current member and required retransmission. Large values of this counter may indicate network congestion or link reliability problems. Click to break down the number of outbound dropped segments by each group member in the table at the bottom of the page.

Tinygrams

Number of tinygrams sent by the current member. This indicates that the TCP payload is being segmented inefficiently, resulting in more packets on the network. Click to break down the number of outbound tinygrams by each group member in the table at the bottom of the page.

Learn more about tinygrams on the ExtraHop blog

Nagle Delays

Number of Nagle delays sent by the current member. This indicates connection delays due to a bad interaction between Nagle's Algorithm and delayed ACKs. Click to break down the number of outbound Nagle's delays by each group member in the table at the bottom of the page.

Learn more about Nagle delays on the ExtraHop Forum

Zero Window

Number of zero window advertisements sent by all members in the current group. A zero window indicates the connection has stalled because the peer member cannot handle the rate of data sent. Click to break down the number of outbound zero window advertisements by each group member in the table at the bottom of the page.

Slow Starts

Number of slow starts sent by the current member. This indicates that TCP slow start congestion avoidance has reduced connection throughput. Click to break down the number of outbound slow starts by each group member in the table at the bottom of the page.

Rcv Wnd Throttles

Number of times the advertised receive window of the current member limits the throughput of the connection. Click to break down the number of outbound received window throttles by each group member in the table at the bottom of the page.

Snd Wnd Throttles

Number of send window throttles. This indicates that the TCP congestion avoidance on the current member might be too conservative. Click to break down the number of outbound send window throttles by each group member in the table at the bottom of the page.

SYNs w/o Timestamps

Number of SYNs without the TCP timestamp option sent by all members of the current group. Click to break down the number of outbound SYNs without timestamps by each group member in the table at the bottom of the page.

SYNs w/o SACK

Number of SYNs without the TCP SackOK option sent by all members of the current group. Click to break down the number of outbound SYNs without the TCP SackOK option by each group member in the table at the bottom of the page.

RTOs

Number of retransmission timeouts caused by congestion as members of the current group were sending data to their peers. Click to break down the number of outbound RTOs by each group member in the table at the bottom of the page.

Retransmissions

Number of times data is resent by the current member. Click to break down the number of outbound retransmissions by each group member in the table at the bottom of the page.

Out of Order

Number of packets sent by the member where the TCP sequence number did not match the sequence number that the Discover appliance was expecting. The reordering may have been introduced at the member itself or by an intermediate member. This can result in reduced connection throughput, increased processing load on the peer member, and additional ACK packets on the network. Click to break down the number of outbound retransmissions by each group member in the table at the bottom of the page.

Bad Congestion Control

Number of events with bad congestion control, which occurs when the system receives RTOs with in-flight data greater than twice the prior congestion window. This indicates that the current member is sending too much data, resulting in network congestion and dropped packets. Click to break down the number of outbound bad congestion control events by each group member in the table at the bottom of the page.

TCP Flow Stalls

Number of events in which the group was not responsive. Click to break down the number of non-responsive events by each group member in the table at the bottom of the page.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

L7

ExtraHop appliances collect metrics about L7 activity.

L7 devices page

Packets In

Displays how applications contribute to the total incoming packet count for the device. Click the chart to display a larger version. Click an application listed in the legend to list the devices sending

or receiving the traffic for that protocol in the table at the bottom of the page. Click the chart to zoom into and select a particular region.

Packets Out

Displays how applications contribute to the total outgoing packet count for the device. Click the chart to display a larger version. Click an application listed in the legend to list the devices sending or receiving the traffic for that protocol in the table at the bottom of the page. Click the chart to zoom into and select a particular region.

Bytes In

Displays how applications contribute to the total incoming byte count for the device. Click the chart to display a larger version. Click an application listed in the legend to list the devices sending or receiving the traffic for that protocol in the table at the bottom of the page. Click the chart to zoom into and select a particular region.

Bytes Out

Displays how applications contribute to the total outgoing byte count for the device. Click the chart to display a larger version. Click an application listed in the legend to list the devices sending or receiving the traffic for that protocol in the table at the bottom of the page. Click the chart to zoom into and select a particular region.

Peer Devices

Click an application listed in the legend to list the devices sending or receiving the traffic for that protocol in the table at the bottom of the page. The protocol metrics appear in a table with the following headings:

IP Address

The IP address of the corresponding device.

Host

The host of the corresponding device.

Device

A link to the corresponding device. For local devices, the link leads to that device. For remote devices, the link leads to the gateway device through which the requests were routed.

Packets In

The number of packets sent from the peer device to the current device for the selected protocol in the area chart.

Packets Out

The number of packets sent from the current device to the peer device for the selected protocol in the area chart.

Bytes In

The number of bytes sent from the peer device to the current device for the selected protocol in the area chart.

Bytes Out

The number of bytes sent from the current device to the peer device for the selected protocol in the area chart.

Note: A category labeled OTHER may appear in the legend to represent traffic that is not TCP/UDP and fails to classify as an L7 protocol. The OTHER category may also represent TCP/UDP traffic that fails to classify as an L7 protocol and fails to add an L4 p:port identifier.

The Bytes In and Bytes Out charts display activity for the top 10 protocols. To view information about other protocols, click the **Details** node in the page navigation panel.

To isolate a single protocol, mouse over the protocol in the legend or click the protocol to select it. When you select a protocol, the table displays a list of devices with activity from that protocol. Click a device in the table to view detailed L7 protocol metrics for that device.

To deselect the protocol and view all the top protocols in the chart again, click the selection in the legend again or click the table title below the charts.

L7 devices packets page
The Packets In and Packets Out area charts display the packet rate (in packets per second) for the selected device over the given time interval.

L7 devices throughput page
The Bytes In and Bytes Out area chart displays the throughput rate (in bits per second) over the selected time interval.

L7 devices turn timing page
A TCP turn is a complete change in direction of TCP payload data being delivered. In order to clearly detect this, the change in data direction must occur only after the TCP ACK is received for all the data in the prior direction, either by a bare TCP ACK or by a TCP ACK within returned data (a "piggybacked" ACK)

If the TCP ACK is not received for all the data, it is less likely to be a true application-level turn and is not counted as a turn. This means if a turn does not appear in the Discover appliance, data sent and received is likely to be overlapping.

The Protocols table displays the timing components for all application turns associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

Protocol

Auto-classified L7 protocol or a TCP/UDP port.

Turns

Number of TCP turns observed due to this protocol in the selected time period. Click the number of turns to display turn timing information over time for a specific protocol.

Network In (ms)

The time in milliseconds before the payload was received by the server. A large Network In value relative to the average application turn time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Processing Time (ms)

The time in milliseconds between the time the payload was received by the server and the time the payload was sent back. A large server processing time relative to the average application turn time indicates application delay.

Network Out (ms)

The time in milliseconds before the server finished sending the payload back. A large Network Out value relative to the average application turn time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

Breakdown

Click a value in the `Turns` column. If the **Response Time** drop-down list is set to **Breakdown**, the dialog box displays the overtime view of the following components:

Network In

The time in milliseconds before the payload was received by the server.

Processing Time

The time in milliseconds between the time the payload was received by the server and the time the payload was sent back.

Network Out

The time in milliseconds before the server finished sending the payload back.

Distribution

Click a value in the Turns column. If the **Response Time** drop-down list is set to **Distribution**, the dialog box displays the overtime view of the following components:

Network In

The time in milliseconds before the payload was received by the server.

Process

The time in milliseconds between the time the payload was received by the server and the time the payload was sent back.

Network Out

The time in milliseconds before the server finished sending the payload back.

Payload Size In

Displays the range of request sizes for all application turns associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values.

Payload Size Out

Displays the range of response sizes for all application turns associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values.

Turns

Number of TCP turns observed due to this protocol in the selected time period. Click the number of turns to display turn timing information over time for a specific protocol.

Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

L7 devices details page

The Protocols table lists all the protocols detected on this device and associated packet and byte counts. Click a protocol in the table to see the list of devices associated with that protocol.

To filter the list of protocols visible in the table, enter a search string in the Filter text box. The list filters automatically as search characters are entered.

L7 networks page

The L7 Protocols sub-page displays metrics for OSI Layer 7 traffic by packet count and throughput (total bytes). It also provides metrics on the top devices sending or receiving network traffic. The page includes the following information:

Packets by Protocol

Displays the packet rates for the top 10 protocols on the network.

Bytes by Protocol

Displays the throughput for the top 10 protocols on the network.

Protocols

Displays the devices sending and receiving traffic for the specified protocol.

L7 networks packets page

The Packets area chart displays how applications contribute to the total packet count on the network. In the chart, Date identifies the date and time for the data point on the graph that is currently being viewed. Packets displays the packet rate for the protocol at the given data point on the area chart, and the color block identifies the associated protocol name.

L7 networks throughput page

The Bytes by Application area chart displays how applications contribute to the total byte count on the network. In the chart, Date identifies the date and time for the data point on the graph that is currently being viewed. Bytes identifies the throughput for the data point that is currently being viewed in the area chart, and the color block identifies the associated protocol name.

L7 networks details page

The L7 Protocols Details page provides a complete list of protocols, and the packet and byte count for each.

L7 groups page

Protocol

The name of the protocol present in the group.

Packets In

The total incoming packet count for the protocol.

Packets Out

The total outgoing packet count for the protocol.

Bytes In

The total incoming byte count for the protocol.

Bytes Out

The total outgoing byte count for the protocol.

LDAP

ExtraHop appliances collect metrics about L7 activity.

LDAP applications page

LDAP Applications Tooblar

The LDAP application toolbar includes the following controls:

Errors

The chart shows the number of LDAP errors. Mouse over the chart to view a summary of a specific time or date. The table lists LDAP error messages and the number of times each occurred.

DNs

The chart shows the number of Distinguished Name (DN) messages transferred. The table displays the list of DN messages and the count associated with each DN message.

Users

The chart shows the number of requests from all users. Mouse over the chart to view a summary of a specific time or date. The table lists users and the request count associated with each user.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the **Request Bytes** counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

LDAP Metrics

Contains the following metrics:

Requests

The number of requests received.

Responses

The number of responses received.

Errors

The number of LDAP errors for the selected time interval.

Plain

The number of plain-text LDAP messages exchanged.

SASL

The number of encrypted LDAP messages exchanged.

Messages

Displays the LDAP messages for the selected time interval, such as BindRequest, BindResponse, UnbindRequest, SearchRequest, SearchResultDone and others. In the LDAP Server view, click the message counter to display clients that issued these messages. In the LDAP Client view, click the message counter to display servers that returned these messages.

Error Codes

Displays the LDAP errors for each LDAP error code within the selected time interval, such as invalidCredentials for LDAP error 49. Click the error counter to display devices that experienced these errors. For detailed error information, click **Errors**.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is

denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

LDAP devices page

LDAP Device Toolbar

The LDAP device toolbar includes the following controls:

LDAP Metric Type

Displays metrics for the current device acting as a LDAP client or LDAP server.

Errors

Displays a detailed list of error messages sent to or received by the current device over the specified time interval.

Servers

When acting as a LDAP client, displays a chart showing the total number of responses compared to processing time during the selected time interval.

Clients

When acting as a LDAP server, displays a chart showing the total number of requests compared to processing time during the selected time interval.

Records

Displays results for records that match the selected metric source and protocol.

LDAP Client

If you select **Client** for the LDAP Metric Type, the Discover appliance displays the following metrics. Click to display the list of servers from which responses were sent.

Requests

Specifies the number of LDAP requests for the selected time interval.

Responses

Specifies the number of responses that the device received when acting as an LDAP client.

Errors

Specifies the number of LDAP errors for the selected time interval.

Plain

Specifies the number of plain-text messages exchanged when the device is acting as an LDAP client.

SASL

Specifies the number of encrypted messages exchanged when the device is acting as an LDAP client.

LDAP Server

If you select **Server** for the LDAP Metric Type, the Discover appliance displays the following metrics. Click to display the list of servers from which responses were sent.

Requests

Specifies the number of requests that the device received when acting as an LDAP server.

Responses

Specifies the number of responses that the device sent when acting as an LDAP server.

Errors

Specifies the number of LDAP errors for the selected time interval.

Plain

Specifies the number of plain-text messages exchanged when the device is acting as an LDAP server.

SASL

Specifies the number of encrypted messages exchanged when the device is acting as an LDAP server.

Messages

Displays the LDAP messages for the selected time interval, such as BindRequest, BindResponse, UnbindRequest, SearchRequest, SearchResultDone and others. In the LDAP Server view, click the message counter to display clients that issued these messages. In the LDAP Client view, click the message counter to display servers that returned these messages.

Error Codes

Displays the LDAP errors for each LDAP error code within the selected time interval, such as invalidCredentials for LDAP error 49. Click the error counter to display devices that experienced these errors. For detailed error information, click **Errors**.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Server Processing Time

Displays the number of LDAP protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red data points to list the peer devices associated with the errors at this point in time. Click and drag across the chart to select a particular region.

Server Processing Time

Displays the median transaction time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing-time metrics. Click and drag across the chart to select a particular region.

Processing Time Distribution

Displays a histogram of times it took the server to process requests. Move the cursor over each bar to display the time range it represents and the number of requests in this bin.

LDAP devices timing page

The timing charts draw data from the **Time Interval** drop-down list on the navigation toolbar.

LDAP Metric Type

Select **Client** or **Server** to display statistics for the current device acting as an LDAP client or server, respectively.

Records

Displays results for records that match the selected metric source and protocol.

SearchResultDone

This chart displays the median transaction time in milliseconds for SearchResultDone LDAP messages as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the transaction-time metrics.

LDAP groups page

LDAP Groups Toolbar

The LDAP groups toolbar includes the following controls:

LDAP Metric Type

Displays metrics for members in the current group acting as a LDAP client or server, respectively.

Errors

Displays a detailed list of error messages sent to or received by members in the current group over the specified time interval.

Records

Displays results for records that match the selected metric source and protocol.

LDAP Client

If you select **Client** for the LDAP Metric Type, the Discover appliance displays the following metrics. Click the counter to break down the responses by group members in the table at the bottom of the page.

Requests

Specifies the number of LDAP requests for the selected time interval.

Responses

Specifies the number of LDAP responses for the selected time interval.

Errors

Specifies the number of LDAP errors for the selected time interval.

Plain

Specifies the number of LDAP plain-text messages for the selected time interval.

SASL

Specifies the number of LDAP encrypted messages for the selected time interval.

LDAP Server

If you select **Server** for the LDAP Metric Type, the Discover appliance displays the following metrics. Click the counter to break down the responses by group members in the table at the bottom of the page.

Requests

Specifies the number of LDAP requests for the selected time interval.

Responses

Specifies the number of LDAP responses for the selected time interval.

Errors

Specifies the number of LDAP errors for the selected time interval.

Plain

Specifies the number of LDAP plain-text messages for the selected time interval.

SASL

Specifies the number of LDAP encrypted messages for the selected time interval.

Messages

Displays the LDAP messages for the selected time interval, such as `BindRequest`, `BindResponse`, `UnbindRequest`, `SearchRequest`, `SearchResultDone` and others. Click to break down the message by group member in the table at the bottom of the page.

Error Codes

Displays the LDAP errors for each LDAP error code within the selected time interval, such as `invalidCredentials` for LDAP error 49. Click to break down the message by group member in the table at the bottom of the page.

LDAP groups processing time page

LDAP Metric Type

Displays metrics for members in the current group acting as a LDAP client or server, respectively.

Records

Displays results for records that match the selected metric source and protocol.

Server Processing Time

Displays median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

Memcache

ExtraHop appliances collect metrics about Memcache activity.

Memcache applications page

Memcache Application Toolbar

The Memcache application toolbar includes the following controls:

Errors

The chart shows the number of Memcache errors. Mouse over the chart to view a summary of a specific time or date. The table lists Memcache error messages and the number of times each occurred.

Hits

The chart shows the total count for Memcache hits (values returned from the server to the client in response to "get" requests). Mouse over the chart to view a summary of a specific time or date. The table lists Memcache keys and the total count associated with each.

Misses

The chart shows the total count for Memcache misses ("get" requests for which the specified key was not found). Mouse over the chart to view a summary of a specific time or date. The table lists Memcache keys and the total count associated with each.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the **Request Bytes** counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

Memcache Metrics

Contains the following metrics:

Requests

The number of Memcache requests.

No-Replies

The number of Memcache requests for which a response was not necessarily expected, and none was received.

Responses

The number of Memcache responses.

Hits

The number of items matched and returned in response to Memcache GET requests.

Misses

The number of items requested but not received in response to Memcache GET requests. Misses are counted even if the server did not explicitly inform the client of the miss (for example, if the GET was a quiet request).

Errors

The number of errors sent by the Memcache server in response to client requests. Some responses other than the default response are not considered errors because they are usually expected to occur during normal operation. For example, the `NOT_FOUND` status code is not considered an error. In the Memcache text protocol analysis, only `ERROR`, `CLIENT_ERROR`, and `SERVER_ERROR` responses are considered errors.

Methods

Displays the Memcache methods for the selected time interval.

Status Codes

The status code section displays the HTTP status codes for the selected time interval. Click the number next to each status code to display a list of URIs associated with each status code.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Cache Hits and Misses

Displays the number of hits and misses as a function of time over the selected time interval.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

Memcache devices page

Memcache Device Toolbar

The Memcache device toolbar includes the following controls:

Memcache Metric Type

Displays metrics for the current device acting as a Memcache client or server.

Errors

Displays the list of error messages sent to or received by the current device over the selected time interval.

IP Address Memcache Metrics

Click the counters next to individual Memcache metrics to show the IP Address Memcache Metrics for Memcache peer devices. For Memcache servers, the peer devices are Memcache clients. For Memcache clients, the peer devices are Memcache servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS hostname of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

Memcache Server

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests that the device received when acting as a Memcache server.

No-Replies

Specifies the number of requests sent for which a response was not necessarily expected, and none was received when the device is acting as a Memcache server.

In the Memcache text protocol, when a client sends a request with the "noreply" keyword, the server performs the requested action but never sends a reply, and the Discover appliance records a no-reply.

In the Memcache binary protocol, some kinds of requests are known as "quiet" requests. One example is the "get quietly" (getq) command: if the specified key is found, the server sends a response containing the corresponding value and the Discover appliance records a response; otherwise, the server sends nothing and the Discover appliance records a no-reply.

If the server is responding and the Discover appliance is receiving a high-quality data feed, the number of requests should equal the number of responses plus the number of no-replies.

Responses

Specifies the number of responses that the device sent when acting as a Memcache server.

Hits

Specifies the number of items matched and that the device sent in response to "get" commands when acting as a Memcache server.

Misses

Specifies the number of items requested but not sent in response to get commands when the device is acting as a Memcache server. Misses are counted even if the server did not explicitly inform the client of the miss (for example, if the get was a quiet request).

Errors

Specifies the number of errors sent by the Memcache server in response to client requests. Some responses other than the default response are not considered errors because they are usually expected to occur during normal operation. For example, the `NOT_FOUND` reply code is not considered an error. In the Memcache text protocol analysis, only `ERROR`, `CLIENT_ERROR`, and `SERVER_ERROR` responses are considered errors.

Memcache

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests that the device sent when acting as a Memcache client.

No-Replies

Specifies the number of requests sent for which a response was not necessarily expected, and none was received when the device is acting as a Memcache client.

In the Memcache text protocol, when a client sends a request with the "noreply" keyword, the server performs the requested action but never sends a reply, and the Discover appliance records a no-reply.

In the Memcache binary protocol, some kinds of requests are known as "quiet" requests. One example is the "get quietly" (getq) command: if the specified key is found, the server sends a response containing the corresponding value and the Discover appliance records a response; otherwise, the server sends nothing and the Discover appliance records a no-reply.

If the server is responding and the Discover appliance is receiving a high-quality data feed, the number of requests should equal the number of responses plus the number of no-replies.

Responses

Specifies the number of responses that the device received when acting as a Memcache client.

Hits

Specifies the number of items matched and that the device received in response to "get" commands when acting as a Memcache client.

Misses

Specifies the number of items requested but not received in response to get commands when the device is acting as a Memcache client. Misses are counted even if the server did not explicitly inform the client of the miss (for example, if the get was a quiet request).

Errors

Specifies the number of errors sent by the Memcache server in response to client requests. Some responses other than the default response are not considered errors because they are usually expected to occur during normal operation. For example, the `NOT_FOUND` reply code is not considered an error. In the Memcache text protocol analysis, only `ERROR`, `CLIENT_ERROR`, and `SERVER_ERROR` responses are considered errors.

Commands

Breakdown of individual commands, organized into meaningful groups. For example, in this area the reply code "get" represents the get and gets commands in the Memcache text protocol and the `get`, `getq`, `getk`, and `getkq` commands in the Memcache binary protocol. In the Memcache text protocol analysis, if a single get or gets command includes multiple keys, a "get" is counted for each of those keys.

Reply Codes

Breakdown of reply codes. In the Memcache binary protocol analysis, for reply codes other than the default `NO_ERROR`, the command that produced the reply is also provided here as part of the reply code.

`EH_DIAG_*` reply codes are present in early access deployments of Memcache to allow ExtraHop engineers to gather additional information about the performance of the module.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Value Sizes

Displays statistical summaries of size distributions for the following values:

Stored

Displays the range of stored value sizes for all transactions associated with the current device. Mouse-over to display the five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean stored value size for each peer device.

Retrieved

Displays the range of retrieved value sizes for all transactions associated with the current device. Mouse-over to display the five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean retrieved value size for each peer device.

Cache Hits and Misses

Displays the number of hits and misses as a function of time over the selected time interval.

Key Access Time Breakdown

Shows key access time as a function of time over the selected time interval. Key access time is the time from the last byte of the client's "get" command to the first byte of the corresponding value returned from the server. Therefore, this metric is recorded only for cache hits. The line chart displays the median and first/third quartiles of the key access time over time.

Key Access Time Distribution

Shows a histogram distribution of key access time over the selected time period. Key access time is the time from the last byte of the client's "get" command to the first byte of the corresponding value returned from the server. Therefore, this metric is recorded only for cache hits.

Memcache groups page

Memcache Groups Toolbar

The Memcache groups toolbar includes the following controls:

Memcache Metric Type

Displays metrics for members in the current group acting as a Memcache client or server, respectively.

Errors

Displays the list of error messages sent to or received by members in the current group over the selected time interval.

Memcache Server

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of Memcache requests.

No-Replies

Specifies the number of requests for which a response was not necessarily expected, and none was received.

In the Memcache text protocol, when a client sends a request with the "noreply" keyword, the server performs the requested action but never sends a reply, and the Discover appliance records a no-reply.

In the Memcache binary protocol, some kinds of requests are known as "quiet" requests. One example is the "get quietly" (getq) command: if the specified key is found, the server sends a response containing the corresponding value and the Discover appliance records a response; otherwise, the server sends nothing and the Discover appliance records a no-reply.

If the server is responding and the Discover appliance is receiving a high-quality data feed, the number of requests should equal the number of responses plus the number of no-replies.

Responses

Specifies the number of Memcache responses.

Hits

Specifies the number of values returned from the server to the client in response to "get" commands.

Misses

Specifies the number of "get" commands for which the specified key was not found. Misses are counted even if the server did not explicitly inform the client of the miss (for example, if the get was a "quiet" request).

Errors

Specifies the number of errors sent by the Memcache server in response to client requests. Some responses other than the default response are not considered errors because they are usually expected to occur during normal operation. For example, the `NOT_FOUND` reply code is not considered an error. In the Memcache text protocol analysis, only `ERROR`, `CLIENT_ERROR`, and `SERVER_ERROR` responses are considered errors.

Memcache

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of Memcache requests.

No-Replies

Specifies the number of requests for which a response was not necessarily expected, and none was received.

In the Memcache text protocol, when a client sends a request with the "noreply" keyword, the server performs the requested action but never sends a reply, and the Discover appliance records a no-reply.

In the Memcache binary protocol, some kinds of requests are known as "quiet" requests. One example is the "get quietly" (getq) command: if the specified key is found, the server sends a response containing the corresponding value and the Discover appliance records a response; otherwise, the server sends nothing and the Discover appliance records a no-reply.

If the server is responding and the Discover appliance is receiving a high-quality data feed, the number of requests should equal the number of responses plus the number of no-replies.

Responses

Specifies the number of Memcache responses.

Hits

Specifies the number of values returned from the server to the client in response to "get" commands.

Misses

Specifies the number of "get" commands for which the specified key was not found. Misses are counted even if the server did not explicitly inform the client of the miss (for example, if the get was a "quiet" request).

Errors

Specifies the number of errors sent by the Memcache server in response to client requests. Some responses other than the default response are not considered errors because they are usually expected to occur during normal operation. For example, the `NOT_FOUND` reply code is not considered an error. In the Memcache text protocol analysis, only `ERROR`, `CLIENT_ERROR`, and `SERVER_ERROR` responses are considered errors.

Commands

Breakdown of individual commands, organized into meaningful groups. For example, in this area the reply code "get" represents the get and gets commands in the Memcache text protocol and the `get`, `getq`, `getk`, and `getkq` commands in the Memcache binary protocol. In the Memcache text protocol analysis, if a single get or gets command includes multiple keys, a "get" is counted for each of those keys. Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Reply Codes

Breakdown of reply codes. In the Memcache binary protocol analysis, for reply codes other than the default `NO_ERROR`, the command that produced the reply is also provided here as part of the reply code.

`EH_DIAG_*` reply codes are present in early access deployments of Memcache to allow ExtraHop engineers to gather additional information about the performance of the module.

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

MongoDB

ExtraHop appliances collect metrics about MongoDB activity.

MongoDB applications page

MongoDB Application Toolbar

The MongoDB application toolbar includes the following controls:

Errors

The chart shows the number of MongoDB errors. Mouse over the chart to view a summary of a specific time or date. The table lists MongoDB error messages and the number of times each occurred.

Methods

The chart shows the total count compared to the mean time (ms). Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists methods, count, total time, and mean time (ms) associated with each method. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Users

The chart shows the number of responses and errors from all users. Mouse over the chart to view a summary of a specific time or date. The table lists users and the number of responses and errors associated with each user.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Database

Displays application metrics by database.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted from the application within the selected time interval. Selecting **By Client IP** in the drop-down list while mousing over the **Request Bytes** counter shows which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

MongoDB Metrics

Contains the following metrics:

Requests

The number of MongoDB requests.

Responses

The number of MongoDB responses.

Errors

The number of errors sent or received within the selected time interval.

Methods

Displays the methods MongoDB uses to authenticate clients. Click the counter to display additional per-client or per-server IP address details.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a

large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

MongoDB devices page

MongoDB Device Toolbar

The MongoDB device toolbar includes the following controls:

MongoDB Metric Type

Displays metrics for the current device acting as a MongoDB client or MongoDB server.

Errors

Displays the list of error messages sent to or received by the current device over the selected time interval.

Methods

Displays the list of methods and associated bytes sent and received by the current device for the selected time interval. Methods are broken out by key parameters, such as the accessed file name.

Users

Displays the list of users accessing the MongoDB server and associated bytes sent and received for the selected time interval.

MongoDB Details specifies the type of additional MongoDB information displayed. Moving the cursor over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays MongoDB metrics by IP address.

By Database

Displays MongoDB metrics by database.

For example, MongoDB Requests is a top-level metric showing how many requests were received by the MongoDB server during the selected time frame. Selecting By IP in the drop-down list while moving the cursor over the `HTTP Requests` counter shows which IP addresses originated these requests. Selecting By Database from the drop-down list while moving the cursor over the HTTP Requests counter shows which databases were accessed by the requestors.

IP Address MongoDB Metrics

Click the counters next to individual MongoDB metrics to show the IP Address MongoDB Metrics for MongoDB peer devices. For MongoDB servers, the peer devices are MongoDB clients. For MongoDB clients, the peer devices are MongoDB servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

MongoDB Server

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests that the device received when acting as a MongoDB server.

Responses

Specifies the number of responses that the device sent when acting as a MongoDB server.

Errors

Specifies the number of errors sent by the MongoDB server.

Requests Aborted

Specifies the number of requests that the device began to receive but did not receive completely when acting as a MongoDB server.

Responses Aborted

Specifies the number of responses that the device began to send but did not send completely when acting as a MongoDB server.

MongoDB Client

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests that the device sent when acting as a MongoDB client.

Responses

Specifies the number of responses that the device received when acting as a MongoDB client.

Errors

Specifies the number of errors sent by the MongoDB client.

Requests Aborted

Specifies the number of requests that the device began to send but did not send completely when acting as a MongoDB client.

Responses Aborted

Specifies the number of responses that the device began to receive but did not receive completely when acting as a MongoDB client.

Methods

Displays the methods MongoDB uses to authenticate clients. Click the counter to display additional per-client or per-server IP address details.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Request Size

Displays the range of request sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Response Size

Displays the range of response sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

MongoDB devices timing page

Request Transfer Time

Displays a histogram of times it took to transfer requests from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Processing Time

Displays a histogram of times it took the server to process requests. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Response Transfer Time

Displays a histogram of times it took to transfer the response from the server to the client. Mouse over each bar to display the time range it represents and the number of requests in this bin.

MongoDB groups page

MongoDB Groups Toolbar

The MongoDB groups toolbar includes the following controls:

MongoDB Metric Type

Displays metrics for the current group acting as a MongoDB client or MongoDB server, respectively.

Errors

Displays the list of error messages sent to or received by the current group over the selected time interval.

Methods

Displays the list of methods and associated bytes sent and received by the current group for the selected time interval. Methods are broken out by key parameters, such as the accessed file name.

Users

Displays the list of users accessing the MongoDB server and associated bytes sent and received for the selected time interval.

MongoDB Server

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests sent by the MongoDB server.

Responses

Specifies the number of responses sent by the MongoDB server.

Errors

Specifies the number of errors sent by the MongoDB server.

Requests Aborted

Specifies the number of incomplete MongoDB requests sent from the current member.

Responses Aborted

Specifies the number of incomplete MongoDB responses received by the current member.

MongoDB Client

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests sent by the MongoDB server.

Responses

Specifies the number of responses sent by the MongoDB server.

Errors

Specifies the number of errors sent by the MongoDB server.

Requests Aborted

Specifies the number of incomplete MongoDB requests sent from the current member.

Responses Aborted

Specifies the number of incomplete MongoDB responses received by the current member.

Methods

Displays the methods MongoDB uses to authenticate clients. Click the counter to display additional per-client or per-server IP address details.

MSRPC

ExtraHop appliances collect metrics about Microsoft Remote Procedure Call (MSRPC) activity.

MSRPC devices page

MSRPC Device Toolbar

The MSRPC device toolbar includes the following controls:

RPC Metric Type

Display metrics for devices acting as an MSRPC client or server.

Packets

The Packets line chart displays the incoming and outgoing RPC packet rate (packets per second) over the selected time interval. Current and Max are the current and maximum packet rates. Total is the total number of packets over the selected time interval.

Throughput

The Throughput line chart displays the incoming and outgoing RPC throughput (bits per second) over the selected time interval. Current and Max are the current and maximum throughputs. Total is the total number of bytes transferred over the selected time interval.

RPC Metrics

Click the RPC Metrics section to display the list of RPC-specific metrics for the selected time interval. Click the counter to display additional per-client or per-server IP address details.

Examples of RPC-specific metrics include:

Rejected Binds (naks)

Number of binds rejected by the server. Rejected binds occur when a server sends and receives bind updates from a peer server out of order.

Failed EPM Binds

Number of failed End-Point Mapper binds.

Orphaned

Number of times when client aborts request in progress.

Faults

Number of "fault" PDUs returned.

Canceled Operation

Number of canceled operations.

Fragmented Responses

Number of fragmented responses.

Authentication Types

The Authentication Types section displays the authentication methods and protection levels in format "AuthMethod - ProtLevel". Click the counter to display additional per-client or per-server IP address details.

Protection levels include:

None

No protection level

Connect

Authenticates only when the client establishes a relationship with a server.

Call

Authenticates only at the beginning of each remote procedure call when the server receives the request.

Packet

Authenticates only that all data received is from the expected client. Does not validate the data itself.

Packet Integrity

Authenticates and verifies that none of the data transferred between client and server has been modified.

Packet Integrity

Includes all previous levels, and ensures clear text data can only be seen by the sender and the receiver.

MSRPC devices interfaces page

MSRPC device toolbar

The MSRPC device toolbar includes the following controls:

RPC Metric Type

Display metrics for devices acting as an MSRPC client or server.

Packets In by Interface

The Packets In by Interface area chart displays the incoming RPC packet rate (packets per second) over the selected time interval.

Packets Out by Interface

The Packets Out by Interface area chart displays the outgoing RPC packet rate (packets per second) over the selected time interval.

Bytes In by Interface

The Bytes In by Interface area chart displays the incoming RPC throughput (bits per second) over the selected time interval.

Bytes Out by Interface

The Bytes Out by Interface area chart displays the outgoing RPC throughput (bits per second) over the selected time interval.

RPC interfaces include:

- AD Setup
- Netlogon
- Exchange MAPIAD Replication
- AD Backup
- LSA (Local Security Authority)
- DCOM RIUnknown
- DCOM RIUnknown2
- System.Activator
- ID Resolver
- SAMR (Security Account Manager Remote)
- SCM (Service Control Manager)
- Srvsvc
- Remote Registry
- Exchange NSPI (Name Service Provider Interface)
- AD XDS (Active Directory Exchange Directory Service)
- EPM (Endpoint Mapper)
- Exchange RFR (Exchange Referral)
- File Replication
- File Replication Ex
- File Replication v1.0
- RPC Management

MSRPC devices interfaces packets page

MSRPC Device Toolbar

The MSRPC device toolbar includes the following controls:

RPC Metric Type

Display metrics for devices acting as an MSRPC client or server.

Packets In

Displays the incoming RPC packet rate (in packets per second) for the selected device over the given time interval.

Packets Out

Displays the incoming RPC packet rate (in packets per second) for the selected device over the given time interval.

MSRPC devices interfaces throughput page

MSRPC Device Toolbar

The MSRPC device toolbar includes the following controls:

RPC Metric Type

Display metrics for devices acting as an MSRPC client or server.

Bytes In

Displays the incoming RPC throughput (in bits per second) for the selected device over the given time interval.

Bytes Out

Displays the outgoing RPC throughput (in bits per second) for the selected device over the given time interval.

Multicast

ExtraHop appliances collect metrics about Multicast activity.

Multicast devices page

The Multicast sub-page for a device displays metrics for multicast and broadcast traffic on the network.

Well-known multicast groups include:

- IEEE Spanning Tree (STP)
- Address Resolution Protocol (ARP)
- IPv6 Neighbor Discovery Protocol (NDP)
- Cisco Discovery Protocol (CDP)
- Cisco Shared Spanning Tree Protocol (CSSTP)
- Alternate Spanning Multicast (ALTSM)
- Router Information Protocol (RIP)
- Network Time Protocol (NTP)
- OSPFMPLSInter Switch Link (ISL)
- Cisco VLAN Bridge (CVB)
- DHCP client (DHCP_CLIENT)
- DHCP server (DHCP_SERVER)
- NETBIOS Name Service (NETBIOS_NS)
- NETBIOS Datagram Service (NETBIOS_DGM)
- Multicast DNS (MDNS)
- Hot Standby Router Protocol (HSRP)
- Uncategorized L2 broadcast (L2BCAST)

Other multicast groups are represented using the numeric form of the group address, protocol, and L4 port.

Packet Count by Group

The Packet Count by Group bar chart displays the packet count for each of the top-ten multicast groups in which the selected device participates.

Byte Count by Group

The Byte Count by Group bar chart displays the byte count for each of the top-ten multicast groups in which the selected device participates.

Multicast Groups

The Multicast Groups table displays the multicast group, packet group, and byte count for the selected device.

Multicast networks page

The Multicast sub-page displays metrics for multicast and broadcast traffic on the network. Well-known multicast groups include:

Well-known multicast groups include:

- IEEE Spanning Tree (STP)
- Address Resolution Protocol (ARP)
- IPv6 Neighbor Discovery Protocol (NDP)
- Cisco Discovery Protocol (CDP)
- Cisco Shared Spanning Tree Protocol (CSSTP)
- Alternate Spanning Multicast (ALTSM)
- Router Information Protocol (RIP)
- Network Time Protocol (NTP)
- OSPFMPLSInter Switch Link (ISL)
- Cisco VLAN Bridge (CVB)
- DHCP client (DHCP_CLIENT)
- DHCP server (DHCP_SERVER)
- NETBIOS Name Service (NETBIOS_NS)
- NETBIOS Datagram Service (NETBIOS_DGM)
- Multicast DNS (MDNS)
- Hot Standby Router Protocol (HSRP)
- Uncategorized L2 broadcast (L2BCAST)

Other multicast groups are represented using the numeric form of the group address, protocol, and L4 port.

Packet Count by Groups

Displays the packet count for each of the top-ten multicast groups.

Byte Count by Groups

Displays the byte count for each of the top-ten multicast groups.

Multicast Groups

Displays the multicast group, packet group, and byte count for the selected device.

Multicast networks details page

Multicast

Lists all multicast groups detected on the network and associated packet and byte counts. Click a multicast group to view the devices sending or receiving the traffic for that multicast group.

Multicast networks top groups page

Top Groups (Packets)

Displays how multicast groups contribute to the total packet count on the network. Click a multicast group listed in the legend to list the devices sending or receiving the traffic for that protocol in the Multicast table below.

Top Groups (Bytes)

Displays how multicast groups contribute to the total byte count on the network. Click a multicast group listed in the legend to list the devices sending or receiving the traffic for that protocol in the Multicast table below.

Multicast

Displays the devices sending or receiving the traffic for the selected protocols.

NFS

ExtraHop appliances collect metrics about Network File System (NFS) activity.

NFS devices page

Note: Where file name detail is presented, the Discover appliance displays both the file path and mount point, if available. The prefix '`...`' indicates that either the mount point or part of the path is not available. This may occur in instances when the capture process was restarted after the "`mount`" or a "`cd`" command was issued, or when the commands were lost due to desyncs.

NFS Device Toolbar

The NFS device toolbar includes the following controls:

NFS Metric Type

Displays metrics for the current device acting as a NFS client or NFS server.

Errors

Displays the list of error messages sent to or received by the current device over the selected time interval.

Methods

Displays the list of methods and associated bytes sent and received by the current device for the selected time interval. Methods are broken out by key parameters, such as the accessed file name.

Users

Displays the list of users accessing the NFS server and associated bytes sent and received for the selected time interval.

Files

Displays the list of files accessed and associated bytes sent and received for the selected time interval. The access time indicates the time to access a file on an NFS partition and is measured by timing non-pipelined commands for every `READ` and `WRITE`.

IP Address NFS Metrics

Click the counters next to individual NFS metrics to show the IP Address NFS metrics for NFS peer devices. For NFS servers, the peer devices are NFS clients. For NFS clients, the peer devices are NFS servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

<Metric value>

Displays the value for the selected metric.

NFS Server

Click the counter next to each metric to display additional IP address details.

Responses

Specifies the number of responses that the device sent when acting as an NFS server.

Errors

Specifies the number of errors sent by the NFS server.

Retransmissions

Specifies the number of NFS requests for which the retransmission timer expired and the request was retried when the device is acting as an NFS server.

Reads

Specifies the number of NFS read requests that the device received when acting as an NFS server.

Writes

Specifies the number of NFS write requests that the device received when acting as an NFS server.

TCP

Specifies the number of NFS requests that the device made over TCP when acting as an NFS server. All versions of NFS can use TCP, and NFSv4 requires it.

UDP

Specifies the number of NFS requests that the device made over UDP when acting as an NFS server. NFSv2 and NFSv3 can use the User Datagram Protocol (UDP) to provide a stateless network connection between the client and server.

Aborts

Specifies the number of incomplete requests that the device received when acting as an NFS server.

NFS Client

Click the counter next to each metric to display additional IP address details.

Responses

Specifies the number of responses that the device received when acting as an NFS client.

Errors

Specifies the number of errors sent by the NFS client.

Retransmissions

Specifies the number of NFS requests for which the retransmission timer expired and the request was retried when the device is acting as an NFS client.

Reads

Specifies the number of NFS read requests that the device sent when acting as an NFS client.

Writes

Specifies the number of NFS write requests that the device sent when acting as an NFS client.

TCP

Specifies the number of NFS requests that the device made over TCP when acting as an NFS client. All versions of NFS can use TCP, and NFSv4 requires it.

UDP

Specifies the number of NFS requests made over UDP when the device is acting as an NFS client. NFSv2 and NFSv3 can use the User Datagram Protocol (UDP) to provide a stateless network connection between the client and server.

Aborts

Specifies the number of incomplete requests that the device sent when acting as an NFS client.

Authentication Methods

Displays the methods NFS uses to authenticate clients. Click the counter to display additional per-client or per-server IP address details.

Versions

Displays the versions of NFS traffic being processed over the selected time interval.

Status Codes

Displays the list of status codes sent to or received by the current device over the selected time interval.

Methods

Displays the NFS methods for the selected time interval.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Request Size

Displays the range of request sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Response Size

Displays the range of response sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Read, Write, and FSInfo Bytes

Displays the total bytes per device transmitted within the selected time interval. Mouse over the graph to see the byte count for each metric at a specific moment in time.

NFS devices timing page

Request Transfer Time

Displays a histogram of times it took to transfer requests from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Access Time

Displays a histogram of file access times. Move the mouse pointer over each bar to display the time range it represents and the number of requests in this bin.

Response Transfer Time

Displays a histogram of times it took to transfer responses from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

NFS groups page

Note: Where file name detail is presented, the Discover appliance displays both the file path and mount point, if available. The prefix '`...`' indicates that either the mount point or part of the path is not available. This may occur in instances when the capture process was restarted after the "`mount`" or a "`cd`" command was issued, or when the commands were lost due to desyncs.

NFS Groups Toolbar

The NFS groups toolbar includes the following controls:

NFS Metric Type

Displays metrics for members in the current group acting as a NFS client or server, respectively.

Errors

Displays the list of error messages sent to or received by members in the current group over the selected time interval.

Methods

Displays the list of methods and associated bytes sent and received by members in the current group during the selected time interval. Methods are broken out by key parameters, such as the accessed file name.

Users

Displays the list of users accessing the NFS servers in this group and associated bytes sent and received for the selected time interval.

Files

Displays the list of files accessed and associated bytes sent and received for the selected time interval.

Records

Displays results for records that match the selected metric source and protocol.

NFS Server

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Responses
Specifies the number of responses sent by the NFS server.

Errors
Specifies the number of errors sent by the NFS server.

Retransmissions
Specifies the number of NFS request retransmissions.

Reads
Specifies the number of read operations requested from the NFS server.

Writes
Specifies the number of write operations requested from the NFS server.

TCP
Specifies the number of NFS connections performed over TCP.

UDP
Specifies the number of NFS connections performed over UDP.

Aborts
Specifies the number of NFS operations aborted from this server.

NFS Client
Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Responses
Specifies the number of responses received by the NFS client.

Errors
Specifies the number of errors sent by the NFS client.

Retransmissions
Specifies the number of NFS request retransmissions.

Reads
Specifies the number of read operations requested by the NFS client.

Writes
Specifies the number of write operations requested by the NFS client.

TCP
Specifies the number of NFS connections performed over TCP.

UDP
Specifies the number of NFS connections performed over UDP.

Aborts
Specifies the number of NFS operations aborted to this client.

Authentication Methods
Displays the methods NFS uses to authenticate clients.

Versions
Displays the versions of NFS traffic being processed over the selected time interval.

Status Codes
Displays the list of status codes sent to or received by the current group over the selected time interval.

Methods
Displays the NFS methods for the selected time interval.

NFS groups processing time pages

NFS Metric Type

Displays metrics for members in the current group acting as a NFS client or server, respectively.

Records

Displays results for records that match the selected metric source and protocol.

Server Access Time

Shows median server access time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

PCoIP

ExtraHop appliances collect metrics about PC over IP (PCoIP) activity.

PCoIP devices page

PCoIP Device Toolbar

The PCoIPdevice toolbar includes the following controls:

Metric Type

Displays metrics for devices acting as a PCoIP client or PCoIPserver.

Messages In

Inbound PCoIP messages.

Audio

Number of audio messages received in the selected time interval.

Other

Number of other messages received in the selected time interval.

USB

Number of USB messages received in the selected time interval.

Video

Number of video messages received in the selected time interval.

Messages Out

Outbound PCoIP messages.

Audio

Number of audio messages sent in the selected time interval.

Other

Number of other messages sent in the selected time interval.

USB

Number of USB messages sent in the selected time interval.

Video

Number of video messages sent in the selected time interval.

Launches

Displays the number of launches for the selected time interval.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Bytes In by Channel

Displays the breakdown of incoming throughput by virtual channel.

Bytes Out by Channel

Displays the breakdown of outgoing throughput by virtual channel.

PCoIP groups page

PCoIP Groups Toolbar

The PCoIP groups toolbar includes the following controls:

Metric Type

Click the **Metric Type** drop-down list and select either **Client** or **Server** to display metrics for members in the current group acting as an PCoIP client or PCoIPserver, respectively.

Messages In

Inbound PCoIP messages.

Audio

Number of audio messages received in the selected time interval.

Other

Number of other messages received in the selected time interval.

USB

Number of USB messages received in the selected time interval.

Video

Number of video messages received in the selected time interval.

Messages Out

Outbound PCoIP messages.

Audio

Number of audio messages sent in the selected time interval.

Other

Number of other messages sent in the selected time interval.

USB

Number of USB messages sent in the selected time interval.

Video

Number of video messages sent in the selected time interval.

RTCP

ExtraHop appliances collect metrics about Real-Time Transport Control Protocol (RTCP) activity.

RTCP applications page

Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Sender IP

Displays application metrics by the sender IP addresses.

By Receiver IP

Displays application metrics by the receiver IP addresses.

By Canonical Name

Displays device metrics by canonical name.

The RTCP sub-page provides RTCP information about an application.

L2-L4 Metrics

Contains the following metrics:

L2 Bytes

The number of L2 bytes associated with RTCP transactions.

Packets

The number of packets associated with RTCP transactions.

RTCP Messages

Contains the following metrics:

Sender Report Messages

The number of packets transmitted by the sender from the beginning of the transmission to the time this sender report packet was generated.

Sender Report Drops

The number of packets that were lost by the sender since the beginning of reception.

Receiver Reports Messages

The number of packets transmitted by the receiver from the beginning of the transmission to the time this receiver report packet was generated.

Receiver Report Drops

The number of packets that were lost by the receiver since the beginning of reception.

Message Types

The number of RTCP records broken down by message type.

Packets Lost

The number of packets lost broken down by sender and receiver.

Sender Report Jitter

An estimate of the statistical variance of the RTP packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

Receiver Report Jitter

An estimate of the statistical variance of the RTP packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

RTCP devices page

Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays device metrics by IP address.

By Canonical Name

Displays device metrics by canonical name.

For device metrics, the RTCP page includes the following data:

RTCP In

Contains the following metrics:

Sender Report Messages

The number of incoming packets transmitted by the sender from the beginning of the transmission to the time this sender report packet was generated.

Sender Report Drops

The number of incoming packets that were lost by the sender since the beginning of reception.

Receiver Report Messages

The number of incoming packets transmitted by the receiver from the beginning of the transmission to the time this receiver report packet was generated.

Receiver Report Drops

The number of incoming packets that were lost by the receiver since the beginning of reception.

RTCP Out

Contains the following metrics:

Sender Report Messages

The number of outgoing packets transmitted by the sender from the beginning of the transmission to the time this sender report packet was generated.

Sender Report Drops

The number of outgoing packets that were lost by the sender since the beginning of reception.

Receiver Report Messages

The number of outgoing packets transmitted by the receiver from the beginning of the transmission to the time this receiver report packet was generated.

Receiver Report Drops

The number of outgoing packets that were lost by the receiver since the beginning of reception.

Message Types In

The number of incoming RTCP records broken down by message type.

Message Types Out

The number of outgoing RTCP records broken down by message type.

Packets Lost

The number of packets that were lost since the beginning of reception.

Sender Report Jitter In

An estimate of the statistical variance of the incoming packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

Sender Report Jitter Out

An estimate of the statistical variance of the outgoing packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

Receiver Report Jitter In

An estimate of the statistical variance of the incoming packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

Receiver Report Jitter Out

An estimate of the statistical variance of the outgoing packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

RTCP groups page

RTCP In

Contains the following metrics:

Sender Report Messages

The number of incoming packets transmitted by the sender from the beginning of the transmission to the time this sender report packet was generated.

Sender Report Drops

The number of incoming packets that were lost by the sender since the beginning of reception.

Receiver Report Messages

The number of incoming packets transmitted by the receiver from the beginning of the transmission to the time this receiver report packet was generated.

Receiver Report Drops

The number of incoming packets that were lost by the receiver since the beginning of reception.

RTCP Out

Contains the following metrics:

Sender Report Messages

The number of outgoing packets transmitted by the sender from the beginning of the transmission to the time this sender report packet was generated.

Sender Report Drops

The number of outgoing packets that were lost by the sender since the beginning of reception.

Receiver Report Messages

The number of outgoing packets transmitted by the receiver from the beginning of the transmission to the time this receiver report packet was generated.

Receiver Report Drops

The number of outgoing packets that were lost by the receiver since the beginning of reception.

Message Types In

The number of incoming RTCP records broken down by message type.

Message Types Out

The number of outgoing RTCP records broken down by message type.

RTP

ExtraHop appliances collect metrics about Real-Time Transport Protocol (RTP) activity.

RTP applications page

Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Sender IP

Displays application metrics by the sender IP addresses.

By Receiver IP

Displays application metrics by the receiver IP addresses.

By Codec

Displays application metrics by codec.

The RTP sub-page provides the following RTP information about an application.

L2-L4 Metrics

Contains the following metrics:

L2 Bytes

The number of L2 bytes associated with RTP transactions.

Packets

The number of packets associated with RTP transactions.

RTP Metrics

Contains the following metrics:

Messages

The number of messages associated with RTP transmissions.

Drops

The number of packets associated with RTP transmissions which were lost in transit.

Duplicates

The number of duplicate messages associated with RTP transmissions.

Out of Order

The number of packets associated with RTP transmissions where the sequence number did not match the sequence number that the Discover appliance was expecting. The reordering may have been introduced at the point of origin or an intermediary. This may result in decreased call quality.

RTP Messages by Codec

The number of RTP messages broken down by codec.

Throughput

The throughput (in bits per second) over the selected time interval.

Message Metrics

The number of drops, duplicates, and out of order messages associated with RTP transmissions over the selected time interval.

Jitter

An estimate of the statistical variance of the RTP packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

MOS

The mean opinion score calculated for packets associated with RTP transmissions.

RTP devices page

Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays device metrics by IP address.

By Codec

Displays device metrics by codec.

For device metrics, the RTP page includes the following data:

RTP In

Contains the following metrics:

Messages

The number of incoming messages associated with RTP transmissions.

Drops

The number of incoming packets associated with RTP transmissions which were lost in transit.

Duplicates

The number of incoming duplicate messages associated with RTP transmissions.

Out of Order

Number of incoming packets associated with RTP transmissions where the sequence number did not match the sequence number that the Discover appliancewas expecting. The reordering may have been introduced at the point of origin or an intermediary. This may result in decreased call quality.

RTP Out

Contains the following metrics:

Messages

The number of outgoing messages associated with RTP transmissions.

Drops

The number of outgoing packets associated with RTP transmissions which were lost in transit.

Duplicates

The number of outgoing duplicate messages associated with RTP transmissions.

Out of Order

Number of outgoing packets associated with RTP transmissions where the sequence number did not match the sequence number that the Discover appliancewas expecting. The reordering may have been introduced at the point of origin or an intermediary. This may result in decreased call quality.

RTP In by Codec

Displays the number of RTP packets in by codec as a function of time over the selected time interval.

RTP Out by Codec

Displays the number of RTP packets out by codec as a function of time over the selected time interval.

Throughput

Displays the number of RTP packets transmitted as a function of time over the selected time interval.

Message Metrics In

The number of incoming drops, duplicates, and out of order messages associated with RTP transmissions over the selected time interval.

Message Metrics Out

The number of outgoing drops, duplicates, and out of order messages associated with RTP transmissions over the selected time interval.

Jitter In

An estimate of the statistical variance of the incoming RTP packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

Jitter Out

An estimate of the statistical variance of the outgoing RTP packets' interarrival time, measured in timestamp units and expressed as an unsigned integer.

MOS In

The mean opinion score calculated for incoming packets associated with RTP transmissions.

MOS Out

The mean opinion score calculated for outgoing packets associated with RTP transmissions.

RTP groups page

RTP In

Contains the following metrics:

Messages

The number of incoming messages associated with RTP transmissions.

Drops

The number of incoming packets associated with RTP transmissions which were lost in transit.

Duplicates

The number of incoming duplicate messages associated with RTP transmissions.

Out of Order

Number of incoming packets associated with RTP transmissions where the sequence number did not match the sequence number that the Discover appliancewas expecting. The reordering may have been introduced at the point of origin or an intermediary. This may result in decreased call quality.

RTP Out

Contains the following metrics:

Messages

The number of outgoing messages associated with RTP transmissions.

Drops

The number of outgoing packets associated with RTP transmissions which were lost in transit.

Duplicates

The number of outgoing duplicate messages associated with RTP transmissions.

Out of Order

Number of outgoing packets associated with RTP transmissions where the sequence number did not match the sequence number that the Discover appliancewas expecting. The reordering may have been introduced at the point of origin or an intermediary. This may result in decreased call quality.

SIP

ExtraHop appliances collect metrics about Session Initiation Protocol (SIP) activity.

SIP applications page

SIP Application Toolbar

The SIP application toolbar includes the following controls:

Errors

The chart shows the number of SIP errors. Mouse over the points to view a summary of a specific time or date. The table lists the URIs in error and the number of times an error occurred.

Initiators

Displays the list of initiators establishing connections over the selected time interval.

URIs

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists the URIs, number of responses, total time (ms), and processing time (ms) associated with each URI. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Methods

The chart shows the total count compared to the processing time (ms). Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists methods, total time (ms), and processing time (ms) associated with each method. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Initiator

Displays application metrics by initiator.

By URI

Displays application metrics by URI.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

SIP Metrics

Contains the following metrics:

Requests

The number of SIP requests.

Responses

The number of SIP responses.

Response Errors

The number of SIP errors for the selected time interval.

Methods

Displays the SIP methods for the selected time interval.

Status Codes

The status code section displays the HTTP status codes for the selected time interval. Click the number next to each status code to display a list of IP addresses associated with each status code.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Server Processing Time

Displays the median server processing time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing time metrics. Click and drag across the chart to select a particular region.

SIP devices page

SIP Device Toolbar

This SIP device toolbar includes the following controls:

SIP Metric Type

Displays metrics for the current device acting as a SIP client or SIP server.

Errors

Displays a detailed list of error messages sent to or received by the current device over the specified time interval.

Clients or Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

Records

Displays results for records that match the selected metric source and protocol.

Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays device metrics by IP address.

By URI

Displays device metrics by URI.

SIP Client

If you select **Client** for the SIP Metric Type, the Discover appliance displays the following metrics:

Requests

Specifies the number of requests that the device sent when acting as a SIP client. Click the counter to display the list of servers to which requests were sent.

Responses

Specifies the number of responses that the device received when acting as a SIP client. Click the counter to display the list of servers from which the responses were received.

Response Errors

Specifies the number of response errors for the selected time interval when acting as a SIP client. Click the counter to display the list of servers associated with the errors.

SIP Server

If you select Server for the SIP Metric Type, the Discover appliance displays the following metrics:

Requests

Specifies the number of requests that the device received when acting as a SIP server. Click the counter to display the list of clients from which requests were received.

Responses

Specifies the number of responses that the device sent when acting as a SIP server. Click the counter to display the list of clients to which the responses were sent.

Response Errors

Specifies the number of response errors for the selected time interval when acting as a SIP server. Click the counter to display the list of clients associated with the errors.

Methods

Displays the SIP methods for the selected time interval.

Status Codes

The status code section displays the HTTP status codes for the selected time interval. Click the number next to each status code to display a list of URIs associated with each status code.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Server Processing Time

Displays the median server processing time in milliseconds as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the processing time metrics. Click and drag across the chart to select a particular region.

SIP groups page

SIP Groups Toolbar

This SIP groups toolbar includes the following controls:

SIP Metric Type

Displays metrics for the current member acting as a SIP client or SIP server.

Records

Displays results for records that match the selected metric source and protocol.

Methods

Displays the SIP methods for the selected time interval. Click the counter next to the method to break it down by group members in the table.

Status Codes

The status code section displays the status codes for the selected time interval. Click the number next to each status code to display a list of members associated with each status code in the table.

SMPP

ExtraHop appliances collect metrics about Short Message Peer-to-Peer (SMPP) activity.

SMPP devices page

SMPP Device Toolbar

The SMPP device toolbar includes the following controls:

SMPP Metric Type

Displays statistics for devices acting as an SMPP client or server.

Errors

Click the **Errors** button to display the list of error messages sent to or received by devices over the selected time interval.

IP Address SMPP Metrics

Click the counters next to individual SMPP metrics to show the IP Address SMPP Metrics for SMPP peer devices. For SMPP servers, the peer devices are SMPP clients. For SMPP clients, the peer devices are SMPP servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

SMPP Server

Click the counter next to each metric to break it down by devices in the table at the bottom of the page.

Requests

Number of requests that the device received when acting as an SMPP server (SMSC).

Responses

Number of responses that the device sent when acting as an SMPP server (SMSC).

Errors

Number of SMPP errors for the selected time interval.

SMPP Client

Click the counter next to each metric to break it down by device in the table at the bottom of the page.

Requests

Number of requests that the device sent when acting as an SMPP client (ESME).

Responses

Number of responses that the device received when acting as an SMPP client (ESME).

Errors

Number of SMPP errors for the selected time interval.

Inbound Messages

The Inbound Messages section displays the inbound SMPP message types for the selected time interval. Refer to the SMPP specification for a comprehensive list of message types.

Transaction Status

Displays the status codes returned by the SMPP server (SMSC) for requests sent to it by the SMPP client (ESME).

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown (ms)

Displays the area chart containing median request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

SMPP devices timing page

Request Transfer Time

Displays a histogram of times it took to transfer requests from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Processing Time

Displays a histogram of times it took the server to process requests. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Response Transfer Time

Displays a histogram of times it took to transfer the response from the server to the client. Mouse over each bar to display the time range it represents and the number of requests in this bin.

SMPP groups page

SMPP Groups Toolbar

The SMPP groups toolbar includes the following controls:

SMPP Metric Type

Click the **Metric Type** drop-down list and select either **Client** or **Server** to display statistics for members in the current group acting as an SMPP client or server, respectively.

Errors

Click the **Errors** button to display the list of error messages sent to or received by members in the current group over the selected time interval.

SMPP Server

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of SMPP requests for the selected time interval.

Responses

Number of SMPP responses for the selected time interval.

Errors

Number of SMPP errors for the selected time interval.

SMPP Client

Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Number of SMPP requests for the selected time interval.

Responses

Number of SMPP responses for the selected time interval.

Errors

Number of SMPP errors for the selected time interval.

Inbound Messages

The Inbound Messages section displays the inbound SMPP message types for the selected time interval. Refer to the SMPP specification for a comprehensive list of message types.

Outbound

The Outbound Messages section displays the outbound SMPP message types for the selected time interval. Refer to the SMPP specification for a comprehensive list of message types.

Transaction Status

Displays the status codes returned by the SMPP server (SMSC) for requests sent to it by the SMPP client (ESME).

SMPP groups processing time page

Server Processing Time

Shows median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

SMTP

ExtraHop appliances collect metrics about Simple Mail Transfer Protocol (SMTP) activity.

SMTP applications page

SMTP Applications Page

The SMTP application toolbar includes the following controls:

Senders

The chart shows bytes transferred compared with message size. Mouse over points to view a summary of message size. The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists sender domains (HELO or EHLO command argument), bytes transferred, and mean message sizes.

Recipients

The chart shows bytes transferred compared with message size. Mouse over points to view a summary of message size. The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists recipient email addresses (RCPT TO command argument), bytes received, and mean message sizes.

Sender Domains

The chart shows bytes transferred. Mouse over the points to view a summary of a specific time or date. The table lists sender domains and the bytes transferred for each.

Errors

The chart shows the number of SMTP errors that occurred. Mouse over the points to view a summary of a specific time or date. The table lists SMTP error messages and the number of times each occurred.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

SMTP Metrics

Contains the following metrics:

Requests

The number of SMTP requests.

Responses

The number of SMTP responses.

Errors

The number of responses by error for the application.

Sessions

The number of SMTP sessions.

Encrypted Sessions

The number of encrypted SMTP sessions.

Methods

Contains metrics for SMTP commands.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

SMTP devices page

SMTP Device Toolbar

The SMTP device toolbar includes the following controls:

SMTP Metric Type

Displays metrics for the current device acting as an SMTP client or SMTP server.

Senders

Displays the list of sender email addresses (`MAIL FROM` command argument), bytes sent, and mean message sizes for the selected time interval.

Recipients

Displays the list of recipient email addresses (`RCPT TO` command argument), bytes received, and mean message sizes for the selected time interval.

Sender Domains

Displays the list of sender domains (`HELO` or `EHLO` command argument) and bytes transferred for the selected time interval.

Errors

Displays the list of error messages sent to or received by the current device over the selected time interval. 4xx and 5xx SMTP responses are considered errors.

Records

Displays results for records that match the selected metric source and protocol.

IP Address SMTP Metrics

Click the counters next to individual SMTP metrics to show the IP Address SMTP Metrics for SMTP peer devices. For SMTP servers, the peer devices are SMTP clients. For SMTP clients, the peer devices are SMTP servers.

IP Address

Represents the IP address of the peer device.

Host

Represents the DNS host name of the peer device determined by passive analysis of the DNS traffic.

Device

Provides a link to the corresponding peer device. For local peer devices, the link leads to that device. For remote peer devices, the link leads to the gateway device through which the requests were routed.

SMTP Client

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests that the device sent when acting as an SMTP client.

Responses

Specifies the number of responses that the device received when acting as an SMTP client.

Errors

Specifies the number of 4xx and 5xx SMTP responses received by the SMTP client.

Sessions

Specifies the number of sessions that the device participated in when acting as an SMTP client.

Encrypted Sessions

Specifies the number of encrypted sessions that the device participated in when acting as an SMTP client.

Requests Aborted

Specifies the number of requests that the device began to send but did not send completely when acting as an SMPT client. Click to display the list of servers to which incomplete requests were sent.

Responses Aborted

Specifies the number of responses that the device began to receive but did not receive completely when acting as an SMTP client. Click to display the list of servers from which incomplete responses were sent.

SMTP Server

Click the counter next to each metric to display additional IP address details.

Requests

Specifies the number of requests that the device received when acting as an SMTP server.

Responses

Specifies the number of responses that the device sent when acting as an SMTP server.

Errors

Specifies the number of 4xx and 5xx SMTP responses sent by the SMTP server.

Sessions

Specifies the number of sessions that the device participated in when acting as an SMTP server.

Encrypted Sessions

Specifies the number of encrypted sessions that the device participated in when acting as an SMTP server.

Requests Aborted

Specifies the number of requests that the device began to receive but did not receive completely when acting as an SMTP server. Click to display the list of clients from which incomplete requests were sent.

Responses Aborted

Specifies the number of responses that the device began to send but did not send completely when acting as an SMTP server. Click to display the list of clients to which incomplete responses were sent.

Methods

Displays the SMTP methods for the selected time interval. Methods include standard SMTP methods as well as Microsoft Exchange specific methods. Click the counter to display additional per-client or per-server IP address details.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Request Size

Displays the range of request sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Response Size

Displays the range of response sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

SMTP Throughput

Displays the incoming and outgoing SMTP throughput (bytes per second) over the selected time interval. Click and drag across the chart to select a particular region.

SMTP groups page

SMTP Groups Toolbar

The SMTP groups toolbar includes the following controls:

SMTP Metric Type

Displays metrics for members in the current group acting as an SMTP client or server, respectively.

Senders

Displays the list of sender email addresses (MAIL FROM command argument), bytes sent, and mean message sizes for the selected time interval.

Recipients

Displays the list of recipient email addresses (RCPT TO command argument), bytes received, and mean message sizes for the selected time interval.

Sender Domains

Displays the list of sender domains (HELO or EHLO command argument) and bytes transferred for the selected time interval.

Errors

Displays the list of error messages sent to or received by members in the current group over the selected time interval. 4xx and 5xx SMTP responses are considered errors.

Records

Displays results for records that match the selected metric source and protocol.

SMTP Client

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of requests sent by the SMTP client.

Responses

Specifies the number of responses received by the SMTP client.

Errors

Specifies the number of 4xx and 5xx SMTP responses received by the SMTP client.

Sessions

Specifies the number of sessions initiated by `HELO` or `EHLO`.

Encrypted

Specifies the number of encrypted sessions initiated by `STARTTLS`.

Requests Aborted

Specifies the number of incomplete SMTP requests sent from all members of the current group.

Responses Aborted

Specifies the number of incomplete SMTP responses received by all members of the current group.

SMTP Server

Click the counter next to the metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of requests received by the SMTP server.

Responses

Specifies the number of responses sent by the SMTP server.

Errors

Specifies the number of 4xx and 5xx SMTP responses sent by the SMTP server.

Sessions

Specifies the number of sessions initiated by `HELO` or `EHLO` commands.

Encrypted

Specifies the number of encrypted sessions initiated by `STARTTLS`.

Requests Aborted

Specifies the number of incomplete **SMPT** requests received by all members of the current group.

Responses Aborted

Specifies the number of incomplete SMPT responses sent from all members of the current group.

Methods

Displays the SMTP methods for the selected time interval. Methods include standard SMTP methods as well as Microsoft Exchange specific methods, such as BDAT, XEXCH, and EXPS. Click the counter next to the method to break it down by group members in the table at the bottom of the page.

SSL

ExtraHop appliances collect metrics about Secure Sockets Layer (SSL) activity.

SSL applications page

SSL Applications Toolbar

The SSL application toolbar includes the following controls:

Certificates

The chart shows the total number of certificates assigned compared with the request and response bytes. Mouse over points to view a summary of a specific time or date.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By Certificate

Displays application metrics by certificate.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

Session Metrics

Contains the following metrics:

Connected

The number of times an SSL handshake was successfully completed.

Resumed

The number of times an SSL session was resumed successfully by reusing a session ID or session ticket.

Decrypted

The number of SSL sessions decrypted.

Aborted

The number of SSL sessions that did not proceed past the SSL handshake.

Renegotiated

Specifies the number of times an SSL session was renegotiated successfully after SSL connection setup.

Compressed

The number of SSL sessions using compression.

SSLv2 Compatible Hello

The number of SSL sessions for which the private key was available, enabling their decryption.

Sessions by Version

The number of times a session used a particular SSL version.

Cipher Suites

Displays the number of times various cryptographic ciphersuites for SSL data transfer have been negotiated by the application.

For example, `TLS_RSA_WITH_AES_256_CBC_SHA` indicates:

- TLS (Transport Layer Security) is used as the cryptographic encapsulation transport
- RSA (the Rivest-Shamir-Adelman Public Key method RSA) is used for the asymmetric cryptographic session setup
- AES (Advanced Encryption Standard, formerly Rijndael) block cipher is used in 256-bit blocks
- CBC (Cipher Block Chaining) is used between subsequent AES-256 blocks
- SHA (Secure Hash Algorithm) is used in the HMAC (Hash Message Authentication Code) to ensure SSL record integrity

For each cipher suite, click the counter to break it down by group members in the table below.

Alerts

Displays the breakdown of alert types sent or received by the current application during the SSL connection. This section displays unencrypted alerts gathered during the SSL handshake and any alerts that were decrypted by the Discover appliance. Alert messages can be exchanged during other stages of the SSL connection. The handshake metrics display the number of times alerts were exchanged during the SSL handshake. The Warning-Close Notify metric displays the number of times various alert types were sent or received by the application.

SSL Metrics

The SSL Metrics line chart displays the rate of new and resumed SSL connections as a function of time over the selected time interval. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

SSL devices page

SSL Device Toolbar

The SSL device toolbar includes the following controls:

SSL Metric Type

Displays metrics for the current device acting as an SSL client or server.

Certificates

Displays the X.509 subject field, key type, key size, expiration dates, and number of times each certificate has been accessed.

Records

Displays results for records that match the selected metric source and protocol.

Session Details

Displays the session details.

Connected

Specifies the number of times current device successfully completed an SSL handshake.

Resumed

Specifies the number of times a prior SSL session was resumed successfully by reusing a previously negotiated session ID.

Decrypted

Specifies the number of SSL decrypted records encountered.

Aborted

Specifies the number of times the current device did not proceed past an SSL handshake.

Renegotiated

Specifies the number of times an SSL session was renegotiated successfully after SSL connection setup.

Compressed

Specifies the number of SSL compressed records encountered.

SSLv2

Specifies the number of times an SSLv2 hello was sent by a client.

Sessions by Version

Displays the number of times a particular SSL version was used in communication.

Cipher Suites

Displays the number of times various cryptographic ciphersuites for SSL data transfer have been negotiated by this device. For example, `TLS_RSA_WITH_AES_256_CBC_SHA` indicates:

- TLS (Transport Layer Security) is used as the cryptographic encapsulation transport.
- RSA (the Rivest-Shamir-Adelman Public Key method RSA) is used for the asymmetric cryptographic session setup.
- AES (Advanced Encryption Standard, formerly Rijndael) block cipher is used in 256-bit blocks.
- CBC (Cipher Block Chaining) is used between subsequent AES-256 blocks.
- SHA (Secure Hash Algorithm) is used in the HMAC (Hash Message Authentication Code) to ensure SSL record integrity.

Records by Content Type

Displays the number of records for each content type in the specified time interval.

Alert

Specifies the number of messages with an Alert content type (21), used to signal unexpected events.

Application Data

Specifies the number of messages with an Application content type (23), used to send SSL data.

Change Cipher

Specifies the number of messages with a ChangeCipherSpec content type (20), used to signal the beginning and end of encrypted content.

Handshake

Specifies the number of messages with a Handshake content type (22), used to establish the SSL connection.

Heartbeat

Specifies the number of messages with a Heartbeat content type (24) that were sent by a client or server.

Note: The detection of heartbeat messages can indicate an attempted Heartbleed exploit. To monitor exploit attempts, download and install the Heartbleed Bundle from the ExtraHop website.

Alerts

Displays the breakdown of alert types sent or received by the current device. This section displays unencrypted alerts gathered during the SSL handshake and any alerts that were decrypted by the Discover appliance. Alert messages can be exchanged during other stages of the SSL connection. The total number of alert messages exchanged is recorded in the Records by Content Type section, Alert metric.

SSL Metrics

Displays the rate of new SSL connections initiated to (if SSL server) or from (if SSL client) the current device over the selected interval, and the rate of resumed connections. Click and drag across the chart to select a particular region.

Record Size

Displays the five-number summary (low, twenty-fifth percentile, median, seventy-fifth percentile, and high) of the size of SSL records being sent by the current device. The SSL specification mandates a maximum 14KB record size; however certain commercial SSL stacks are known to violate this limit, sometimes resulting in compatibility problems.

SSL groups page

SSL Groups Toolbar

The SSL groups toolbar includes the following controls:

SSL Metric Type

Displays metrics for members of the current group acting as an SSL client or server, respectively.

Certificates

Display the various X.509 subject fields, expiration dates, and number of times each certificate has been accessed.

Records

Displays results for records that match the selected metric source and protocol.

Session Details

For each detailed metric, click the counter to break it down by group members in the table at the bottom of the page.

Connected

Specifies the number of times the current member successfully completed an SSL handshake.

Resumed

Specifies the number of times a prior SSL session was resumed successfully by reusing a previously negotiated session ID.

Decrypted

Specifies the number of SSL decrypted records encountered.

Aborted

Specifies the number of times the current member did not proceed past an SSL handshake.

Renegotiated

Specifies the number of times SSL session was renegotiated successfully after SSL connection setup.

Compressed

Specifies the number of SSL compressed records encountered.

SSLv2 Compatible Hello

Specifies the number of times an SSLv2 hello was sent by a client.

Sessions by Version

Displays the number of times a particular SSL version was used in communication. For each version, click the counter to break it down by group members in the table at the bottom of the page.

Cipher Suites

Displays the number of times various cryptographic cipher suites for SSL data transfer have been negotiated by this member.

For example, TLS_RSA_WITH_AES_256_CBC_SHA indicates:

- TLS (Transport Layer Security) is used as the cryptographic encapsulation transport.
- RSA (the Rivest-Shamir-Adelman Public Key method RSA) is used for the asymmetric cryptographic session setup.
- AES (Advanced Encryption Standard, formerly Rijndael) block cipher is used in 256-bit blocks.
- CBC (Cipher Block Chaining) is used between subsequent AES-256 blocks.
- SHA (Secure Hash Algorithm) is used in the HMAC (Hash Message Authentication Code) to ensure SSL record integrity.

For each cipher suite, click the counter to break it down by group members in the table at the bottom of the page.

Records by Content Type

For each content type, click the counter to break it down by group members in the table at the bottom of the page.

Alert

Specifies the number of messages with an Alert content type (21), used to signal unexpected events.

Application Data

Specifies the number of messages with an Application content type (23), used to send SSL data.

Change Cipher

Specifies the number of messages with a ChangeCipherSpec content type (20), used to signal the beginning and end of encrypted content.

Handshake

Specifies the number of messages with a Handshake content type (22), used to establish the SSL connection.

Heartbeat

Specifies the number of messages with a Heartbeat content type (24) that were sent by a client or server.

Note: The detection of heartbeat messages can indicate an attempted Heartbleed exploit. To monitor exploit attempts, download and install the Heartbleed Bundle from the ExtraHop website.

Alerts

Displays the number of times various alert types were sent or received by members in this group. This section displays unencrypted alerts gathered during the SSL handshake and any alerts that were decrypted by the Discover appliance. Alert messages can be exchanged during other stages of the SSL connection.

Storage NAS

ExtraHop appliances collect metrics about network-attached storage (NAS) activity.

Storage - NAS applications page

NAS Application Toolbar

The NAS application toolbar includes the following controls:

Errors

The chart shows the number of Storage - NAS errors. Mouse over the points to view a summary of a specific time or date.

The table lists Storage - NAS error messages and the number of times each occurred.

Warnings

The chart shows the number of Storage - NAS warnings.

Files

The chart shows responses compared with access time. Mouse over the points to view a five-number summary of access time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists files, number of responses, and the access time (ms) associated with each file.

Users

The chart shows responses compared with request and response bytes. Mouse over the chart to see summaries of a specific date or time. The table lists users and the number of responses, request bytes, response bytes, and access time associated with each user.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By File

Displays application metrics by file name.

By User

Displays application metrics by user.

For example, Client Bytes is a top-level metric showing how many client bytes were transmitted in and out of the application within the selected time interval. Selecting **By Client IP** in the drop-down list while mousing over the **Client Bytes** counter shows which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

Storage - NAS Metrics

Contains the following metrics:

Responses

The number of NAS responses.

Response Errors

The number of NAS response errors.

Response Warnings

The number of responses with a status code of 4xx.

Reads

The number of NAS read operation requests.

Writes

The number of NAS write operation requests.

FS Info

The number of NAS file system metadata queries.

Locks

The number of NAS lock operation requests.

Access Time

The time to access a file on a CIFS or NFS partition. For CIFS, the access time is measured by timing the first `READ` or `WRITE` on every flow. For NFS, the access time is measured by timing non-pipelined commands for every `READ` and `WRITE`.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Read, Write, and FSInfo Bytes

Displays the total bytes per application transmitted within the selected time interval. Mouse over the graph to see the byte count for each metric at a specific moment in time.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

VLANs

ExtraHop appliances collect metrics about Virtual Local Area Network (VLAN) activity.

VLANs networks page

Top VLANs (Packets)

Displays how VLANs contribute to the total packet count for the network.

Top VLANs (Bytes)

Displays how VLANs contribute to the total byte count for the network.

Click a VLAN in the legend to view an isolated graph of its activity over time.

VLANs networks details page

Click a VLAN listed in the table to list the devices sending or receiving the traffic for that VLAN. The VLAN groups appear in a table with the following headings:

Group

Provides a link to a list of devices in the corresponding VLAN group.

Packets

Represents the total packet count for the currently selected VLAN group.

Bytes

Represents the total byte count for the currently selected VLAN group.

When you click a VLAN group, the VLAN device metrics appear in a table with the following headings:

Device

Provides a link to the corresponding device. For local devices, the link leads to that device. For remote devices, the link leads to the gateway device through which the requests were routed.

Packets In

Represents the incoming packet rate for the currently selected VLAN in the area chart.

Packets Out

Represents the outgoing packet rate for the currently selected VLAN in the area chart.

Bytes In

Represents the incoming byte count for the currently selected VLAN in the area chart.

Bytes Out

Represents the outgoing byte count for the currently selected VLAN in the area chart.

VoIP

ExtraHop appliances collect metrics about voice over IP (VoIP) activity.

VoIP applications page

SIP Invites

Displays the number of SIP invites as a function of time over the selected time interval.

RTP Messages by Codec

Displays the number of RTP messages by codec as a function of time over the selected time interval. Click the chart to view a table with the total number of messages broken down by codec.

VoIP

Displays the number of VoIP packets transmitted as a function of time over the selected time interval.

VoIP devices page

VoIP Device Toolbar

The VoIP device toolbar includes the following controls:

VoIP Metric Type

Displays metrics for the current device acting as an VoIP client or server.

SIP Invites

Displays the number of SIP invites as a function of time over the selected time interval.

RTP In by Codec

Displays the number of RTP messages in by codec as a function of time over the selected time interval. Click the chart to view a table with the total number of messages broken down by codec.

RTP Out by Codec

Displays the number of RTP messages out by codec as a function of time over the selected time interval. Click the chart to view a table with the total number of messages broken down by codec.

VoIP

Displays the number of VoIP packets transmitted as a function of time over the selected time interval.

Web HTTP

ExtraHop appliances collect metrics about Web or Hypertext Transfer Protocol (HTTP) activity.

Web applications page

Web Application Toolbar

The Web application toolbar includes the following controls:

Errors

The chart shows the number of HTTP errors (5xx level responses). Mouse over the points to view a summary of a specific time or date. The table lists HTTP URIs in error and the number of times an error occurred.

URIs

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists HTTP URIs, number of responses, total time (ms), and processing time (ms) associated with each URI. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Referers

The chart shows the number of HTTP referer URIs identified. Mouse over the points to view a summary of a specific time or date. The table lists the HTTP referer URIs and the count associated with each referer.

Clients

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists client IP addresses, the host and device associated with each client, the number of responses from each client, and the total time and processing time for each client.

Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Servers

The chart shows the total number of responses compared to processing time. Mouse over the points to view a five-number summary of processing time (minimum, lower quartile, median, upper quartile, and maximum values). The orange bars represent a confidence interval around the median value bounded by the 25th and 75th percentile values.

The table lists server IP addresses, the host and device associated with each server, the number of responses from each server, and the total time and processing time for each server. Mouse over the orange bars to view the mean time, standard deviation, and count for each metric.

Application Details

Specifies the type of additional application information displayed. IP detail views display directly monitored IP addresses and IP addresses that appear via routed traffic. IP addresses that appear via routed traffic are preceded by the word via. Mousing over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By Client IP

Displays application metrics by the client IP addresses.

By Server IP

Displays application metrics by the server IP addresses.

By URI

Displays application metrics by URI.

For example, Request Bytes is a top-level metric showing how many request bytes were transmitted in and out of the application within the selected time interval. Select **By Client IP** in the drop-down list while mousing over the **Request Bytes** counter to view which client IP addresses originated these requests.

L2-L4 Metrics

Contains the following metrics:

Request L2 Bytes

The number of L2 bytes associated with requests.

Response L2 Bytes

The number of L2 bytes associated with responses.

Request Packets

The number of packets associated with requests.

Response Packets

The number of packets associated with responses.

Request RTOs

Specifies the number of times the client delayed TCP retransmissions and missed server acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Response RTOs

Specifies the number of times the server delayed TCP retransmissions and missed client acknowledgments. A retransmission timeout is a 1-second stall in the TCP connection flow due to excessive retransmissions.

Request Zero Window

Specifies the number of client-side zero window advertisements. A zero window indicates the connection has stalled because the client cannot handle the rate of data the server is sending.

Response Zero Window

Specifies the number of server-side zero window advertisements. A zero window indicates the connection has stalled because the server cannot handle the rate of data the client is sending.

HTTP Metrics

Contains the following metrics:

Requests

The number of HTTP requests.

Requests Aborted

The number of HTTP requests that began transmission but were not sent completely.

Responses

The number of HTTP responses.

Responses Aborted

The number of HTTP responses that began transmission but were not sent completely.

Response Errors

The number of HTTP response errors.

Status Codes

The status code section displays the HTTP status codes for the selected time interval. Click the number next to each status code to display a list of URIs associated with each status code.

Methods

Displays the HTTP request methods for the selected time interval. The HTTP request methods include `GET`, `HEAD`, `POST`, `PUT`, `DELETE`, `TRACE`, `CONNECT`, and `OPTIONS`, as well as dynamic method names. Click to display additional per-URI, per-client IP, or per-server IP details.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median round-trip time, request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Round-Trip Time (ms)

Displays the median round-trip time (RTT) in milliseconds (ms) from the current objects to clients as a function of time over the selected time interval. Vertical dotted lines indicate the upper and lower quartiles (75th and 25th percentiles) of the round-trip time metrics. Click and drag across the chart to select a particular region.

Congestion Requests: Goodput (bps) and RTOs

Displays goodput and RTOs into the object as a function of time over the selected time interval.

Congestion Responses: Goodput (bps) and RTOs

Displays goodput and RTOs out of the object as a function of time over the selected time interval.

Goodput is application-level throughput (the number of useful information bits) and RTOs are retransmission timeouts. The Congestion In and Out graphs show the relationship over time between the rate of good application throughput and RTOs. An increase in RTOs theoretically leads to a decrease in goodput due to TCP back-off and packet retransmissions. It is best to view these charts in a smaller window of time so the metrics taken over time are not rolled up or smoothed out. In a small timeframe (30 minutes or less), one could see a decrease in goodput associated with a large number of RTOs, assuming that most flows on the server during this time frame experience this behavior. If only one or two flows are affected by RTOs, then the decreased goodput correlation may be masked by superficially healthy flows.

HTTP devices page

HTTP Device Toolbar

The HTTP device toolbar includes the following controls:

HTTP Metric Type

Displays metrics for the current device acting as an HTTP client or HTTP server.

Errors

Displays the list of error messages sent to or received by the current device over the selected time interval.

URIs

Displays the list of HTTP URIs, number of responses, total time (ms), and processing time (ms) associated with each URI.

Referers

Displays the list of HTTP referer URLs and the count associated with each referer.

Clients and Servers

Displays the associated client IP addresses when the device is acting as a server, and the associated server IP addresses when acting as a client.

Records

Displays results for records that match the selected metric source and protocol.

Moving the mouse pointer over the counter next to each top-level metric opens a context menu that includes the following options in the drop-down list:

By IP

Displays HTTP metrics by IP addresses.

By Host

Displays HTTP metrics by host name.

By URI

Displays HTTP metrics by URI.

For example, HTTP Requests is a top-level metric showing how many requests were received by the HTTP server during the selected time frame. Selecting **By IP** in the drop-down list while moving the mouse pointer over the HTTP Requests counter shows which IP addresses originated these requests.

Selecting **By URI** from the drop-down list while moving the mouse pointer over the HTTP Requests counter shows which URIs were accessed by the requestors.

HTTP Metrics by IP Address

Click **By IP** in the drop-down list to display the following information in the details table.

IP Address

Represents the HTTP server's IP address.

Host

Represents the DNS host name of the HTTP server determined by passive analysis of the DNS traffic.

Origin Address

Represents the origin address of a client that is connected through a proxy or load balancer.

Device

Provides a link to the corresponding HTTP server device. For local HTTP servers, the link leads to the HTTP server device. For remote HTTP servers, the link leads to the gateway device through which the requests were routed.

<Metric value>

Displays the value for the selected metric.

HTTP Metrics by Host

Click **By Host** in the drop-down list to display the following information in the details table.

HTTP Host

Represents the virtual host as defined in the Host attribute of the HTTP request header.

<Metric value>

Displays the value for the selected metric.

HTTP Metrics by URI

Click **By URI** in the drop-down list to display the following information in the details table.

URI

Represents the full HTTP URI.

<Metric value>

Displays the value for the selected metric.

Processing Time

Represents the time in milliseconds it took to process URIs requested by the currently selected HTTP client. Timing information is expressed as a confidence interval around the mean value bounded by one standard deviation. This metric is available for successful HTTP responses only.

HTTP Client

If you select **Client** for the HTTP Metric Type, the Discover appliance displays the following metrics:

Requests

Specifies the number of requests that the device sent when acting as an HTTP client. Click to display the list of servers to which requests were sent.

Requests Aborted

Specifies the number of requests that the device began to send but did not send completely when acting as an HTTP client. Click to display the list of servers to which incomplete requests were sent.

Pipelined Requests

Specifies the number of pipelined requests that the device sent when acting as an HTTP client. Pipelined requests consist of multiple requests written to the same connection without waiting for the corresponding responses. Click to display the list of servers to which pipelined requests were sent.

Responses

Specifies the number of responses that the device received when acting as an HTTP client. Click to display the list of servers from which the responses were received and per-server response times. This metric also provides the detailed per-server and per-URI processing time information.

Responses Aborted

Specifies the number of responses that the device began to receive but did not receive completely when acting as an HTTP client. Click to display the list of servers from which incomplete responses were sent.

Chunked Transfers

Specifies the number of responses received that used chunked transfer coding when the device is acting as an HTTP client. Click to display the list of servers from which chunked responses were sent.

Compressed Transfers

Specifies the number of responses received that used 'gzip' or 'deflate' content coding when the device is acting as an HTTP client. Click to display the list of servers from which compressed responses were received.

Authed Requests

Specifies the number of HTTP requests that provided an Authorization request header and did not receive a 401 status code in the response. Click to display the list of servers to which authorized requests were sent.

HTTP Server

If you select **Server** for the HTTP Metric Type, the Discover appliance displays the following metrics:

Requests

Specifies the number of requests that the device received when acting as an HTTP server. Click to display the list of clients from which requests were received.

Requests Aborted

Specifies the number of requests that the device began to receive but did not receive completely when acting as an HTTP server. Click to display the list of clients from which incomplete requests were sent.

Pipelined Requests

Specifies the number of pipelined requests that the device received when acting as an HTTP server. Pipelined requests consist of multiple requests written to the same connection without waiting for the corresponding responses.

Responses

Specifies the number of responses that the device sent when acting as an HTTP server. Click to display the list of clients to which the responses were sent and per-

client response times. This metric also provides the detailed per-server and per-URI processing time information.

Responses Aborted

Specifies the number of responses that the device began to send but did not send completely when acting as an HTTP server. Click to display the list of clients to which incomplete responses were sent.

Chunked Transfers

Specifies the number of responses sent that used chunked transfer coding when the device is acting as an HTTP server. Click to display the list of clients to which chunked responses were sent.

Compressed Tranfers

Specifies the number of responses sent that used 'gzip' or 'deflate' content coding when the device is acting as an HTTP server. Click to display the list of clients to which compressed responses were sent.

Authed Requests

Specifies the number of HTTP requests that provided an Authorization request header and did not receive a 401 status code in the response. Click to display the list of clients from which authorized requests were sent.

Status Codes

Displays the HTTP response status codes for the selected time interval. Click to display additional per-client or per-server details.

Methods

Displays the HTTP request methods for the selected time interval. The HTTP request methods include `GET`, `HEAD`, `POST`, `PUT`, `DELETE`, `TRACE`, `CONNECT`, and `OPTIONS`, as well as dynamic method names. Click to display additional per-client or per-server details.

Transactions Metrics

Transaction metrics display the timing components for all transactions associated with the current device. Timing components are expressed as a confidence interval around the median value bounded by the 25th and 75th percentile values. Mouse over each component to display a five-number statistical summary.

ReqXfer

Request transfer time. The time in milliseconds before the request was received by the server. A large ReqXfer value relative to the total transaction time indicates network delay. If the request size is large, some network delay due to transfer time is expected.

Process

Server processing time. The time in milliseconds between the time the request was received by the server and the time the response was sent. A large server processing time indicates application delay.

RspXfer

Response transfer time. The time in milliseconds before the server finished sending the response. A large RspXfer relative to the total transaction time indicates network delay. If the response size is large, some network delay due to transfer time is expected.

RTT

TCP round-trip time in milliseconds. Large round-trip time indicates that network latency is high.

Click the **Transaction Metrics** graph to display a chart showing responses compared to mean processing time during the selected time interval. The table below contains the total and mean time for each response.

Request Size

Displays the range of request sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Response Size

Displays the range of response sizes for all transactions associated with the current device. The five-number summary includes the minimum, lower quartile, median, upper quartile, and maximum values. Click to display the mean request size for each peer device.

Transactions Per Second

Displays the number of protocol transactions per second as a function of time over the selected time interval. The chart is annotated with red data points to indicate errors. The volume of errors is denoted by the height of red bars under the chart. Click the red dot to see the number of errors that occurred at that time. Click and drag across the chart to select a particular region.

Response Time Breakdown

Displays the area chart containing median request transfer time, server processing time, and response transfer time over time in milliseconds. Click and drag across the chart to select a particular region.

Conetnt Types

Displays the relative frequencies of HTTP response content types. Click to display relative frequencies within a content-type category.

HTTP devices timing page

Request Transfer Time

Displays a histogram of times it took to transfer requests from the client to the server. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Processing Time

Displays a histogram of times it took the server to process requests. Mouse over each bar to display the time range it represents and the number of requests in this bin.

Response Tranfer Time

Displays a histogram of times it took to transfer the response from the server to the client. Mouse over each bar to display the time range it represents and the number of requests in this bin.

HTTP groups page

HTTP Groups Toolbar

The HTTP groups toolbar includes the following controls:

HTTP Metric Type

Displays metrics for members in the current group acting as an HTTP client or HTTP server, respectively.

Errors

Displays the list of error messages sent to or received by the current member over the selected time interval.

URIs

Displays the list of HTTP URIs, number of responses, total time (ms), and processing time (ms) associated with each URI.

Referers

Displays the list of HTTP referer URLs and the count associated with each referer.

Records

Displays results for records that match the selected metric source and protocol.

HTTP Client

If you select **Client** for the HTTP Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of HTTP requests sent from all members of the current group.

Requests Aborted

Specifies the number of incomplete HTTP requests sent from all members of the current group.

Pipelined Requests

Specifies the number of HTTP/1.1 pipelined requests sent from all members of the current group. Pipelined requests consist of multiple requests written to the same connection without waiting for the corresponding responses.

Responses

Specifies the number of HTTP responses received by all members of the current group.

Responses Aborted

Specifies the number of incomplete HTTP responses received by all members of the current group.

Chunked Transfers

Specifies the number of HTTP/1.1 responses that made use of chunked transfer-coding received by all members of the current group.

Compressed Tranfers

Specifies the number of HTTP/1.1 responses that made use of gzip or deflate content coding.

Authed Requests

Specifies the number of HTTP requests that provided an Authorization request header and did not receive a 401 status code in the response.

HTTP Server

If you select **Server** for the HTTP Metric Type, the Discover appliance displays the following metrics. Click the counter next to each metric to break it down by group members in the table at the bottom of the page.

Requests

Specifies the number of HTTP requests received by all members of the current group.

Requests Aborted

Specifies the number of incomplete HTTP requests received by all members of the current group.

Pipelined Requests

Specifies the number of HTTP/1.1 pipelined requests received by all members of the current group. Pipelined requests are when multiple requests are written to the same connection without waiting for the corresponding responses.

Responses

Specifies the number of HTTP responses sent from all members of the current group.

Responses Aborted

Specifies the number of incomplete HTTP responses sent from all members of the current group.

Chunked Tranfers

Specifies the number of HTTP/1.1 responses that made use of chunked transfer-coding sent from all members of the current group.

Compressed

Specifies the number of HTTP/1.1 responses that made use of gzip or deflate content coding.

Authed Requests

Specifies the number of HTTP requests that provided an Authorization request header and did not receive a 401 status code in the response.

Status Codes

Displays the HTTP response status codes for the selected time interval. Click the counter next to each status code to break it down by group members in the table at the bottom of the page.

Methods

Displays the HTTP request methods for the selected time interval. The HTTP request methods include `GET`, `HEAD`, `POST`, `PUT`, `DELETE`, `TRACE`, `CONNECT`, and `OPTIONS`, as well as dynamic method names. Click the counter next to each status code to break it down by group members in the table at the bottom of the page.

HTTP groups processing time page

HTTP Group Toolbar

The HTTP group toolbar includes the following controls:

HTTP Metric Type

Displays metrics for members in the current group acting as an HTTP client or HTTP server, respectively.

Records

Displays results for records that match the selected metric source and protocol.

Server Processing Time

Shows median server processing time over the selected time interval for each member in the group. The five-number summary, which includes the minimum, lower quartile, median, upper quartile, and maximum values, is displayed by hovering over a bar.

Records

Records are structured flow and transaction information about events on your network. After you pair an ExtraHop Discover appliance to an ExtraHop Explore appliance, you can generate and send record information to the Explore appliance for storage and retrieval.

Note: Records are not visible in the ExtraHop Web UI until after the Explore appliance connection is established.

Records enable you to see the details of your summary metrics to identify potential problems. For example, if you had fifty HTTP 503 errors, you could view details about those errors by querying the records stored on the Explore appliance. The records would contain specific information about each individual HTTP transaction, which might reveal the underlying problem.

Types of records

Each record has a single record type that can be linked to a record format. Record formats determine how records are queried and displayed in the ExtraHop Web UI. You can generate and send two types of records: flow records and transaction records.

Flow records

Flow records show communication between two devices over an (L3) IP protocol. To generate and send flow records to an Explore appliance, you must enable the Automatic Flow Records setting in the ExtraHop Discover appliance Admin UI. After this setting is enabled, stored flow records can be queried from the local Discover appliance or the Command appliance.

For more information, see the Automatic Flow Records section in the ExtraHop Admin UI Guide.

Transaction records

Transaction records show details from individual messages or transactions over L7 protocols. There are three types of supported protocols: transactional (such as HTTP, CIFS, and NFS), message-based (such as ActiveMQ, DNS, and DHCP), and connection-based (such as SSL and ICA).

To generate transaction records for protocol metrics, you must write triggers for the protocol events (such as HTTP responses) that you want to capture and store. After the triggers are written, stored transaction records can be queried from the local Discover appliance or Command appliance.

For more information, see the ExtraHop Trigger API Reference.

Querying records

After records are sent to an Explore appliance, you can query and display the information stored in those records about any object on your network. In addition, you can save record queries to run at a later time.

You can query records that are stored in the Explore appliance from four areas in the ExtraHop Web UI:

- Click **Metrics**, and then click **Record Queries**. This page displays a list of saved queries and enables you to create new queries.
- Type a search term in the global search box at the top of the screen and click Global Records Query to start a query across all stored records.
- Click the Records icon  from the panel of Action icons on an application or device protocol page that has built-in record formats. This option queries for records that match the selected metric source and protocol.

- Click the Records icon in the left-hand column from any drill-down metrics page. This option queries for records that match the selected metric source, protocol, and detailed stat value.

Viewing Query Results

All queried records can be displayed in table or verbose view. The record information displayed in table view is determined by the record format on the local Discover or Command appliance. Built-in record formats are available for most major protocols, which cannot be modified. However, you can create custom record formats to display custom record types that you have sent to the Explore appliance through a trigger.

For more information about creating custom record formats, see the Record Formats section.

Create a record query

You can create a new query to view records stored on the Explore appliance.

1. In the left pane of the Metrics page, click **Record Queries**.
2. In the content pane, click **New Query**.
   When you click **New Query**, the content pane displays chart and table information about all of the records stored on the Explore appliance during the time period specified in the Global Time Selector. You can view and filter record results through the fields and controls in the left pane and content pane.
3. To save your query, in the upper right hand corner of the navigation bar, click **Save Query as...**. Type a name for your query and then click **Save**.

Record queries page

Left Pane

Record Type

Contains all record formats that are defined in **System Settings** > **Record Format**. Record types listed in italics were created on another Discover or Command appliance on your network. To filter your query results by record type, select one or more record types from the list.

Group By

Contains all of the attributes that are in the selected record type. If you select an attribute from this list, the table in the content pane shows a count for each unique value in the set of query results. If you have selected **Any Type** from the **Record Type** drop-down list, the Refined Results list shows all of the available record types. The **Group By** option only applies to the table view.

Refine Results

Displays the top filterable values and their count for common fields. You can click on a value to add the field as a filter.

Content Pane

Chart Summary

Displays a count of records found during the time period specified in the Global Time Selector.

Query bar

Contains all of the fields that you can filter results by. To filter results, select a field and an operator, enter a search term, and then click Add Filter. Your selections appear in the Filter bar, located above Chart Summary. You can remove filters by clicking the x next to the filter name. For the best search results, always select a specific field from the list instead of Any Field.

Table View Button

Displays the query results in a table. The table only displays columns for records that have the display setting enabled in the record format.

Verbose View Button

Displays the query results as a block of field-value pairs. Verbose View is useful for inspecting all of the data in a record, including fields that do not have the display setting enabled in the record format and which do not appear in Table View.

Columns

Contains all of the fields that can be displayed in the table.

View a saved query

After you save a record query, you can view and run the query at a later time.

1. Click **Metrics**.
2. Click **Record Queries**.
   The Record Queries page appears in the content pane with a list of saved queries.
3. (Optional) To run a saved query, click the name of a query.

   You can run saved queries to view records at different times (such as peak hours) or during notable events (after a new server is added). You can start with a saved query, modify the selected filter criteria, and save the existing query with a unique name. You can also create a new query, add filter criteria, and save the new query.

   Saved queries apply the global time interval when they are run instead of any interval you might have selected when running the initial query. The selections you make for your view (table or verbose) and column preferences are saved.

Delete a saved query

From the Record Queries page, in the **Saved Query** list, click the name of the query you want to delete.

View query properties

1. From the Record Queries page, in the **Saved Query** list, click the name of the query.
2. Click the command menu in the upper right hand corner of the screen and click **Query Properties**.
   The following query properties appear:

   Name

   The name the query was saved as.

   Owner

   The name of the user who created the query.

   Description

   A user-provided description of the query.

   Query Time Limit

   The maximum amount of time that a query search is processed.

   Note: Long running queries can affect Explore appliance performance, regardless of the Query Time Limit value.

Modify query properties

1. From the Record Queries page, in the **Saved Query** list, click the name of the query.
2. Click the command menu in the upper right hand corner of the screen and click **Query Properties**.
3. Modify the query properties, and then click **Save**.

Export query results

1. From the Record Queries page, in the **Saved Query** list, click the name of the query.

2. Click the command menu in the upper right hand corner of the screen and click either **Export to Excel** or **Export to CSV**.

Record formats

Record formats determine how records are queried and displayed in the ExtraHop Web UI. There are built-in record formats for most major protocols, which cannot be modified. However, you can create custom record formats to display records that you have sent to the ExtraHop Explore appliance through a trigger.

Custom record formats can be created, updated, or deleted. You can create a new custom record format, or you can build a custom record format by copying and modifying an existing record format.

For example, if you want to generate and send transaction details for HTTP that are not part of the built-in record format, you can write a trigger to send that custom record information to the Explore appliance. Then, you can create a custom HTTP record format, assign unique names for the record format and record type, and then add the transaction details you want to display.

For more information, see the Records section and the ExtraHop Trigger API Reference.

View record formats

The Record Format Settings page displays a list of all built-in and custom record formats that are available on your local ExtraHop Discover or Command appliance. You can select a record format from the left pane and the parameters and schema for that record format displays in the right pane.

Record formats consist of the following settings:

Display Name

The name displayed for the record format in the Web UI. If there is no record format for the record, the record type is displayed.

Author

(Optional) The author of the record format. All built-in record formats display `ExtraHop` as the author.

Record Type

A unique alphanumeric name that identifies the type of information contained in the associated record format. The record type links the record format with the records that are sent to the EEA. Built-in record formats have a record type that begins with a tilde (~). Custom record formats cannot have a record type that begins with a tilde (~).

Schema on Read

A JSON-formatted array with at least one object, which consists of a field name and value pair. Each object describes a field in the record and each object must have a unique combination of name and data type for that record format. You can create the following objects for a custom record format:

name

The name of the field.

display_name

The display name for the field. If the `display_name` field is empty, the `name` field is displayed.

description

(Optional) Descriptive information about the record format. This field is limited to the Record Format Settings page and is not displayed in any record query.

default_visible

(Optional) If set to `true`, this field displays in the Web UI as a column heading by default in table view.

facet

(Optional) If set to `true`, facets for this field display in the Web UI. Facets are a short list of the most common values for the field that can be clicked to add a filter.

data_type

The abbreviation that identifies the type of data stored in this field. The following data types are supported:

Data Type	Abbreviation
ExtraHop application ID (string)	`app`
Boolean value	`b`
ExtraHop device ID (string)	`dev`
An IPv4 address in dotted-quad format. Greater or less than filters are supported.	`addr4`
An IPv6 address. Only string-oriented filters are supported.	`addr6`
Number (integer or floating point)	`n`
String	`s`

meta_type

The sub-classification of the data type that further determines how the information is displayed in the Web UI. The following meta-types are supported for each of the associated data types:

Data Type	Meta Type
String	• `user`
Number	• `bytes` • `count` • `expiration` • `milliseconds` • `packets` • `timestamp`

Create a new custom record format

1. Click the Systems Settings icon in the top toolbar.
2. Click **Record Formats**.
3. Click **New Record Format**.
4. In the Parameters section, type the following information:

Display Name

The display name for your record format.

Author (Optional)

The name of the record format author.

Record Type

An alphanumeric identifier to match records with your record format.

5. In the Schema on Read section, add information about each of the record fields.

   You can type or paste directly into the Schema on Read editor.

   Note: The following words are reserved and should not be entered as a field name: flowID, ex, client, server, sender, receiver. The following characters should not be included in the name of a field: period (.), colon (;), square brackets ([and]). An exception is raised by the commitRecord global function if any of these words or characters are included in a field name.
6. Click **Create**.
7. (Optional) To see which fields are available for a record, write a trigger to generate and send that record data, and then wait for the records to appear in verbose mode. Then, add a record format with the field information you discovered earlier.

Copy a record format

1. Click the Systems Settings icon in the top toolbar.
2. Click **Record Formats**.
3. Select the name of the record format you want to base the custom record format on.
4. In the Schema on Read section, select and copy the text to your clipboard.
5. Click **New Record Format**.
6. In the Parameters section, type the following information:

   Display Name

   The display name for your record format.

   Author (Optional)

   The name of the record format author.

   Record Type

   An alphanumeric identifier to match records with your record format.
7. In the Schema on Read section, replace the sample text with the one you copied from the existing record, and edit the text as needed.
8. Click **Create**.

Modify a custom record format

1. Click the Systems Settings icon in the top toolbar.
2. Click **Record Formats**.
3. Select the name of the custom record format you want to modify.
4. Edit as needed, and then click **Update**.

Delete a custom record format

1. Click the Systems Settings icon in the top toolbar.
2. Click **Record Formats**.
3. Select the name of the custom record format you want to delete, and then click **Delete**.
4. In the Delete pop-up window, click **Confirm**.

Alerts

The Discover appliance associates a baseline value with every metric it collects and allows users to set alerts on these metrics. Alerting makes it easy to inform teams when there are network, device, or application anomalies or Software License Agreement (SLA) violations. Alerts can send an email message or an SNMP trap to the appropriate people in an organization.

The Discover appliance has two different types of alerts, threshold and trend. Threshold-based alerts are triggered when a monitored metric crosses a defined value in a time period and are well suited to SLA-violation monitoring. Trend-based alerts are triggered when a network statistic is outside of the normal trend learned by the system and are well suited for metrics such as errors where meaningful thresholds are difficult to define. Because trend-based alerts need historical data to define a trend, these alerts are triggered approximately three or four days after the Discover appliance has been installed.

Note: To learn more about alerts, view the following training modules:

- Intro to Alerts
- Configure your first alert

View alerts

To view alerts that have been triggered:

1. Click **Alerts**.
2. Click the Global Time Selector to specify a time interval for viewing all fired alerts in that period.
3. To view details about a specific alert, click the name of the alert in the Alert column.
   The Alert Details window appears.

For threshold alerts, the Alert Details window displays information about the alert.

Name
The name of the alert.

Expression
The metric, time interval, operator, and sensitivity that were defined when the alert was created.

Value
The value of the metric at the time the alert fired. This is used for comparison against the alert expression.

Description
The optional user-defined description of the alert.

For trend alerts, the Trend Alert Details window includes the following:

Name
The name of the alert.

Alert Conditions
The type of alert, time interval, operator, and/or percentage of the trend that were defined when the alert was created.

View at Time of Alert
Displays the alert graph from when the alert was fired.

View Current State
Displays the alert graph of the current trend state of the alert.

Configuring alerts

Discover appliances include built-in alerts that are available by default. However, you can also create custom alerts to inform you of when specific events occur on your network.

The Alerts page provides the following information about each alert:

Name
Specifies the name assigned to the alert. Click the alert name to modify an existing alert.

Author
Specifies the creator of the alert. Alerts loaded by default have the author "ExtraHop."

Metrics
Shows a compact alert definition statement.

Command appliance
Specifies whether the alert was created on the Command appliance or locally on the node.

Description
Provides a space for an optional, user-defined description.

Status
Reflects whether the alert is assigned to any objects.

Unassigned
Alert is not assigned to any objects.

Assigned
Alert is assigned to some objects.

Assigned to all
Alert is assigned to all applicable objects.

Create an alert

Before creating an alert, it is important to determine the metric that you want to monitor, the alert threshold for that metric, and the recipient or recipients of notifications for the alert. An alert is triggered only when a threshold is passed. For example, if a threshold is too low, then one alert will be sent when the threshold is passed and no more alerts will be sent.

Note: All external notification alerts are sent in UTC regardless of the time zone set in the Discover appliance Admin UI.

To create a new alert:

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. On the Alerts page, click the **New** button.
4. Enter the following information in the **Alert Settings** tab.

Alert Settings
Provides configuration settings to define the alert name and the alert expression.

Trend Settings
Provides configuration settings to select the time, lookback, and weight of trend-based alerts.

Description
Provides a space for an optional, user-defined description.

Exclusion Intervals
Displays the exclusion intervals assigned to this alert.

Notifications

Provides configuration settings to identify the email groups that should be notified when this alert fires.

Assignments

Displays where in the system this alert has been manually assigned to a device or group.

5. When you are finished entering the alert settings, click **OK** to save the alert and exit the Alert Configuration dialog box.

Copy an alert

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. In the Alerts table, select the checkbox next to the alert(s) that you want to copy and use as a template for defining a new alert.
4. Click the **Copy** button.

The name of the copied alert is generated automatically by appending the word "(copy)" to the original name.

Assign an alert

Alerts can be assigned to applications, devices, and groups.

1. On the page for an application, device, or group, click the **Alerts** tab.
2. Click the add icon to the left of the **Alerts** field.
3. In the Assign Alerts dialog box, select the device alerts that you want to show in the network capture.
4. In the **Filter** text box, provide an optional filter string to filter the list of alerts by name.
5. Click **OK**.

Enable an alert

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. In the Alerts table, select the checkbox next to the alert that you want to enable.
4. Click the **Enable** button.

Disable an alert

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. In the Alerts table, select the checkbox next to the alert(s) that you want to disable.
4. Click the **Disable** button.

Delete an alert

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. In the Alerts table, select the checkbox next to the alert that you want to delete.
4. Click the **Delete** button.

View alert settings

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.

3. (Optional) Click the drop-down list and select one of the following options:

   All Alerts

   Displays alerts that were created on both the Command appliance and the node.

   Command appliance Alerts

   Displays alerts that were created on the Command appliance.

   Local Alerts

   Displays alerts that were created on the node.

   Note: Alerts created on the Command appliance automatically sync with all of its nodes, but alerts created on individual nodes do not sync with the Command appliance. It is best practice to manage all alerts from the Command appliance rather than the nodes.

4. To view settings for an individual alert, click the name of the alert.

Alert settings

You can view the following alert settings:

The **Alert Settings** tab contains the following fields:

Name

Specifies a name for the alert.

Author

Specifies the creator of the alert. The author is set by the Discover appliance based on the user name for manually created objects or the imported bundle, or set manually by the user. On bundle export, you can specify an author to override authors of any local objects included in the bundle. Alerts loaded by default have the author "ExtraHop."

Disable Alert

Specifies whether the alert is disabled.

Alert Type

Specifies the type of alert.

Threshold

Specifies the threshold value and the threshold time interval used in the alert expression. When a trend alert is configured, the value in the **Threshold** field operates as a logical AND so both the trend and the threshold must be met in order to for the alert to fire. Threshold values for most metrics are integers. For metrics that are collected by the Discover appliance as ratios, a decimal value can be specified in the **Threshold** field. The **Interval** drop-down list specifies the units for the threshold value. **Per Second** specifies that the alert is evaluated over the 30-second interval and then divided by 30 to obtain a per-second value, which is compared against the threshold. Metrics that are collected by the Discover appliance as measurements of time or as sizes do not use the time interval unit.

Detail

For threshold alerts, select the **Top-level** or **Detail** radio button.

Top-level

Specifies summary metrics for an object, such as a device, application, or capture. A threshold alert on a top-level metric identifies the total and compares it to the threshold. Examples of top-level metrics include HTTP Requests and HTTP Responses.

Detail

Specifies metrics for an object keyed by other criteria, such as an IP address. A threshold alert on a detailed metric identifies the count for each keyed item and determines whether any of the counts are over the threshold. Threshold

alerts for detailed metrics may fire multiple times if more than one is over the threshold.

Alert settings allow alerting only on certain criteria. Other criteria may require you to first create a custom metric using the Metric Catalog. For example, to send alerts about metrics containing a specific URI, you can record a custom metric and then alert on the value of the custom metric. For more information, refer to Metric Catalog.

Trend

Trend-based alerts are triggered when a network statistic is outside of the normal trend learned by the system. Trend-based alerts are well suited for metrics such as errors where meaningful thresholds are difficult to define. Trend-based alerts need historical data to define a trend, so these alerts will fire once the Discover appliance has collected enough data to establish a baseline.

Metric

Specifies the metric that this alert is associated with.

To select a metric, click the gear icon to the right of this field.

1. The Select Metric dialog box appears.
2. Expand the **application**, **capture** or **device** nodes to locate the metric that you want to use in the alert expression.
3. In the **Key pattern** text box, specify additional information about the metric to refine your search. The input is interpreted as a regular expression and must use Perl-Compatible Regular Expression (PCRE) syntax. Refer to PCRE documentation for more information.
4. Select the metric and click **OK**.

Ratio

Click the **Ratio** checkbox and select another metric using the gear icon.

Dataset and Sampleset Settings

The following settings are available only for trend alerts with dataset and sampleset metrics:

Merge

Merges all the datasets and applies the trending function to one big dataset.

A 30-second cycle has a single dataset for each 30-second interval, so a 30-minute interval has 60 datasets. You can generate a trendline from these datasets in one of the following ways:

- Take the mean/median/nth percentile of each dataset, and perform a trend calculation on this value. For example, you might want to take the moving average (trend function) of the 95th percentile of processing time.
- Merge all of the datasets together into one big dataset, and perform a trend calculation on this value. For example, you might want to merge the datasets, then take the trimean (trend function) of the combined dataset.

Mean

Takes the mean of each dataset.

Percentile

Allows you to set a percentile value of datasets.

Standard Deviation

Calculates the normal deviation compared to the current trend alert using the same standard deviation parameters as the trend. These parameters can be absolute or relative, and population or sample. Normalization displays the standard deviation relative to mean. Click the **Normalization** drop-down list and select one of the following options.

Absolute

Displays the standard deviation as a constant.

Relative to Mean

Displays the standard deviation relative to the mean.

Note: If the trend alert is not a standard deviation, it is calculated as an absolute sample.

Firing Mode

Specifies the criteria under which the trigger is fired. This selection may affect the behavior of assigned geomaps.

Edge-Triggered

The alert fires once when the threshold is breached and will fire again only when the metric goes below the threshold and breaches it again. For persistent problems (i.e., the threshold is breached continuously), a red dot appears on the geomap when the problem first occurs, but not continuously if the condition is only breached once.

Level-Triggered

The alert fires continuously while the condition is breached at the specified re-firing interval. If you are creating alerts to view in geomaps, ExtraHop recommends configuring level-triggered alerts that re-fire at the same interval or more frequently as defined in the geomap, which is 30 minutes by default.

Alert When

Specifies the criteria for sending the alert.

For trend alerts, the following options are available in the first drop-down list:

mean

Specifies the mean value of the alert.

median

Specifies the median value of the alert.

25th percentile

Specifies the 25th percentile value of the alert.

75th percentile

Specifies the 75th percentile value of the alert.

count (total)

Specifies the count or total alerts as an absolute value.

std. deviation

Calculates the normal deviation compared to the current alert.

ANY

Fires the alert when any of the following conditions are present.

ALL

Fires the alert when all of the following conditions are present.

NONE

Fires the alert when none of the following conditions are present.

The next drop-down menu specifies the time frame to collect the data. The next drop-down menu contains the following options:

>

Greater than

<

Less than

<=

Less than or equal to

>=

Greater than or equal to

==

Equal to

If applicable, specify a number to represent the percent of the trend, an absolute number, the number of trends per second, or the number of trends per minute, and select that choice from the final drop-down list. Selecting 100 percent of trend causes the value to overlap with the trend. If you want to alert on a condition 50 percent above the trend, then enter 150. If you want to alert on a condition 50 percent below the trend, then enter 50.

The **Trend Settings** tab contains the following fields:

Window

Specifies the calculation window for the trend.

Same Hour of Week

Calculates the trend within a specified 1-hour window each week.

Same Hour of Day

Calculates the trend within a specified 1-hour window each day.

Minute Rolling Average

Calculates the trend based on the average of the data gathered each minute within a specified amount of time from the present time.

Hour Rolling Average

Calculates the trend based on the average of the data gathered each hour within a specified amount of time from the present time.

Lookback

Specifies the number of minutes of lookback.

Weighting Model

Specifies the weighting model.

Mean

Specifies the manner in which to calculate the average.

Linear Average

Calculates the average with all data points weighted equally.

Single Exponential

Calculates the average with the most recent data points weighted more heavily.

Double Exponential

Calculates the average with the most recent data points weighted the most heavily.

For linear averages, the most recent value is weighted at 1 times the oldest value by default. For single and double exponential means, enter a number to weight the most recent value.

Percentile

Specifies the percentile value used as a basis for creating the trend.

Percentile

Records the trend using data points from a user-specified percentile.

Min Value

Records the lowest data point gathered during the time interval.

Max Value

Records the highest data point gathered during the time interval.

Regression

Linear

Calculates steadily increasing trends based on previous trends that are equally incremental.

2nd Degree Polynomial

Calculates exponentially accelerating trends by projecting a curve using the equation

```
y = ax2 + bx + c
```

Standard Deviation

Calculates the normal deviation compared to the current trend.

Type

Uses a sample-based or population-based standard deviation.

Normalization

Displays the standard deviation relative to mean.

Note: If a trend is a standard deviation, its associated alerts use the same parameters as the trend. If the trend is not a standard deviation, then the alert is calculated as "sample" and "absolute".

Static Value

Calculates a static value based on the number you enter, and is useful to plot constant lines for SLAs.

Time Delta

Uses the oldest trend, resulting in a time delta option based on the lookback window.

Trimean

Calculates the weighted average of the 25th, 50th, and 75th percentile values.

Winsorized Mean

Replaces the most outlying values with the highest and lowest remaining values. Values above the 90th percentile become the same value as the 90th and values below the 10th percentile become the same value as 10th.

The **Exclusion Intervals** tab shows all the defined exclusion intervals that can be applied to alerts. For more information, refer to Exclusion Intervals.

The **Notifications** tab contains the following fields:

Severity

Specifies the level of severity required to send email notifications. The color next to each option is reflected in the geomap(s) and summary widget(s) associated with this alert.

Send SNMP Trap

Specifies whether notifications are sent to an SNMP listener. Users with administration privileges can configure the SNMP listener in the Discover applianceAdministration UI.

Email notification groups

A list of defined email groups that can receive alert notifications. The Default group is checked by default. Users with administration privileges can configure additional groups in the Discover appliance Administration UI.

Additional email addresses

Specifies the email addresses that should receive notification when this alert fires.

Additional metrics in emails (one per line)

Specifies additional metrics to include in the notification email. Paste the metric names into the window or click the **Find metric...** button to search for a metric.

Find metric...

Enables you to select a metric from a list of all possible metrics.

The **Description** tab provides a space for an optional, user-defined description of the alert.

The **Assignments** tab contains the following items:

Note: Assigning trend alerts to more than 1000 devices may impact system performance

Assign to All

Specifies that the alert should be assigned to all devices, current as well as devices discovered in the future.

Assignments

Displays where in the system the alert has been manually assigned to a device or group. To manually disassociate the alert from a device or group, click the delete symbol next to the device or group name. (This field does not show when the alert was assigned by clicking the **Assign To All** checkbox.)

Remove All Assignments

Removes all manually-added devices and groups from the alert.

Exclusion intervals

Exclusion Intervals define a time in which alerts are suppressed. For example, if you do not want to be notified about alerts after hours or on the weekends, an exclusion interval can suppress the alert during that time period.

Create an exclusion interval

To configure an exclusion interval for an alert:

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. On the Alerts page, click the **Exclusion Intervals** tab.
4. On the toolbar, click **New**. The Exclusion Interval Configuration dialog box opens with the following tabs:

   Interval Settings

   Provides configuration settings to define the exclusion interval.

   History

   Contains a list of changes.

Copy an exclusion interval

To copy an exclusion interval:

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. On the Alerts page, click the **Exclusion Intervals** tab.
4. In the Exclusion Intervals table, click the checkbox next to the interval that you want to copy and use as a template for defining a new exclusion interval.
5. Click the **Copy** button.

The name of the copied exclusion interval is generated automatically by appending the word "(copy)" to the original name.

Assign an alert

Alerts can be assigned to applications, devices, and groups.

1. On the page for an application, device, or group, click the **Alerts** tab.
2. Click the add icon to the left of the **Alerts** field.
3. In the Assign Alerts dialog box, select the device alerts that you want to show in the network capture.
4. In the **Filter** text box, provide an optional filter string to filter the list of alerts by name.
5. Click **OK**.

Delete an exclusion interval

To delete exclusion intervals from every assigned alert:

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. On the Alerts page, click the **Exclusion Intervals** tab.
4. In the table of exclusion intervals, click the checkbox next to the interval (or intervals) that you want to remove from the system.
5. Click the **Delete** button.

View exclusion intervals

To view a list of exclusion intervals:

Note: Exclusion intervals created on the Command appliance automatically sync with all of its nodes, but exclusion intervals created on individual nodes do not sync with the Command appliance.

1. Click the **System Settings** icon in the top toolbar.
2. In the Settings pop-up window, click **Alerts**.
3. On the Alerts page, click the **Exclusion Intervals** tab.
4. Click the drop-down list and select one of the following options:

All Intervals

Displays exclusion intervals that were created on both the Command appliance and the node.

Command appliance Intervals

Displays exclusion intervals that were created on the Command appliance.

Local Intervals

Displays exclusion intervals that were created on the node.

The Exclusion Intervals table contains the following information:

Name

Specifies the name of the exclusion interval.

Description

Provides a space for an optional, user-defined description.

Command appliance

Specifies whether the exclusion interval was created on the Command appliance or locally on the node.

Type

Specifies the type of exclusion interval. Options include:

One-time

Specifies an exclusion period that occurs only once from a designated start time (date and time) to a designated end time (date and time).

Daily

Specifies an exclusion period that occurs every day from a designated starting hour to a designated ending hour.

Weekly

Specifies an exclusion period that occurs every week from a designated start time (day and time) to a designated end time (day and time).

Note: Time intervals are excluded in one-hour blocks.

Exclusion interval settings

You can view the following exclusion interval settings:

The **Interval Settings** tab contains the following fields:

Name

Enter a descriptive name for the new exclusion interval.

Description

Provides a space for an optional, user-defined description.

Assign to All

Click the **Alerts** checkbox, the **Trends** checkbox, or both checkboxes to assign this exclusion interval to all alerts and/or trends in the Discover appliance.

Exclude

Specifies the time frame for the exclusion interval.

From

Sets a one-time exclusion interval.

Every day

Sets a daily exclusion interval.

Every week from

Sets a weekly exclusion interval.

The **History** tab contains the following columns.

Change

Displays the change that was made to the exclusion interval.

Author

Displays the author of the change.

Timestamp

Displays when the change was made.

Trouble groups

The Discover appliance automatically generates trouble groups based on network traffic. Trouble groups represent a collection of devices that meet specific criteria indicating potential problems.

The Trouble Groups table includes the following group information:

Name

Specifies the name of the trouble group.

Count

Identifies the number of devices that belong to this group.

Refer to the specific trouble group sections for the criteria that defines that group.

View trouble groups

To view details about the devices in a trouble group:

1. Click **Alerts**.
2. Click **Trouble Groups**.
3. In the Name column, click the trouble group name to view the list of devices in the group.
4. On the device list page, click the device name to view device-level statistics.

   When you click a device name from this page, the Discover appliance Web UI redirects to the Devices page and opens the device statistics page.

Aborted HTTP/DB transactions

Aborted HTTP/DB transactions indicate a high level of aborts during active HTTP or database transactions. Aborts are generally initiated by clients, so this may indicate that the server hangs on the response or does not complete the response in a timely manner.

Criteria	Check for high levels of Requests Aborted or Responses Aborted
Devices	Devices that show HTTP or DB server activity and are not gateways or load balancers
Update	Hourly
Remedial Actions	For HTTP transactions, check for URLs that take along time to process. For database transactions, check for long-running stored procedures

ADC SNAT pool too small

ADC SNAT pool too small indicates that a connection failed to initiate because the current device interpreted the SYN as belonging to a previous connection.

Criteria	Check for any PAWS-Dropped-SYNs (In)
Devices	Known ADCs only (based on MAC address OID lookup)
Update	Hourly
Remedial Actions	On the BIG-IP Application Delivery Controller (ADC), the SNAT pool size should be increased

ADC TCP connection throttling

ADC TCP connection throttling indicates that the connections are stalling in the Application Delivery Controller (ADC) and it is unable to keep up with the rate of data sent.

Criteria	Check for Zero Windows (Out) as a factor of the number of established connections
Devices	Known ADCs only (based on MAC address OID lookup)
Update	Hourly

Remedial Actions	On the BIG-IP Application Delivery Controller (ADC), the proxy_buffer_high setting in the TCP profile should be increased

Database server backups

Database server backups are caused by backups taking place over CIFS, NFS, or Veritas on active database servers.

Criteria	Detect large amount of storage traffic exchanged from the server
Devices	Devices that show CIFS, NFS, or TCP port 13724 activity (Veritas) and are not gateways or load balancers
Update	Every 30 minutes
Remedial Actions	Throttle down backups and schedule them during times with lower traffic

DNS missing entries

DNS missing entries may indicate a service availability problem.

Criteria	Compare DNS NXDOMAINS responses with the total number of responses
Devices	Devices that show DNS server activity and are not gateways or load balancers
Update	Hourly
Remedial Actions	If these queries are intended, add an entry to DNS. If not, find the clients making erroneous DNS requests and configure them to stop making these requests

Excessive CIFS metadata queries

Excessive CIFS metadata queries indicate a high level of file metadata queries compared to read/write activity (or "goodput") on a CIFS server.

Criteria	Compare FSInfo to the number of Read and Write bytes
Devices	Devices that show CIFS server activity and are not gateways or load balancers
Update	Hourly
Remedial Actions	Check clients that generate large numbers of CIFS for configuration issues that would cause them to perform an overly high level of directory scans

Excessive HTTP authorizations

Excessive HTTP authorizations should be checked for large numbers of HTTP authorization errors, which may indicate break-in attempts.

Criteria	Check for 401 errors and compare them with the number of valid responses

Devices	Devices that show HTTP server activity and are not gateways or load balancers
Update	Hourly
Remedial Actions	Log these HTTP authorization errors, as they may be break-in attempts

HTTP broken links

HTTP broken links indicate that a resource has been moved or deleted but the document might still points to the old location.

Criteria	Check for 404s and compare it with the number of valid responses
Devices	Devices that show HTTP server activity and are not gateways or load balancers
Update	Hourly
Remedial Actions	Track down the source of 404s

Path MTU mismatch

Path MTU mismatch displays the list of devices for which path MTU mismatch was detected. These devices are not respecting the Fragmentation Needed ICMP announcements.

Criteria	Check for ICMP type 3 code 4
Devices	All devices
Update	Hourly
Remedial Actions	Check documentation for devices that are not respecting path MTU announcements for configuration options

Problematic TCP offloading engine

Problematic TCP offloading engine. Indicates that the current device is sending too much data resulting in network congestion and dropped packets. This behavior has been seen with a number of TCP offloading engines.

Criteria	Check for Bad Congestion Control (Out)
Devices	NICs known to have problems (based on MAC address OID lookup)
Update	Hourly
Remedial Actions	Turn off TCP offloading

Server TCP connection throttling

Server TCP connection throttling is caused by server running out of buffer or CPU resources and throttling network connections as a result.

Criteria	Check for the Zero Windows (Out) as a factor of the number of established connections
Devices	Devices that are servers and are not gateways or load balancers

Update	Every 30 minutes
Remedial Actions	Check buffer sizes and CPU, and increase those resources, if necessary

SPAN oversubscription

SPAN oversubscription indicates that data coming over the SPAN port is incomplete. This can happen to data being dropped at the SPAN port due to oversubscription or microbursts.

Criteria	Compare the desyncs to the number established connections
Devices	All devices
Update	Daily
Remedial Actions	Filter down data coming over the SPAN port or use a larger capacity SPAN port

SSL Key Size < 2048

SSL key size < 2048 indicates a 1024-bit SSL key. In 2010, 1024-bit public keys have been declared insecure by NIST. As a result, certificate authorities are moving to 2048-bit keys.

Criteria	Check for SSL public key size less than 2048 bits
Devices	Devices that show SSL server activity and are not gateways
Update	Hourly
Remedial Actions	Deploy 2048-bit keys in place of potentially insecure ones

Virtual packet loss

Virtual packet loss indicates that a virtual instance is overwhelmed and cannot send packets out in a timely fashion. TCP interprets delayed ACKs as packet loss and sends less data.

Criteria	Check for large numbers of RTOs coming from devices within virtualized environments
Devices	Virtualized devices (based on MAC address OID lookup)
Update	Hourly
Remedial Actions	Provide more hardware resources to stressed VMs

Reports

The Reports page displays a list of defined reports in the Discover appliance and provides controls to manage the reports. After you have created reports, you can add metric data from specific devices, applications, and networks to any report by clicking **Add to Report** on the page-level toolbar.

The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

View a report

1. Click **Systems Settings** in the top toolbar.
2. Click **Reports**.

Create a report

1. On the Reports page, click **New**.

   The Report Configuration window opens with the following fields:

   Report name

   Enter a descriptive name for the new report.

   Description

   Provides a space for an optional description of the new report.

   Report items

   When modifying the report later, click the **Delete** icon next to any items that you do not want to appear in the report.
2. (ExtraHop Command appliance only) Click the **Email Schedule** tab and complete the following fields:

   Schedule

   Select a radio button to specify the frequency with which to send the report.

   Email groups

   Select a predefined email group that should receive the report.

   Additional emails

   Enter individual email addresses not associated with a predefined email group.
3. (ExtraHop Command appliance only) Click the **Email Footers** tab and complete the following fields:

   Email Footers

   Select the radio button to include email footers.

   HTML Footer

   Enter your email footer using HTML markup.

   Text Footer

   Enter a text-only version of your email footer.
4. Click **OK** to save the new report.
5. Navigate to relevant sections of the Discover appliance Web UI, click **Add to Report**, and select the report you created to add items to the report.

Generate a test report

1. Click the report that you want to view.
2. In the Report Configuration dialog box, click the **View report for** drop-down list and select the time interval for specifying the scope of the capture data that is displayed in the report.
3. Click **Generate**.

Copy reports

1. In the Reports table, select the checkbox next to the report(s) that you want to copy and use as a template for defining a new report.
2. Click the **Copy** button.

The name of the copied report is generated automatically by appending the word "(copy)" to the original name.

Delete reports

1. In the Reports table, select the checkbox next to the report(s) that you want to delete.
2. Click the **Delete** button.

Customizing ExtraHop appliances

You can customize ExtraHop Appliances by creating custom groups, devices, pages, and metrics. You can also create new flex grids and geomaps. ExtraHop enables you to specify which devices are in full analysis by managing device limits. You can also apply tags to groups and devices to associate objects that share common characteristics.

Custom groups

A custom group is a user-defined group of devices on the Discover appliance. There are two types of custom groups: static and dynamic.

A static custom group is a user-defined grouping of devices. Once you create a static custom group, you must manually add devices to it. A new static custom group does not have any devices assigned to it. To populate this new group with devices, you must select the devices to add. For information about adding a device to a group, refer to Devices Page.

A dynamic custom group manages the collection of devices programmatically based on the criteria specified by the user. The criteria used to populate the dynamic group is a substring match. The substring can be a host name, IP address, MAC address, or any of the other defined device criteria. A new dynamic custom group is populated automatically with devices that match the device criteria.

For example, it is possible to define a dynamic custom group that includes all devices in which the host name contain the substring extrahop. For this rule, devices with names such as www.extrahop.com and extrahop.net match the criteria specified in the rule and are included in the dynamic group, if they are present on the network.

Create a static custom group

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the **Select Action** drop-down list and select **Add**.
4. In the Add Custom Group dialog box, in the **Name** text box, enter a name for the new static custom group.
5. For the **Group Type** option, select **Static**.
6. In the **Description** text box, add a brief description for the new custom group.
7. Click **OK**.
8. (Optional) Add a device to the group.
   a) Click **Metrics**.
   b) Click **Devices**.
   c) Select the checkbox next to the name of the device you want to add.
   d) Click the **Select Action** drop-down list and select **Add to Group**.

Create a dynamic custom group

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the **Select Action** drop-down list and select **Add**.
4. In the Add Custom Group dialog box, in the **Name** text box, enter a name for the new dynamic custom group.
5. For the **Group Type** option, select **Dynamic with criteria**.

6. Click the drop-down list and select one of the following options:

   ip address

   Specify the device IP address. The IP address criteria can include CIDR notation in IP address/subnet prefix length format. For example, 10.10.0.0/16 for IPv4 networks or2001:db8::/32 for IPv6 networks.

   name

   Specify the device name. Criteria can include the DHCP name, NETBIOS name, or DNS name.

   mac address

   Specify the device MAC address.

   tag

   Specify the user-defined device tag.

   type

   Specify a device type from the drop-down list. Criteria for each device type includes the following:

   Activity

   Includes the metric types that were active in the selected time interval. For example, the HTTP server metric a search for "http_server" returns devices with HTTP server metrics and any other device with the custom type set to http_server.

   Device type

   Includes Gateway, Firewall, Load Balancer, File Server, and Custom Device.

   Class

   Includes Node, Remote, Custom, and Pseudo.

   vendor

   Matches a substring in the device vendor name as determined by the MAC OID lookup.

   vlan

   Matches a substring in the device Virtual Local Area Network (VLAN) tag. VLAN information is extracted from VLAN tags, if the traffic mirroring process preserves them on the mirror port.

7. In the text box, enter the characters that you want to use for the substring match.

   If the search string value starts and ends with a forward slash (/), the portion of the input between the slashes is interpreted as a regular expression. The regular expression must be written in PostgreSQL syntax. Refer to PostgreSQL documentation for more information.

   If you are using the criteria type, select a type from the drop-down list.

8. In the **Description** text box, add a brief description for the new custom group.
9. Click **OK**.

Managing custom groups

You can manage custom groups on your Extrahop appliance.

View a custom group on a Command appliance

1. Click **Metrics**.
2. Click **Custom Groups**.
3. On the Custom Groups page, click the **All Groups** drop-down list and select one of the following options:

   All Groups

   Displays custom groups that were created on both the Command appliance and the node.

   Command appliance Groups

   Displays custom groups that were created on the Command appliance.

Local Groups

Displays custom groups that were created on the node.

Note: Custom groups created on the Command appliance automatically sync with all of its nodes, but custom groups created on individual nodes do not sync with the Command appliance.

4. In the Custom Groups table, click the custom group that you want to view.

View a custom group on a Discover appliance

1. Click **Metrics**.
2. Click **Custom Groups**.
3. In the Custom Groups table, click the custom group that you want to view.

Modify a custom group name

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a custom group.
4. On the Custom Group page, click the edit icon to the right of the **Name** field.
5. In the text area, enter a new name for the custom group.
6. Click **OK**.

Modify custom group criteria

To modify the criteria of a dynamic custom group:

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a dynamic custom group.
4. On the Custom Group page, click the edit icon to the right of the **Criteria** field.
5. Click the **Criteria** drop-down list and select one of the following options:

any

Matches a substring in any device element.

ip address

Matches a substring in the device IP address. The IP address criteria can include CIDR notation in IP address/subnet prefix length format. For example, 10.10.0.0/16 for IPv4 networks or 2001:db8::/32 for IPv6 networks.

name

Matches a substring in the device name.

node

(Command appliance only) Matches a substring in the node name.

mac address

Matches a substring in the device MAC address.

tag

Matches a substring in the user-defined device tag.

type

Matches a substring to a specified device attribute type. When you select **type**, the **Find** text box becomes a drop-down list. In the **Find** drop-down list, select from the following:

Activity

Includes the metric types that were active in the selected time interval. For example, selecting **HTTP Server** returns devices with HTTP server metrics, and any other device with the custom type set to **HTTP Server**.

Device type

Includes Gateway, Firewall, Load Balancer, File Server, and Custom Device.

Class

Includes Node, Remote, Custom, and Pseudo.

vendor

Matches a substring in the device vendor name as determined by the Organizationally Unique Identifier (OUI) lookup.

vlan

Matches a substring in the device Virtual Local Area Network (VLAN) tag. VLAN information is extracted from VLAN tags, if the traffic mirroring process preserves them on the mirror port.

6. In the text box, enter the characters that you want to use for the substring match.

   If the search string value starts and ends with a forward slash (/), the portion of the input between the slashes is interpreted as a regular expression. The regular expression must be written in PostgreSQL syntax. Refer to PostgreSQL documentation for more information.
7. Click **OK**.

Modify a custom group description

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a custom group.
4. On the Custom Group page, click the edit icon to the right of the **Description** field.
5. In the text area, enter a description for the custom group.
6. Click **OK**.

Assign an alert

Alerts can be assigned to applications, devices, and groups.

1. On the page for an application, device, or group, click the **Alerts** tab.
2. Click the add icon to the left of the **Alerts** field.
3. In the Assign Alerts dialog box, select the device alerts that you want to show in the network capture.
4. In the **Filter** text box, provide an optional filter string to filter the list of alerts by name.
5. Click **OK**.

Assign an alert to a custom group

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a custom group.
4. On the Custom Group page, click the **Alerts** tab.
5. Click the add icon to the left of the **Alerts** field.
6. In the Assign Alerts dialog box, select the custom group alert that you want to show in the network capture.
7. Click **OK**.

Remove an alert from a custom group

1. Click **Metrics**.

2. Click **Custom Groups**.
3. Click the name of a custom group.
4. On the Custom Group page, click the **Alerts** tab.
5. Click the delete icon to the left of the alert that you want to remove.

Assign a trigger to a custom group

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a custom group.
4. On the Custom Group page, click the **Triggers** tab.
5. Click the **+** icon next to Triggers.
6. Select the checkbox next to the trigger you want to assign to the group.
7. Click **OK**.

Assign a custom group to a geomap

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a custom group.
4. Go to the Custom Group page and click the **Geomaps** tab.
5. Click the **+** icon next to Geomaps and select the checkbox next to the geomap(s) you want to associate with the custom group.
6. Click **OK**.

Remove a custom group from a geomap

1. Click **Metrics**.
2. Click **Custom Groups**.
3. Click the name of a custom group.
4. Go to the Custom Group page and click the **Geomaps** tab.
5. Click the delete icon to the left of the geomap that you want to remove.

Custom devices

By default, all IP addresses outside the locally-monitored broadcast domains are aggregated at one of the incoming routers. To identify the devices behind these routers, you can create custom devices to enable reporting on these devices.

Custom devices in version 5.0 take the place of pseudo devices, which were configured through the Discover appliance Administration UI in pre-4.0 versions. ExtraHop recommends using custom devices to track traffic associated with remote IP addresses or subnets in version 5.0. Unlike pseudo devices, you do not need Administrator privileges to configure custom devices.

If you have created pseudo devices in pre-4.0 versions, they will remain on your Discover appliance until you migrate them to custom devices.

Create custom device

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Custom Devices**.
3. On the Custom Device Settings page, click **New**.

   The Custom Device Configuration dialog box opens.

4. Complete the following fields:

   Name

   Specifies the name assigned to the custom device.

   ExtraHop ID

   Specifies the unique name of the custom device.

   Author

   Specifies the creator of the custom device.

   Description

   Provides a space for an optional, user-defined comment.

   Match Criteria

   Provides a way to group all devices that match user-defined criteria.
5. Click the **Add Criteria** button to add a specific IP address, L4 port, or VLAN.

The custom device appears in the Custom Devices table. When you search for this device on the All Devices page, the device appears with an icon denoting that it is a custom device.

Delete custom device

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Custom Devices**.
3. On the Custom Device Settings page, select the checkbox next to the custom device(s) that you want to delete.
4. Click **Delete**.

Enable a custom device

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Custom Devices**.
3. On the Custom Device Settings page, select the checkbox next to the custom device(s) that you want to enable.
4. Click **Enable**.

Disable a custom device

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Custom Devices**.
3. On the Custom Device Settings page, select the checkbox next to the custom device(s) that you want to disable.
4. Click **Disable**.

The selected custom devices will become inactive and not appear in the left panel.

Migrate pseudo devices to custom devices

If you have created pseudo devices in pre-4.0 versions, they will remain on your Discover appliance until you migrate them to custom devices. If pseudo devices are in use and working effectively, there is no requirement to migrate them to custom devices.

You should consider migrating pseudo devices to custom devices if:

- Router devices must be able to report their full compliment of traffic. (Pseudo devices based on traffic flowing through router devices causes those statistics to be subtracted from the counts of the router device.)

- Pseudo device configurations must be updated regularly and the more streamlined process of maintaining custom devices is preferred.
- Custom devices for remote subnets need to be created by users who are not ExtraHop administrators.
- Definitions of remote subnet devices need to be more granular than a simple IP subnet definition. (For example, the definition should be based on port number or VLAN.)
- The work of migrating from pseudo devices to custom devices outlined below is acceptable.

1. Define new custom devices matching those currently defined for existing pseudo devices.
2. Remove the old pseudo devices.

   Note: When performing steps 1 and 2 above, you should consider the appliance's licensed device limitations. For up to 24 hours, both devices for a remote subnet (the custom device and the pseudo device) will be considered active and count against device limits. When appliances are operating near their maximum licensed device count, it may be desirable to perform the migration in multiple stages to ensure continuous coverage.
3. Existing pseudo devices may have some of the associations in the following list. Update the definition or assignment to reference the new custom devices instead of the pseudo devices.

   - Group definitions
   - Device tags
   - Alert assignments
   - Custom page assignments
   - Flex grid assignments
   - Geomap assignments
   - Trigger assignments

View custom devices

To view custom devices:

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Custom Devices**.

The Custom Devices table provides the following information about custom devices:

Name
Identifies the name assigned to the custom device.

ExtraHop ID
Identifies the unique name of the custom device.

Author
Specifies the creator of the custom device.

Description
Provides a space for an optional, user-defined description.

Status
Specifies whether the device is enabled or disabled.

Device limits

The Discover appliance can discover, monitor, and analyze a large number of devices in your network. A device limit ensures that the Discover appliance operates efficiently when there are too many devices.

The device limit for your appliance is determined by the license you acquired. The device limit is the total number of devices that can be in full analysis. Full analysis means thatDiscover appliance will collect complete metrics for that device.

Note: Pseudo devices and custom devices are included in the device count for full analysis.

Limited Analysis

If the device limit is exceeded, and there are too many devices for the Discover appliance to fully analyze, some devices will be placed in limited analysis. Only L2 and L3 metrics will be collected for devices in limited analysis.

In this scenario, the Discover appliance determines which devices will be in full analysis and limited analysis. Devices that are currently active and were discovered earliest get a higher priority for receiving full analysis.

If you have a high-priority device that you want to make sure does not fall into limited analysis, you can add that device to a whitelist and ensure that it will remain in full analysis. For more information, see the Add a Device to the Whitelist section.

View Device Limit Details

The Device Limits page provides the following information about device counts:

Device Safety Limit
The total number of devices the Discover appliance can fully analyze.

Active Devices in Limited Analysis
The total number of devices being analyzed minus the device safety limit. When the number of devices is less than the device safety limit, this number is 0.

Devices in Whitelist
The number of devices that were manually added to the whitelist. All devices in the whitelist are designated to receive full analysis regardless of activity, but a device may not appear in the device list if it has not received traffic within the specified time interval.

The Device Limits page provides the following information about each device:

Name
Specifies the primary name the device.

MAC Address
Specifies the unique identifier of the device network interface.

VLAN
Specifies the VLAN for the device.

IP Address
Specifies the primary IP address the device uses to communicate on the network.

Discovery Time
The time the device was first discovered on the Discover appliance.

Description
Provides a space for an optional, user-defined description.

View device limits

1. Click **System Settings** in the top toolbar.
2. Click **Device Limits**.
3. Click the drop-down list to view a list of devices.

   Full Analysis
   The total set of devices receiving full analysis, including devices that were manually added to the whitelist and active devices receiving full analysis before the device safety limit is reached.

Limited Analysis

Devices receiving limited analysis after the device safety limit is reached.

Whitelist

Devices that were manually added to the whitelist.

Eligible for Licensing

Devices that contribute to the licensed server count, including devices receiving full and limited analysis that are not L2-only devices.

Note: It is possible to blacklist devices based on their unique MAC addresses by modifying the Discover appliance's running config file. Contact your administrator to blacklist devices.

Add a device to the whitelist

1. On the Device Limits page, select one or more devices from the list.
2. Click **Add to Whitelist**.

Tip: An efficient way to add devices to the whitelist is to search for servers with a certain type of activity. For instance, search by **Type** and click the drop-down arrow in the menu to select an activity type. For example, select **Activity: HTTP Servers** to find all servers with HTTP activity, and add them allto the whitelist. Other activity types include `db_server`, `dns_server`, `ldap_server`, `cifs_server`, `nfs_server`, and more.

Remove a device from the whitelist

1. On the Device Limits page, select one or several devices from the list.
2. Click **Remove from Whitelist**.

Bundles

Bundles enable you to create and save a set of system customizations for the Discover appliance or ExtraHop Command appliance. The following system customizations can be saved as part of a bundle:

- Alerts
- Applications
- Custom pages
- Dashboards
- Dynamic groups
- Flex grids
- Geomaps
- Triggers

After creating a bundle, you can download the file in .json format and share the file with other users.

Essentials bundle

Along with the built-in Activity and Network dashboards, the ExtraHop system ships with the Essentials bundle. This bundle provides a set of customizations that are designed to readily display common and related network metrics through a series of Essentials dashboards.

Although the Essentials bundle is pre-installed, you must apply the bundle before you can view the Essentials dashboards. In addition, some of the dashboards require that you enable triggers on the system to view all of the metrics in the pre-defined widgets.

Note: The Essentials bundle is designed to have minimal impact to your system, but you should always exercise caution when enabling a trigger.

Apply the essentials bundle

1. Click the System Settings icon, and then click **Bundles**.
2. From the list of bundles, select **Essentials**.
3. Select the checkbox in the lower-right corner of the window to apply included assignments.
4. Click **Apply**, and then click **OK**.
5. (Optional) To view the Essentials dashboards, click **Dashboards**.
   The dashboards are listed in the left pane, under My Dashboards.

Enable triggers for the Essentials bundle

Note: The dashboards for encryption and DNS require that you enable two triggers that ship with the Essentials bundle.

1. Click the System Settings icon, and then click **Triggers**.
2. From the list of triggers, select **AAAA detection on IPV4 networks** and **Encryption Auditing Trigger (Application)**.
3. Click **Enable**.

After these triggers are enabled, your network traffic must be processed before the metrics in the dashboards display any data.

Create a bundle

1. On the Bundles page, click **New**.
2. Complete the following information in the Bundle Settings window:

   Name

   Assign a name to the bundle.

   Author

   Specify the creator of the bundle. Bundles loaded by default have the author "ExtraHop."

   Required Version

   Specify a required version for this bundle. If you try to import a bundle that requires a newer firmware version, a warning message displays in the Actions section of the Bundle Settings window.

   Contents

   Select the system customizations that you want to add to the bundle. Click the arrow to expand the list of available items.

   Description (Optional)

   Type a description about the bundle.
3. Click **OK** to save the bundle.

Modify a bundle

1. On the Bundles page, click the name of a bundle.
   In the View Bundle dialog box, the Actions section is now present. This section appears only after the bundle has been saved to the list.
2. (Optional) To download the bundle to your work station in .json file format, click the **Download** button.
3. (Optional) To choose how to handle imported objects with names that match existing objects, click the Existing Objects drop-down list.
   - To exclude the object from the import, select **Skip**.
   - To modify the preexisting object to match the data of the object being imported, select **Overwrite**.
4. Click the **Edit Raw Data** button to view and edit the source code.

This text box is blank until the bundle has been added to the list.

Note: The bundle preserves the version of the customizations at the time of creation. If you modify an object after creating a bundle, the object inside the bundle remains unchanged. If you want the bundle to contain the modified object, you must create another bundle.

Upload a bundle

1. On the Bundles page, click the **Upload** button.
2. The Load Bundle dialog box appears.
3. In the Load Bundle dialog box, do one of the following:
   - Paste the bundle data directly into the Load Bundle window.
   - Click the **Choose File** button to upload a saved bundle in .json file format.
4. Click the **Upload** button.

Apply a bundle

1. On the Bundles page, click on the name of a bundle.
2. In the Actions section, click the **Apply** button to enable the bundle and see the Bundle Import Status dialog box.

   This dialog box tells you whether the bundle was applied. If you selected **Skip** from the **Existing objects** drop-down list, the dialog box includes a list of skipped objects.

Delete a bundle

1. On the Bundles page, select the checkbox next to the bundle(s) that you want to delete.
2. Click the **Delete** button.

Device tags

Device tags are user-defined annotations that can be attached to a device or a group of devices that share common characteristics. The Device page provides the ability to create and apply device tags to specific devices. The Device Tags page lists all device tags and provides the ability to manage the device tags associated with those devices.

The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

Assign device tags

1. Click **Metrics**.
2. Click **Devices**.
3. In the table, select the device(s) that you want to tag.
4. Click on the name of the device at the top of the left pane.
5. Click the **Tags** tab.
6. Click the green plus icon next to Tag.
   An Add Device Tag window will appear.
7. Type the name of the tag you want to assign to your device.
8. Click **OK**.

Modify device tags

To edit the properties of an existing device tag:

1. Click **System Settings** in the top toolbar.
2. On the Settings page, click **Device Tags**.
3. Click the device tag name.
4. The Tag Properties window opens.
5. In the **Name** field, enter a new name for the device tag.
6. (Optional) To remove device associations from the tag, in the Devices list box, find the device that you want to disassociate from the tag.
7. Click the delete icon.
8. Click **OK**.

Remove device tags

1. Click **Metrics**.
2. Click **Devices**.
3. In the table, click the device name to open its Device page.
4. In the side pane, click the name of the device.
5. Click the **Tags** tab and find the tag that you want to remove in the list of tags.
6. Click the delete icon next to the tag name.

Delete device tags

1. Click **System Settings** in the top toolbar.
2. On the navigation bar, click **Settings**.
3. On the Settings page, click **Device Tags**.
4. On the Device Tags page, click the checkbox next to the device tag name that you want to delete.
5. Click **Delete**.

View device tags

1. Click **System Settings** in the top toolbar.
2. Click **Device Tags**.

Flex grids

Flexible grids (or flex grids) provide a way to create versatile reporting summaries of user-specified metrics across devices, groups, and applications. Flex grids facilitate reporting across remote sites, multiple applications, and tiers of a single application.

Create a flex grid

1. Click the System Settings icon in the top toolbar.
2. Click **Flex Grids**.
3. On the Flex Grids page, click **New**.
4. In the Configure Flex Grid window, in the **Name** field, type a name for the flex grid.
5. Select the **Object Type**.
6. Click the **Add** button to add a metric to be analyzed.
7. In the Select Metric window, select the metric(s) to add.
   The list is populated according to the selected object type.
8. Enter an optional description.
9. Click **OK**.

Assign an object to a flex grid

Devices, groups, and applications can be assigned to a flex grid.

1. Navigate to a page for an object that you want to assign to the grid.
2. In the side pane, click the name of the object.
3. Click the **Flex Grids** tab.
4. Click the **+** button next to **Flex Grids** to add a previously defined flex grid to this object.
5. In the Assign to Flex Grids window, select the checkbox next to one or more flex grids and click **OK**.

Copy a flex grid

1. Click the System Settings icon in the top toolbar.
2. Click **Flex Grids**.
3. On the Flex Grids page, select one or more flex grids from the list.
4. Click **Copy** to create another grid to use as a template for defining a new flex grid.
   The name of the copied grid is generated automatically by appending the word "(copy)" to the original name. This action copies metrics only, not devices, groups, or applications.

Delete a flex grid

1. Click the System Settings icon in the top toolbar.
2. Click **Flex Grids**.
3. On the Flex Grids page, select one or more flex grids from the list.
4. Click **Delete**.

Add a metric to a flex grid

1. Click the System Settings icon in the top toolbar.
2. Click **Flex Grids**.
3. On the Flex Grids page, click the name of a flex grid.
4. Click **Edit**.
5. In the Configure Flex Grid window, click the **Add** button.
6. In the Select Metric window, select the metric to add.
   The list is populated according to the selected object type.
7. Click **OK** to add the metric.
8. Click **OK** to save the changes and exit the Configure Flex Grid window.

Remove a metric from a flex grid

1. Click the System Settings icon in the top toolbar.
2. Click **Flex Grids**.
3. On the Flex Grids page, click the name of a flex grid.
4. Click **Edit**.
5. In the Configure Flex Grid window, select the checkbox next to one or more metrics in the list.
6. Click **Remove**.
7. Click **OK** to save the changes and exit the Configure Flex Grid window.

View flex grids

To view flex grids:

1. Click the System Settings icon in the top toolbar.
2. Click **Flex Grids**.

3. On the Flex Grids page, click a grid in the list to view the capture details.
   The flex grid details page appears. The grid displays the service status for applicable detailed metrics in red, orange, yellow, or green based on the severity of configured alerts. If no alerts are configured for the specified metric, it has no service status.
4. Click a metric in the flex grid table to view a graph of the specific metric activities of that object over time.
5. Click the drop-down list and select **Availability** to view the information in a bar graph.
6. Click the chart to go to the object sub-page that contains the metric.

The Flex Grids page provides a list of flex grids and the following information:

Name

Specifies the name assigned to the flex grid.

Object Type

Specifies the object type.

Description

Specifies the user-defined description applied to the flex grid.

Geomaps

Geomaps show worldwide activity and metrics based on the translation of an IP address or an origin address to a geographical location. You can create multiple geomaps based on any metric recorded by the ExtraHop system and associate the geomaps with any device, group, or application.

View a geomap

To view a geomap:

1. Navigate to the page for a device, group, or application.
2. Click **Geomaps**.
3. Click the name of a geomap.

You can view details about a specific dot on the geomap by clicking on it.

A user-defined, detailed alert can be associated with a device, group, or application on which the geomap is rendered. When you configure an alert for a specific metric, any alert of the same metric type will appear on the geomap. For example, if you configured an alert to fire on HTTP responses, but you configured alerts for HTTP errors and response times as well, all three metrics will appear on the geomap. When an alert fires, the dot on the geomap associated with that alert is colored based on the severity level set in the alert. If multiple alerts fire on the same location, the color of the dot reflects the most severe alert. For more information about severity levels, see Alert Settings.

The resulting colors reflected in the geomaps are as follows.

Gray

No alerts, or only edge-triggered alerts are configured.

Green

No alerts, or only alerts with a severity level of `Debug` and `Informational` have fired.

Orange

At least one alert with a severity level of `Notice` or `Warning` has fired.

Red with spinning edges

At least one alert with a severity level of `Error` or `Critical` has fired.

Red with sonar beacons

At least one alert with a severity level of `Emergency` or `Alert` has fired.

The Firing Mode setting of the associated alert may affect the behavior of the geomap. For more information, refer to Alert Settings.

- If the alert is set to Edge-triggered, the alert fires once when the threshold is breached and will fire again only when the metric goes below the threshold and breaches it again. For persistent problems (i.e., the threshold is breached continuously), a red dot appears on the geomap when the problem first occurs, but not continuously if the condition is only breached once.
- If the alert is set to Level-trigged, the alert fires continuously while the condition is breached at the specified re-firing interval.

ExtraHop recommends configuring level-triggered alerts that re-fire at the same interval or more frequently as defined in the geomap, which is 30 minutes by default.

Geomap details include the following:

Interval

Click the **Interval** drop-down list to select the time interval for which metrics are displayed. You can specify **Last 5 minutes**, **Last 30 minutes**, **Last 6 hours**, **Last 24 hours**, or **Last 7 days**.

Region

Click the Region drop-down list to view alerts in a specific country. The default setting is **World**.

Map

Click the **Map** drop-down list to change the visual appearance of the geomap.

Color

Click the **Color** drop-down list to change the color of the geomap.

Per-County

View a snapshot of the traffic in specific countries during the selected time interval.

Autopilot

In the **Autopilot** box, click the **Start** button. A set of crosshairs appears and displays data about a city represented by a dot on the map. The Region Details dialog box appears, and if the city contains IP addresses that have alerts with a severity level of critical or higher, the Alert Details dialog box appears also. The autopilot feature automatically navigates between the top eight cities that have the highest values of the metric configured for this geomap.

Click the **Stop** button and click a dot to display data about a specific city manually.

Click the **Next** button to display data about the next city with high-volume traffic.

Updater

Contains a timer that counts down to the next automatic update. The updater's countdown changes depending on the time selected in the **Interval** drop-down list.

- If **5 minutes** is selected, the geomap updates once every 30 seconds.
- If **30 minutes** is selected, the geomap updates every 5 minutes.
- If **6 hours**, **24 hours**, or **7 days** is selected, the geomap updates every 60 minutes.

If a user-defined time window was created using the **Interval** drop-down list in the Discover appliance UI, this item appears in the **Interval** drop-down list in the geomap as well.

- If the user-defined interval is less than 5 minutes, the geomap updates once every 30 seconds.
- If the user-defined interval is less than 60 minutes, the geomap updates once every 5 minutes.
- If the user-defined interval is 60 minutes or greater, the geomap updates once every 60 minutes.

Click the **Pause** button to pause the timer.

Click the **Update** button to manually receive an update.

Alert Details

Contains the IP addresses that have notifications and their severity level. The Alert Details dialog box appears only for cities containing IP addresses that have alerts with a severity level of critical or higher.

Region Details

Contains information related to the specific metrics analyzed in the geomap and user activity in that region.

Create a geomap

1. Click the System Settings icon in the top toolbar.
2. In the Settings pop-up window, click **Geomaps**.
3. On the Geomaps page, click the **New** button.

   The Geomap Configuration dialog box opens with the following tabs:

   Geomap Settings

   Provides configuration settings to define the geomap name, metrics, and appearance.

   Description

   Provides a space for an optional, user-defined description of the geomap.

   Assignments

   Displays where in the system the geomap has been manually assigned to a device or group.
4. Select geomap configuration information.
5. Click **OK**.

Copy a geomap

1. Click the System Settings icon in the top toolbar.
2. In the Settings pop-up window, click **Geomaps**.
3. On the Geomaps page, select the checkbox next to the geomap(s) that you want to copy and use as a template for defining a new geomap.
4. Click **Copy**.
   The name of the copied geomap is generated automatically by appending the word "(copy)" to the original name.

Assign to geomap

Applications, devices, and groups can be assigned to a geomap.

1. Navigate to a page for an application, device, or group.
2. Click the **Geomaps** tab.
3. Click the **+** button next to **Geomaps**.
4. Select a geomap, and then click **OK**.

Delete a geomap

1. Click the System Settings icon in the top toolbar.
2. In the Settings pop-up window, click **Geomaps**.
3. In the Geomaps table, select the checkbox next to the geomap(s) that you want to delete.
4. Click **Delete**.

View geomap settings

1. Click the System Settings icon in the top toolbar.

2. In the Settings pop-up window, click **Geomaps**.

   The Geomaps table provides the following information:

   Name

   Specifies the name assigned to the geomap.

   Type

   Specifies whether a device, group, or application is being analyzed in the geomap.

   Metric

   Specifies the metric being analyzed in the geomap.

   Description

   Provides a space for an optional, user-defined description.

3. Click the name of a geomap to view geomap configuration settings.

The Geomap Settings tab contains the following fields:

Name

Specifies the name assigned to the geomap.

Metric

Specifies the application or device metric being analyzed in the geomap. To select the metric associated with the geomap, click the gear icon to the right of the **Metric** field.

Color

Specifies the color scheme of the geomap.

Map Style

Specifies the visual appearance of the geomap.

Region

Specifies which part of the world to view metrics.

Interval

Specifies the time interval for the geomap to refresh.

The **Description** tab provides a space for an optional, user-defined description of the geomap.

The **Assignments** tab contains the following fields:

Assign to All

Specifies that the geomap should be assigned to all devices, current as well as devices discovered in the future.

Assignments

Displays where in the system the geomap has been manually assigned to a device or group. To manually disassociate the geomap from a device or group, click the delete symbol next to the device or group name. This field does not show if the geomap is currently assigned to all.

Remove All Assignments

Removes all manually-added devices and groups from the alert. This field does not show if the geomap is currently assigned to all.

Custom metrics

In addition to built-in protocol metrics, you can build your own custom metrics. Learn more by visiting the following sections:

Triggers

Create a custom metric by designing a trigger, which is a user-defined script. For more information, see the Triggers section.

Metric Catalog

View and manage custom metrics in the Metric Catalog. For more information, see the Metric Catalog section.

Triggers

Application Inspection Triggers are user-defined scripts that perform additional actions during well-defined events, such as:

- HTTP requests or responses
- Database requests or responses
- CIFS or NFS requests or responses

When you create a trigger, the Discover appliance saves the custom metrics within that trigger to the Metric Catalog. In the Metric Catalog, you can rename the metric, add datatypes, and change the description. For more information, see the Metric Catalog section.

You can use a trigger to record a custom metric and display it in custom pages. For more information about using custom pages to display metrics collected from triggers, see the Custom Pages section.

For more information about writing triggers, see the following guides:

- Getting Started with Application Inspection Triggers
- ExtraHop Trigger API

The Triggers page displays the complete list of triggers defined in the Discover appliance. From the Triggers page, you can view and modify trigger properties, define new triggers, copy existing triggers, and delete triggers from the system.

The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

The Triggers page provides the following information about each trigger:

Name

Specifies the name assigned to the trigger.

Author

Specifies the creator of the trigger. Triggers loaded by default have the author, "ExtraHop".

Events

Specifies the event on which the trigger will fire.

Type

Specifies the type of entity being analyzed in the trigger.

Debug Mode

Specifies whether debugging is enabled. If debugging is enabled, debug information will appear in the Runtime Log tab of the trigger.

Command appliance

Specifies whether the trigger was created on the Command appliance or locally on the node. This column is only available on a node.

Description

Provides a space for an optional, user-defined description.

Status

Shows whether the triggers are assigned to devices and the number of assignments.

Create a trigger

1. On the Triggers page, click the **New** button.

   The Trigger Configuration dialog box contains the following tabs:

Configuration

Provides configuration settings to define the trigger.

Editor

Provides a space to enter and edit the source code.

Assignments

Displays where in the system the trigger has been assigned to an object.

Runtime Log

Displays related debugging log information.

Performance

Displays performance metrics generated by the trigger.

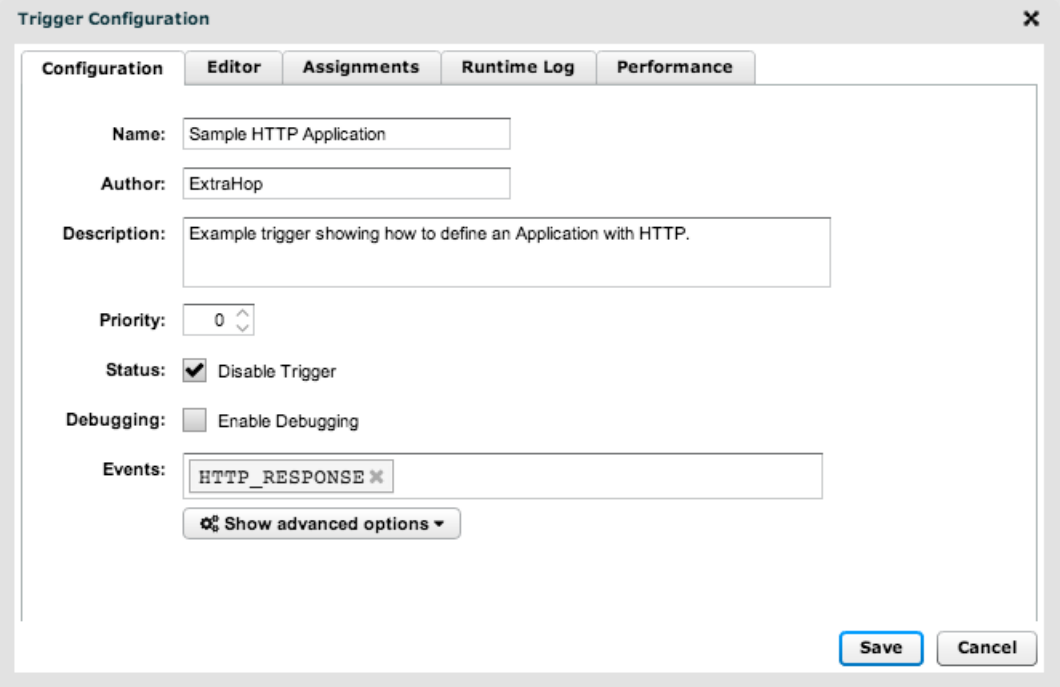

2. You must provide information in the **Trigger Configuration** tab to save the trigger.
   After you click **Save**, all five tabs appear in the Trigger Configuration window.
3. When you are finished entering the trigger configuration information, click **Save** to save the trigger and click **Cancel** to exit the Trigger Configuration window.

Copy a trigger

1. In the Triggers table, select the checkbox next to the trigger(s) that you want to copy and use as a template for defining a new trigger.
2. Click **Copy**.
   The name of the copied trigger is generated automatically by appending the word "(copy)" to the original name.

Enable a trigger

1. In the Triggers table, select the checkbox next to the trigger(s) that you want to enable.
2. Click **Enable**.

Disable a trigger

1. In the Triggers table, select the checkbox next to the trigger(s) that you want to disable.
2. Click **Disable**.

Delete a trigger

1. In the Triggers table, select the checkbox next to the trigger(s) that you want to delete.
2. Click **Delete**.

View triggers

1. Click the **Systems Settings** icon in the top toolbar.
2. Click **Triggers**.
3. (Optional) If you are viewing triggers on an ExtraHop Discover node, you can select one of the following option to filter which triggers are displayed:

   All Triggers

   Displays triggers that were created on both the Command appliance and the node.

   Command appliance Triggers

   Displays triggers that were created on the Command appliance.

   Local Triggers

   Displays triggers that were created on the node.

   Note: Triggers created on the Command appliance automatically sync with all of its nodes, but triggers created on individual nodes do not sync with the Command appliance.

Assign a trigger

Devices and custom groups can be assigned to a trigger.

1. Navigate to a page for a device or custom group.
2. Click the **Triggers** tab.
3. Click the **+** button next to **Triggers**.
4. Select a trigger, and then click **OK**.

View a custom metric

To view a custom metric, use the Metric Catalog.

1. Click **System Settings** in the top toolbar.
2. In the Settings pop-up window, click **Metric Catalog**.

The Metric Catalog page allows you to find an existing metric. You can sort the results by name or discovery time.

The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

When you select a built-in metric, information about that metric displays in the right pane.

Metric catalog

The Metric Catalog enables you to view information about built-in and custom metrics in the ExtraHop system. You also can delete and edit custom metrics through the Metric Catalogs.

With the Metric Catalog, you can search for metrics in the filter text box. The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

When you select a metric, information about that metric displays in the right pane in several sections.

Note: This information is useful for writing API queries and metric variables in a text box widget. For more information, see the Configure a Text Box Widget in Markdown section.

In the Parameter section provides information about the selected metric:

Source

Specifies the source of the metric as **Builtin**, which is a built-in default metric, or **Trigger**, which is a custom metric created by a user.

Metric

Specifies the API parameter name of the selected metric.

Object Type

Specifies the object (application, device, or network) associated with the metric.

Data Type

Specifies the data type associated with the metric.

Type

Specifies whether the metric is a base (top-level) metric or a detail metric.

The Display section provides information about how the selected metric will be displayed in the Web UI:

Name

Specifies the display name of the metric.

Units

Specifies the unit for the metric.

Description

Specifies a description for the metric.

In the Detail Relationships section provides information about whether there are base (top-level) metrics or detail metrics linked to the selected metric.

The REST API Parameter section provides an example of a JSON query structure for the selected metric with API parameters.

For more information, see the Manually Manage Custom Metrics section.

Manually manage custom metrics

When you manually manage a custom metric in the Metric Catalog, you must create a source for it using triggers before you are able to use it.

Note: It is best practice to create a trigger first so the custom metric is created automatically when you save the trigger. If you still want to create the custom metric manually, ensure that each value in the Parameters section matches its corresponding value in the trigger. If the values do not match, then the associated charts will not show data for the custom metric.

There are two options for manually managing custom metrics in the Metric Catalog: configure and delete a custom metric.

Configure a custom metric

Click the command menu next to the **Type to filter** field and select **Create Metric Manually**.

The Parameters section contains the following editable fields:

Source

Specifies the source of the metric, such as "Builtin" or "Trigger."

Builtin

The metric is a built-in, default metric.

Trigger

The metric is a custom metric created by a user. The source of the metric is a custom trigger created by a user. If you are creating a metric manually, this is the setting of the source field.

Metric

Type the name of the metric as listed in the source.

Object Type
Select the object (application, device, or network) associated with the metric.

Data Type
Select the data type associated with the metric.

Units
Select the unit of measure.

Type
Select the radio button next to the type of metric.

Name
Specify the name of the metric.

Units
Select the units for the metric.

Description
Specify a description for the metric.

Detail Relationship: (Optional)
Specify a detail or base metric to add to the metric.

Delete a single custom metric

1. Search for the metric and select it in the list.
2. Click command menu next to the **Type to filter** field and select **Delete Selected Metric**.

Delete multiple custom metrics

1. Search for a common term shared by the custom metrics you want to delete.
   Built-in metrics are excluded from the list.
2. Click the command menu next to the **Type to filter** field.
3. Select the **Custom Metrics Only** checkbox.
4. Select **Delete All Matching Metrics**.
   You can delete up to 1,000 metrics that match the search term even if they are not on the current page.
5. Click **Delete x Metrics** to confirm their deletion.

Custom pages

The Pages page displays a list of configured custom pages and provides controls to manage these pages. Custom pages are user-defined pages that display custom metrics recorded by user-defined triggers or by built-in metrics using flexible charting tools.

The Filter text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

The Pages table provides the following information about custom pages:

Name
Identifies the name assigned to the custom page.

Author
Specifies the creator of the page. Pages loaded by default have the author "ExtraHop."

Description
Provides a space for an optional, user-defined description.

Type

Specifies whether the page is for a device, application, or network.

Status

Specifies the number of assignments for the page.

Create a page

Important: Assigning a custom page with trends to more than 1,000 devices may impact system performance.

1. On the Pages page, click **New**.
   The Page Configuration window opens.
2. On the **Page Settings** tab, complete the following fields:

   Name

   Specifies the name assigned to the page.

   Author

   Specifies the creator of the custom page. The author is set by the Discover appliance based on the user name for manually created objects or the imported bundle, or set manually by the user. On bundle export, you can specify an author to override authors of any local objects included in the bundle. Pages loaded by default have the author "ExtraHop."

   Page Type

   Specifies whether the custom page is for a device, application, or network.

   Description

   Provides a space for an optional, user-defined comment.
3. On the **Assignments** tab, refer to the following:

   Assign to All

   Specifies that the custom page should be assigned to all objects, current as well as discovered in the future.

   Assignments

   Shows where in the system this custom page has been manually assigned to an object. (This field does not show when the page was assigned by clicking the **Assign To All** checkbox.) For more information, see the Assign a Page section.

   Remove All Assignments

   Removes all manually-added objects from the custom page.

Assign a page

Applications, devices, networks, and custom groups can be assigned to a custom page.

1. Navigate to a page for an application, device, network, or group.
2. Click on the name of the application, device, network, or group at the top of the left pane.
3. Click the **Pages** tab.
4. Select a custom page. If there are no pages available, you need to create a custom page. For more information, see the Create a Page section.
5. Click **OK**.

Configure a page

1. On a custom page, click **Edit Page**.
   The custom page toolbar appears.
2. Click **Configure Page**.

The **Page settings** tab displays the following:

Name
Specifies the user-defined name of the page.

Author
Specifies the user-defined author of the page.

Page type
Specifies the system-defined page type.

Description
Specifies the user-defined description of the page.

The **Assignments** tab displays the following fields:

Assign to All
Specifies that the page should be assigned to all applications, devices, or groups.

Assignments
Displays where in the system the page has been manually assigned to an applications, devices, or groups. To manually disassociate the alert from a device or group, click the delete symbol next to the device or group name. (This field does not show when the alert was assigned by clicking the **Assign To All** checkbox.)

Remove All Assignments
Removes all manually-added applications, devices, or groups from the alert.

3. Click **OK** to save the changes and exit the window.

Add a chart to a page

1. On a custom page, click **Edit Page** from the toolbar.
2. Click **Add Chart**.
3. In the Add New Chart window, click the **Title** field and enter a name for the chart.
4. Click the **Metric Source** drop-down list and select **Built-in** or **Trigger**.
5. Click the **Metric Type** drop-down list and select the type of metric to plot.
   The following types are available:

Name	Definition	Sample Metrics
Count	Integer datatype	• Number of packets in the network capture • Number of requests to the HTTP server • Number of errors on the database server
Dataset	A frequency table: For each entry, frequency is the number of times a value has been seen.	• HTTP server request transfer time • HTTP server processing time • HTTP server response transfer time
Sampleset	Mean and standard deviation	• HTTP server processing time detail (timing shown

Name	Definition	Sample Metrics
		when clicking on HTTP server "Responses") • Database server processing time detail (timing shown when clicking on DB server "Methods") • CIFS server access time detail (timing shown when clicking on CIFS "Files") • TCP round trip times (timing shown when clicking on TCP "Accepted" or "Connected")
Snapshot	Integer datatype, representing a snapshot in time	TCP established connections
Maximum	The final value produced from all of the recorded values	Top slowest database statements

6. Click the **Chart Type** drop-down list and select the format in which to view the metrics.
   The icon next to the drop-down list displays a preview of the selected format.

   Note: The heatmap displays only one dataset metric.

7. For line, box, bar, and heatmap charts, enter a label to use on the Y-axis in the **Units Label** field.
8. For line, box, heatmap, and bar charts, click the **Percentiles** drop-down list and select which percentiles to use when plotting uncertainty.
9. For line charts, select the icon with the type of line chart you want to view.
   - Broken line segments from a list of percentiles
   - Dashed bar series from a pair of percentiles
10. For line, box, and bar charts, click the **Use logarithmic scale** checkbox to view the values in non-linear intervals.
11. The Metrics table shows metrics that have been added to this chart.

    Display Name
    Name of the metric as it will appear in the chart.

    Metric Name
    Name of the top-level metric to plot in the chart.

    Detail Metric Name
    Name of the metric to display when clicking on the chart.

12. (Optional) To add a new line to the table, click **New** to add a new line to the table.
    a) Click the **Display Name**, **Metric Name**, or **Detail Metric Name** to edit the values.
       To plot a detail metric as a top-level metric on a custom page, use the following syntax:
       - For built-in detail metrics, use metric name, then "?", then substring to match the key (for instance, `http_server_detail:req?10.10.6.79` or `uri_http_server_detail:req?uri_substring`).
       - For custom detail metrics, use prefix "`custom_detail:`", then metric name, then "?", then substring to match the key (for instance, `custom_detail:custom_detail_metric_name?substring`).

Note: The substring must match one key exactly. If multiple keys are matched, results will be incorrect. You cannot plot sets of sets of detail metrics (for example, HTTP status codes and DB methods) as top-level metrics.

13. To search for built-in and custom metrics available on the current application, click **Find**.

    If **Metric Source** is set to **Built-in**, then the **Available Metrics** list shows all built-in metrics available on the current application. If **Metric Source** is set to **Custom**, then the **Available Metrics** list shows all custom metrics available on the current application. Start typing the metric name in **Metric Name** or **Detail Metric Name** fields and the metrics matching that substring will appear in the **Available Metrics** list.

Remove a chart from a page

1. On the page-level toolbar, click **Edit Page**.
   The configuration gear and delete icons appear in the upper right corner of each chart.
2. Click the **Delete** icon on the chart you want to remove.

Add a trend chart to a page

1. On a custom page, click **Edit Page**.
2. Click **Add Trend Chart**.
3. In the Add Trend Chart window, in the **Title** field, type a name for the chart.
4. In the **Units Label** field, type a label for the units in your chart.
5. To show a logarithmic measurement in your chart, click the **Use logarithmic scale** checkbox.
   If you do not click the checkbox, the chart will show a linear measurement by default.
6. Click the **Add Trend Line** button.
   For new charts, click Yes at the prompt to save the chart. For more information about adding a trend line, see the Trend Configuration section.
7. To manage trend lines, select the checkbox next to a trend name, click the **Select Action** drop-down list, and then select an action.

   Copy
   Copies the trend line.

   Delete
   Deletes the trend line from the list.

   Enable
   Makes the trend line appear on the chart.

   Disable
   Removes the trend line from the chart.

   Reset Data
   Removes previously gathered data from the chart.
8. Click **OK**.

After the Discover appliance has collected enough data to plot a trend line, it appears on the chart:

To isolate lines on the chart, uncheck the checkboxes next to the trends that are not of interest.

Trend configuration

The Trend Configuration window appears when you click **Add Trend Line** while editing a trend chart. From this window, you can do the following:

- Configure the Trend Lines
- Determine the Weighting Model
- Add an optional Multiplier
- Configure Exclusion Intervals

- Enter an optional Description
- View the trend History

Trend Line

The **Trend Line** tab provides configuration settings to define the trend line. The **Trend Line** tab contains the following items:

Name

Specifies a name for the trend.

Author

Specifies the creator of the trend. The author is set by the Discover appliance based on the user name for manually created objects or the imported bundle, or set manually by the user. On bundle export, you can specify an author to override authors of any local objects included in the bundle. Alerts loaded by default have the author "ExtraHop."

Disable Trend

Specifies whether the trend is disabled.

Metric

Specifies the metric associated with the alert. To select a metric, click the gear icon to the right of this field. Trends with dataset and sampleset metrics have the following additional options:

Merge

Merges all the datasets and applies the trending function to one big dataset.

Mean

Takes the mean of each dataset.

Percentile

Allows you to set a percentile value of datasets.

Standard Deviation

Calculates the normal deviation compared to the current trend using the same standard deviation parameters as the trend. These parameters can be absolute or relative, and population or sample. Normalization displays the standard deviation relative to mean. Click the **Normalization** drop-down list and select one of the following options.

Absolute

Displays the standard deviation as a constant.

Relative to Mean

Displays the standard deviation relative to the mean.

Note: If the trend is not a standard deviation, it is calculated as an absolute sample.

Ratio

Click the **Ratio** checkbox if the data will be measured as a ratio.

Window

Specifies the calculation window for the trend.

Same Hour of Week

Calculates the trend within a specified 1-hour window each week.

Same Hour of Day

Calculates the trend within a specified 1-hour window each day.

Minute Rolling Average

Calculates the trend based on the average of the data gathered each minute within a specified amount of time from the present time.

Hour Rolling Average

Calculates the trend based on the average of the data gathered each hour within a specified amount of time from the present time.

Lookback

Specifies the number of minutes of lookback.

Weighting Model

The **Weighting Model** tab provides options for how to weigh the trend line.

Mean

Specifies the manner in which to calculate the average.

Linear Average

Calculates the average with all data points weighted equally.

Single Exponential

Calculates the average with the most recent data points weighted more heavily.

Double Exponential

Calculates the average with the most recent data points weighted the most heavily.

For linear averages, the most recent value is weighted at 1 times the oldest value by default. For single and double exponential means, enter a number to weight the most recent value.

Percentile

Specifies the percentile value used as a basis for creating the trend.

Percentile

Records the trend using data points from a user-specified percentile.

Min Value

Records the lowest data point gathered during the time interval.

Max Value

Records the highest data point gathered during the time interval.

Regression

Linear

Calculates steadily increasing trends based on previous trends that are equally incremental.

2nd Degree Polynomial

Calculates exponentially accelerating trends by projecting a curve using the equation

```
y = ax2 + bx + c
```

Standard Deviation

Calculates the normal deviation compared to the current trend.

Type

Uses a sample-based or population-based standard deviation.

Normalization

Displays the standard deviation relative to mean.

Note: If a trend is a standard deviation, its associated alerts use the same parameters as the trend. If the trend is not a standard deviation, then the alert is calculated as "sample" and "absolute".

Static Value

Calculates a static value based on the number you enter, and is useful to plot constant lines for SLAs.

Time Delta

Uses the oldest trend, resulting in a time delta option based on the lookback window.

Trimean

Calculates the weighted average of the 25th, 50th, and 75th percentile values.

Winsorized Mean

Replaces the most outlying values with the highest and lowest remaining values. Values above the 90th percentile become the same value as the 90th and values below the 10th percentile become the same value as 10th.

Multiplier

The **Multiplier** tab allows you to view a multiple of the trend, which can be useful for banding. The **Multiplier** tab contains the following fields:

Multiplier

Enter a number by which to multiply the trend line.

Offset

Enter a number to display an offset line along the x-axis. Positive numbers offset forward and negative numbers offset backward. Click the drop-down list and select one of the following:

(absolute)

Displays the offset as a constant.

per minute

Calculates the offset per minute.

per hour

Calculates the offset per hour.

Exclusion Intervals

The **Exclusion Intervals** tab shows all the defined exclusion intervals that can be applied to trends. From this page, you can configure exclusion intervals. The Exclusion Intervals table contains the following information:

Name

Specifies the name of the exclusion interval.

Description

Provides a space for an optional, user-defined description.

Type

Specifies the type of exclusion interval. Options include:

One-time

Specifies an exclusion period that occurs only once from a designated start time (date and time) to a designated end time (date and time).

Daily

Specifies an exclusion period that occurs every day from a designated starting hour to a designated ending hour.

Weekly

Specifies an exclusion period that occurs every week from a designated start time (day and time) to a designated end time (day and time).

Description

The **Description** tab provides a space for an optional, user-defined description of the trend.

History

The **History** tab displays changes that have been made to the trend. The table contains the following columns:

Change
Displays the change that was made to the trend.

Author
Displays the author of the change.

Timestamp
Displays when the change was made.

Copy a page

1. On the Pages page, select the checkbox next to the page(s) that you want to copy and use as a template for defining a new page.
2. Click **Copy**.
   The name of the copied page is generated automatically by appending the word "(copy)" to the original name.

Delete a page

1. On the Pages page, select the checkbox next to the page(s) that you want to delete.
2. Click **Delete**.

Enable a page

1. On the Pages page, select the checkbox next to the page(s) that you want to enable.
2. Click **Enable**.

Custom pages that were previously disabled will become active and appear in the left panel.

Disable a page

1. On the Pages page, select the checkbox next to the page(s) that you want to disable.
2. Click **Disable**.

The selected custom pages will become inactive and not appear in the left panel.

View custom pages

1. Click **System Settings** in the top toolbar.
2. Click **Pages**.

Setup, administration, and maintenance

You can view information about ExtraHop system health through the Discover appliance and Command applianceWeb UI. To perform setup, administration, and maintenance tasks, log in to the ExtraHop Admin UI.

Log into the ExtraHop Admin UI

To access the Administration UI of a Discover or Command appliance:

1. Click the System Settings icon in the top toolbar.
2. In the Settings pop-up window, click **Administration**.
3. On the Administration UI Login screen, in the **Username** text box, enter your user name.
4. In the **Password** text box, enter your password.
5. Click **Log In**.

   If you do not know your Administration UI credentials, contact your organization's ExtraHop administrator. For detailed information about the settings that you can configure through the Administration UI, see the ExtraHop Admin UI Guide.

   Note: The default password for Amazon Web Services (AWS) users is the string of numbers after the `-i` in the instance ID.

View system health

The System Health page contains information used by ExtraHop Support to assess the health of the Discover appliance.

To view the System Health page:

1. Click the Systems Settings icon in the top toolbar.
2. In the Settings pop-up window, click **System Health**.
3. (Optional) Generate a system health report for ExtraHop Support.
   a) Click the **Time Interval** drop-down list, and select **Last week**.
   b) Click the PDF button.
   The PDF report will open in a new browser tab.
   c) Save the PDF file, and email it to ExtraHop Support.

System health

The System Settings window contains a table with the following information:

Capture

Incoming Packets Breakdown

The incoming packet rate (packets per second) for the selected time interval. `Current` and `Max` are the current and maximum packet rates. `Total` is the total number of packets over the selected time interval. Packets are broken down by the following criteria:

Analyzed

The number and rate of packets analyzed by the Discover appliance.

Filtered

The number and rate of filtered packets not included in network L2 metrics.

L2 Duplicates

The number and rate of identical Ethernet frames that are counted as duplicated L2 packets.

L3 Duplicates

The number and rate of identical IPv4 packets (TCP and UDP only) that are counted as duplicated L3 packets.

Note: A large packet rate may sometimes result in frames dropped at the span source or a span aggregator.

Note: Deduplication is enabled by default in newly installed (not upgraded) Discover appliances version 4.0 and later. Deduplication is not enabled in older and upgraded versions. Contact ExtraHop Support for more information about changing these settings.

Incoming Throughput Breakdown

The incoming throughput (bits per second) over the selected time interval. `Current` and `Max` are the current and maximum throughputs. `Total` is the total number of bytes transferred over the selected time interval. Throughput is broken down by the following criteria:

Analyzed

The number and rate of throughput analyzed by the Discover appliance.

Filtered

The number and rate of throughput not included in network L2 metrics.

L2 Duplicates

The number and rate of identical Ethernet frames that are counted as duplicated L2 throughput.

L3 Duplicates

The number and rate of identical IPv4 packets (TCP and UDP only) that are counted as duplicated L3 throughput.

RPCAP Packets

Displays the following metrics for the selected time interval:

Encapsulation

The total number of RPCAP encapsulation packets received by the Discover appliance.

Tunnel Eligible

The total number of packets eligible to be forwarded to the Discover appliance.

Tunnel Sent

The total number of RPCAP-tunneled packets forwarded to the Discover appliance.

Tunnel Received

The total number of RPCAP-tunneled packets received by the Discover appliance. Click the graph to view a list of peers.

RPCAP Throughput

Displays the following metrics for the selected time interval:

Encapsulation

The total number of RPCAP encapsulation bytes received by the Discover appliance.

Tunnel Received

The total number of RPCAP-tunneled bytes received by the Discover appliance. Click the graph to view a list of peers.

TCP Desyncs

The rate at which desyncs occurred system-wide during the selected time interval. `Current` and `Max` are the current and maximum rates of desyncs. `Total` is the total number of desyncs that occurred over the selected time interval.

Note: A desync is recorded if synchronization is lost when processing a TCP connection. Large numbers of desyncs might indicate dropped packets on the monitoring interface, SPAN, or network tap.

Trigger Executes

The number of times triggers were executed during the selected time interval.

Trigger Load

The trigger cycles as a percentage of total capture thread time. Mouse over a point on the graph to display the following metrics for the selected time interval:

Load

The trigger cycle load at the selected point in time.

Cycles

The number of cycles used out of the total available cycles.

Executes

The number of executes and the average number of executes per second.

Average per execute

The average number of cycles per execute.

Capture Heap Allocation

The amount of memory dedicated to the network capture.

Trigger Heap Allocation

The amount of memory dedicated to triggers.

External Timestamp

The percentage of packets with an external timestamp read by the Discover appliance based on the total number of packets processed.

Packet Capture Disk Throughput

Displays the following metrics for the selected time interval:

Total

The total number of bytes captured.

Current

The number of bytes captured during the most recent second.

Max

The largest number of bytes captured in a single second during the selected time interval.

Trigger Load by Trigger

The number of cycles used by each trigger.

Trigger Exceptions by Trigger

The number of unhandled exceptions thrown by each trigger. This graph will be empty if no triggers threw an unhandled exception during the selected time interval.

Remote

These graphs capture information about all of the activity on an ExtraHop Discover appliance that is generated by a “Remote.” trigger call or records that are sent to an external server.

Messages are one of the following types of system activity configured through Open Data Streams in the ExtraHop Admin UI:

- HTTP requests, from Remote.HTTP in Triggers
- Kafka messages, from Remote.Kafka in Triggers
- Syslog messages, from Remote.Syslog in Triggers
- MongoDB inserts, updates, and deletes, from Remote.MongoDB in Triggers

- Records sent through commitRecord() in Triggers or through built-in flow records

Messages Sent

The number of messages that were sent to a target system from the ExtraHop Discover appliance, per second. **Current** and **Max** are the current and maximum per-second rates. **Total** is the total number of messages sent for the selected time interval.

Message Throughput

The network throughput or bandwidth of message data that is sent from the ExtraHop Discover appliance in bytes per second.

Message Errors

The errors that are detected during the send process. Mouse over a point on the graph to display the following metrics for the selected time interval:

Send Errors

Indicates that an error occurred during the packet transmission.

Parse Errors

Indicates that there was an error when encoding the message from the trigger code, and the message could not be sent.

Bad Targets

The target name specified in the trigger code was not configured in the Admin UI. For example, Remote.Syslog("target_name").send(...).

Queue Full

Indicates that the remote server cannot keep up with the current message rate or the ExtraHop system cannot send messages fast enough.

Connections

The number of connection attempts sent to the target system.

Connection Attempts

The number of connection attempts.

Connection Errors

The number of errors generated during the connection attempt.

Remote Heap Allocation

The amount of memory dedicated to triggers that are generating message data.

Messages Dropped

The number of messages that were dropped because the internal message queue was full.

Message Queue Length

The number of messages waiting to be sent in the internal message queue.

Datastore

Disk Read Throughput

The disk read throughput rate (reads per second) over the selected time interval. `Current` and `Max` are the current and maximum read rates. `Total` is the total number of reads for the selected time interval.

Disk Write Throughput

The disk write throughput rate (writes per second) over the selected time interval. `Current` and `Max` are the current and maximum write rates. `Total` is the total number of writes for the selected time interval.

Store Read Throughput

The store read throughput rate (reads per second) over the selected time interval. `Current` and `Max` are the current and maximum read rates. `Total` is the total number of reads for the selected time interval.

Store Write Throughput

The store write throughput rate (writes per second) over the selected time interval. `Current` and `Max` are the current and maximum write rates. `Total` is the total number of writes for the selected time interval.

Working Set Size

The datastore write cache working set size for metrics on 1-hour, 5-minute, and 30-second cycles, which equals roughly the total metric types times the number of devices. ExtraHop recommends viewing these metrics over a 24-hour time period.

Metric Size

The metric size distribution.

Active Devices

The active L2, gateway, pseudo, and L3 devices over the selected time interval.

Total Devices

The total number of L2, gateway, pseudo, and L3 devices tracked in the datastore over the selected time interval.

Store Lookback

The estimated datastore lookback for fast (30 seconds), medium (5 minutes), and slow (1 hour) metrics based on write throughput.

Datastore Heap Allocation

The amount of memory dedicated to the datastore.

Datastore Trigger Heap Allocation

The amount of memory dedicated to triggers in the datastore process.

Datastore Trigger Load

The trigger cycles as a percentage of total datastore thread time. Mouse over a point on the graph to display the following metrics for the selected time interval:

Load

The trigger cycle load at the selected point in time.

Cycles

The number of cycles used out of the total available cycles.

Executes

The number of executes and the average number of executes per second.

Average

The average number of cycles per execute.

Datastore Trigger Drops

The number of triggers not executed in the datastore process because the triggers are using too many cycles.

Datastore Trigger Load by Trigger

The number of cycles used by each trigger in the datastore process.

Datastore Trigger Exceptions by Trigger

The number of unhandled exceptions thrown by each trigger in the datastore process. This graph will be empty if no triggers threw an unhandled exception during the selected time interval.

Trends

Performance Overview

The trend process utilization based on the last hour, and the date and time of the last trend recorded.

Trend Details

The total processing time in milliseconds for each trend on the Discover appliance based on the last hour.

View certificates

To view which uploaded certificates are being used for decryption:

1. Click the Systems Settings icon in the top toolbar.
2. In the Settings pop-up window, click **System Health**.
3. Click the **Certificates** button.

The SSL Certificates window contains a table with the following information:

Subject

The certificate name and unique SHA-1 hash function.

Decrypted

The number of sessions that were successfully decrypted.

Unsupported

The number of sessions that could not be decrypted using passive analysis (for example, using DHE key exchange).

Detached

The number of sessions that were not decrypted or only partially decrypted due to desyncs.

Passthrough

The number of sessions that were not decrypted due to hardware errors (for example, exceeding the specifications of SSL acceleration hardware).

Contact us

We value your feedback.

Please let us know how we can improve this document. Send your comments or suggestions to documentation@extrahop.com.

If you need additional help, please contact ExtraHop Support. at or visit the ExtraHop Customer Support Portal at https://www.extrahop.com/support/portal/.

- **Email:** support@extrahop.com
- **Support Portal Website:** https://www.extrahop.com/support/portal/
- **Telephone:**
  - 877-333-9872 (US)
  - +44 (0)203 7016850 (EMEA)
  - +65-31585513 (APAC)

Appendix

Global navigation overview

The ExtraHop Web UI provides a framework of elements that remain static as you move around the system. The information and options in the left and content panes of the Web UI change based on your selections in the top menu.

The following figure identifies both global navigation elements and the areas of the Web UI that will change based on your selection.

Top Menu

The following elements are located across the top of the Web UI.

Dashboards

Provides built-in system dashboards that give you an instant view of the activity on your network. You can also create and share dashboards with other users. For more information, see Navigating Dashboards.

Metrics

Provides access to system metrics sources, group metrics, and record queries. For more information, see Navigating Metrics.

Alerts

Provides access to the Alerts pages. For more information, see Navigating Alerts.

Global Search field

Enables you to type any object or search criteria and find a match on your Discover appliance. If you have an ExtraHop Explore appliance configured, you can also search for saved records.

Community Icon

Launches a new tab in your web browser with information about ExtraHop forums and other external resources.

Help icon

Launches documentation for the page that you are currently viewing.

System Settings

Provides access to system configuration options.

User Icon

Enables you to log in and log out of your Discover appliance or Command appliance, change your password, and access API options.

Navigation bar

The following elements are located across the top of the Web UI, below the top menu.

Pane toggle

Enables you to collapse or expand the left pane.

Global Time Selector

Enables you to determine the global time interval that is applied to all system metrics.

Recent Pages

Enables you to see the most recent pages you visited. Repeated pages are deduplicated and condensed to save space.

Navigation Path

Displays where you are in the system and provides available pivot points so you can search for the same metrics across multiple protocols, devices, or other swappable criteria.

Command menu drop-down

Appears throughout the Web UI and contains context-sensitive commands for the area you are in. For example, when you click the **Dashboards** top menu, the command menu at the end of the navigation bar provides options to view dashboard properties and to create a new dashboard.

Left Pane

The left pane changes based on your selection in the top menu and navigation bar.

- Click to see the Dashboard options
- Click to see the Metrics options
- Click to see the Alerts options

Content Pane

The content pane changes based on a combination of your top menu and left pane selections.0

- Click to see the Dashboard options
- Click to see the Metrics options
- Click to see the Alerts options

Navigating dashboards

Click **Dashboards** to view built-in system dashboards that give you an instant view of the activity on your network. You can also create and share dashboards with other users.

Left Pane

When you select **Dashboards** from the top menu, the left pane displays a dashboard Dock. The dashboard Dock is composed of folders, such as the **Dashboard Inbox**, **System Dashboards**, and **My Dashboards**. These folders contain system dashboards or any dashboards that you create or share. You can create additional folders as needed.

The following fields and controls are available in the **Dashboard** left pane.

Type to filter field

Enables you to limit the displayed list of items.

Dashboard sort buttons

Enables you to switch between ascending, descending, and custom sort views.

Dashboard Inbox

Contains dashboards that other users have sent you.

My Dashoards

Contains dashboards that you create.

System Dashboards

Contains the two default system dashboards, which are Network and Activity.

New Dashboards

Enables you to create a new dashboard.

Command menu button

Enables you to edit the dashboard dock and create a new, empty folder.

Content Pane

When you select **Dashboards** from the top menu, the content pane displays the selected dashboard.

Command Menu

When you select **Dashboards** from the top menu, a command menu appears on the far right of the navigation bar.

The following fields and controls are available in the Dashboard command menu.

Edit Layout

Customize your dashboards.

Dashboard Properties

Edit your dashboard name and access rights.

Share

Share your dashboard with another user.

Print

Send the dashboard you are viewing to a printer.

Modify Sources

Modify the metric sources used in the dashboard.

Copy

Save a duplicate of your dashboard.

Delete

Remove a dashboard from the system.

New Dashboard

Create a new dashboard.

Show Descriptions

Hover tooltips where available.

Presentation Mode

Displays a full-screen view of the metrics on the currently selected dashboard.

Widget Slideshow

Displays a slideshow of widgets within the current window.

Metric Explorer

Enables you to configure widgets to add to a dashboard.

For more information, see Dashboards.

Navigating metrics

Click the Metrics to view all metric sources, group metrics, and saved record queries.

Left Pane

When you select Metrics from the top menu, the left pane displays all of the types of available metrics sources in the system.

The following fields and controls are available in the Metrics left pane.

Type to filter field
Enables you to limit the displayed list of items.

Sources
Enables you to select metrics for applications, devices, and networks.

Groups
Enables you to select metrics for activity groups or to create a custom group.

Records
Enables you to query records and save queries for future use.

Content Pane

When you select **Metrics** from the top menu, the content pane displays the last metric source that you viewed. As you continue to select options from the left pane, the content pane displays lists, charts, and metrics for your selection.

For more information, see Metrics.

Create an alert

Before creating an alert, it is important to determine the metric that you want to monitor, the alert threshold for that metric, and the recipient or recipients of notifications for the alert. An alert is triggered only when a threshold is passed. For example, if a threshold is too low, then one alert will be sent when the threshold is passed and no more alerts will be sent.

Note: All external notification alerts are sent in UTC regardless of the time zone set in the Discover appliance Admin UI.

To create a new alert:

1. Click the System Settings icon in the top toolbar.
2. In the Settings window, click **Alerts**.
3. On the Alerts page, click the **New** button.
4. Enter the following information in the **Alert Settings** tab.

Alert Settings
Provides configuration settings to define the alert name and the alert expression.

Trend Settings
Provides configuration settings to select the time, lookback, and weight of trend-based alerts.

Description
Provides a space for an optional, user-defined description.

Exclusion Intervals
Displays the exclusion intervals assigned to this alert.

Notifications

Provides configuration settings to identify the email groups that should be notified when this alert fires.

Assignments

Displays where in the system this alert has been manually assigned to a device or group.

5. When you are finished entering the alert settings, click **OK** to save the alert and exit the Alert Configuration dialog box.

Add a page to a report

1. Navigate to a metric page.
2. Click **Add to Report**.
3. Select a report.
   - To add the page to an existing report, click the **Add to Report** drop-down list, select an existing report, and click **OK**.
   - To add the page to a new report, click **New Report**. In the Report Configuration window, enter a report name and click **OK**.

Export data to Excel

1. Navigate to a metric page.
2. Right-click any table, chart, or tile on the page and select **Export to Excel**.

Export data to Excel

1. Navigate to a metric page.
2. Right-click any table, chart, or tile on the page and select **Export to Excel**.

Export data to CSV

1. Navigate to a metric page.
2. Right-click any table, chart, or tile on the page and select **Export to CSV**.

Create a PDF of a metric page

1. Navigate to a metric page.
2. Click **PDF**.

Open metrics in the Metric Explorer

Metrics from a metric page can be added to a widget in the Metric Explorer.

1. Navigate to a metric page.
2. Click **Open in Metric Explorer**.

Pin a metric page to a dashboard

1. Navigate to a metric page.
2. Click **Pin to Dashboards**.
   The confirmation dialog box displays the name of the dashboard that the page has been added to.
3. Click **Dashboards**.
4. In the left pane, under My Dashboards, click the name of the dashboard.

Sort metrics

If a metric contains a gear icon in the upper-right corner, the metric can be sorted by key or value.

1. Navigate to a metric page.
2. Click the gear icon on the upper-right corner of a metric.

3. Select either **Sort by Key** or **Sort by Value**.

Navigating alerts

The **Alerts** top menu enables you to view system alerts information.

Left Pane

When you select Alerts from the top menu, the left pane displays available alert information in the system

The following fields and controls are available in the Metrics left pane.

Alert History

Enables you to view detected system alerts.

Trouble Groups

Enables you to view built-in metrics groups that have been identified as having problems.

Content Pane

When you select **Alerts** from the top menu, the content pane displays the latest alerts that you viewed.

For more information, see Alerts.

Time selector

The Time Selector enables you to specify a time interval for the collection and presentation of network data. There are two types of Time Selectors: a Global Time Selector for specifying global time intervals, and a Region Time Selector for specifying region time intervals.

The Global Time Selector is located at the top-left of the navigation bar. The Region Time Selector is located to the top-right of the dashboard region header.

A global time interval is applied across the Discover appliance. Navigating from one area to another will not change the time interval for the metrics you are viewing. This means that the same time interval applies whether you are viewing different metrics across the Web UI or if you are drilling-down to view detailed metrics.

Note: Global time interval information is included at the end of the URL. When copying a URL, make sure that the entire URL is copied to maintain the specified global time interval.

A region time interval is applied by dashboard region and you can set different time intervals per-region. When you add a widget to an existing region, the widget inherits the time interval for that region.

You can apply either a global time interval or a region time interval to a dashboard region. To toggle between time intervals, start by clicking the command menu in the region header. To apply a region time interval, select **Use Region Time Selector**. To apply a global time interval, select **Use Global Time Selector**. When the Region Time Selector disappears from the region header, this indicates that the global time interval is applied to the region.

To specify a global or region time interval:

1. Click the Global Time Selector or the Region Time Selector.
2. From the Time Interval tab, select one of the following options:

   Last 30 minutes

   Displays the last 30 minutes of data collected.

   Last 6 hours

   Displays the last six hours of data collected.

   Last day

   Displays the last 24 hours of data collected.

Last week

Displays the last seven days of data collected.

Last

Displays the data collected within a custom time window. For more information, see the Specify a Time Window section.

Custom time range

Displays the data collected within a fixed time range. For more information, see the Specify a Custom Time Range section.

3. Click **Save**.

You can view metrics with different levels of granularity based on the time interval that you specify. For example, if you specify a time interval of 120 minutes or less you will see metrics in aggregations of 30-seconds, if available. (If 30-second aggregation metrics are unavailable, five-minute or 60-minute aggregation metrics will be displayed depending on availability.) If you specify a time interval between 121 minutes and 24 hours, you will see metrics in aggregations of five-minutes, if available. A time interval that is greater than 24 hours will display 60-minute metrics. If you have an extended datastore that is configured for 24-hour aggregation metrics, a specified time interval of 30 days or longer will display 24-hour metrics. One-second metrics are available for specific network and device-level data when the specified time interval is less than six minutes. For more information, see the L2 Networks Page and L2 Devices Page sections.

Time intervals are preserved across login sessions. The five most recent unique time intervals are also saved in the **Time Selector History** tab.

To select a previous time interval:

1. Click the Global Time Selector or Region Time Selector.
2. Click the **History** tab.
3. Select a time interval. You selection will be applied to the options on the **Time Interval** tab.
4. Click **Save**.

Displaying Running Time and Snapshot Time Intervals

For dashboards and top-level metrics pages—where metrics are polled automatically—you will see the running time for the global time interval displayed in the Global Time Selector.

For a detailed metric page or a records query results page—where metrics are not polled automatically—you will see the snapshot of the global time interval, which includes a blue refresh icon and gray text that indicates when the metric or record query was last polled. To reload the metrics or query for the specified time interval, click the refresh icon in the Global Time Selector display.

Specify a time window

To view metrics that occurred at a specific time, you can use the custom time window option in the Time Selector to specify the number of minutes, hours, days, or years from the present.

To specify a custom time window for a global or region time interval:

1. Click the Global or Region Time Selector and select the **Last** radio button in the **Time Interval** tab.
2. Type the number of units of time.
3. Click the drop-down list and select **minutes**, **hours**, **days**, **weeks**, **months**, or **years**.
4. Click **Save**.

Specify a custom time range

To view metrics that occurred during a specific time, you can specify a custom time range or you can zoom in on a chart.

To specify a custom time range:

1. Click the Global Time Selector or Region Time Selector.
2. From the **Time Interval** tab, and select **Custom Time Range**.
   The drop-down field will display a default time range.
3. Click the drop-down field. A calendar dialog box opens.
4. Click a day to specify the start date for the range. One click will specify a single day. Clicking another day will specify the end date for the range.

   Note: Use the back and forward arrows on the calendar to change the month displayed on the calendar.
5. Click **Save**.

Compare metric deltas

From the Dashboards page, you can compare a single metric across two time intervals.

Note: Delta comparison is only available for dashboards. If you save a comparison and navigate to another area of the Discover appliance, the comparison will be disabled temporarily. When you return to the Dashboards area, the delta comparison you saved will be enabled again.

To create a delta comparison for a dashboard region:

1. Locate the dashboard region containing the metrics you want to compare.
2. Click the Time Selector.
   - If you applied a global time interval to the dashboard region, click the Global Time Selector in the navigation bar.
   - If you applied a region time interval to the dashboard region, click the Region Time Selector in the upper right corner of the region.

   For more information, see the Time Selector topic.
3. Click **Compare** in the Time Interval tab.
4. In the Delta Comparison pane, select the time interval to use in the comparison or enter your own custom ending time.

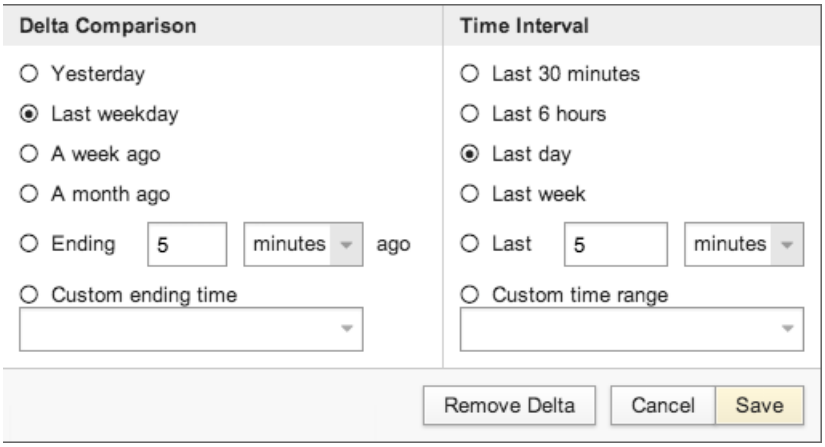

5. Click Save.

   On the dashboard, a new chart is overlaid onto the previous chart displaying the metrics for the new time interval.

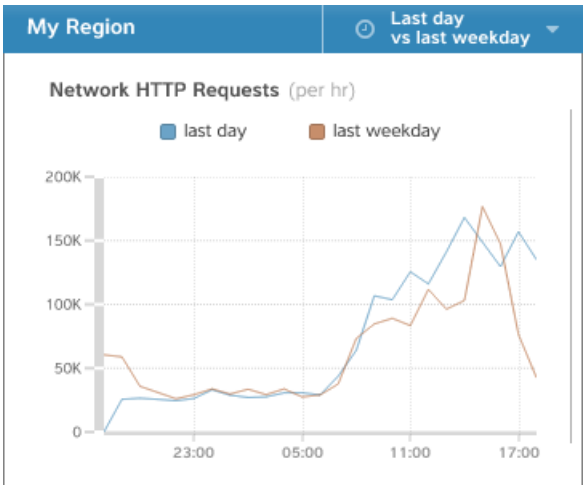

Zoom in on a time range

You can click-and-drag across a region in a line chart to zoom in and specify a custom time range in the Time Selector. For example, if you observe a spike in a chart, you can click-and-drag across the spike to zoom in on the activity that occurred in that time range.

Note: This option is only available for time-series charts. It is not available for bar charts, text widgets, or tables.

If you are zooming in on a chart within a dashboard region that has a region time interval applied to it, this time range will become the region time interval for every widget in that region (unless you have applied a global time interval to that dashboard region). The ability to zoom in on a time range is useful for observing other metric activity that occurred in that same time range. For more information, see the Time Selector section.

If the specified time range is valid it appears green. If the specified time range is less than one minute, the range is invalid and appears red.

Note: Data might not be available for the zoomed time range.

1. Click and drag your mouse across the chart to select a time range.
2. Release the mouse button. The graph is redrawn to the specified time range.

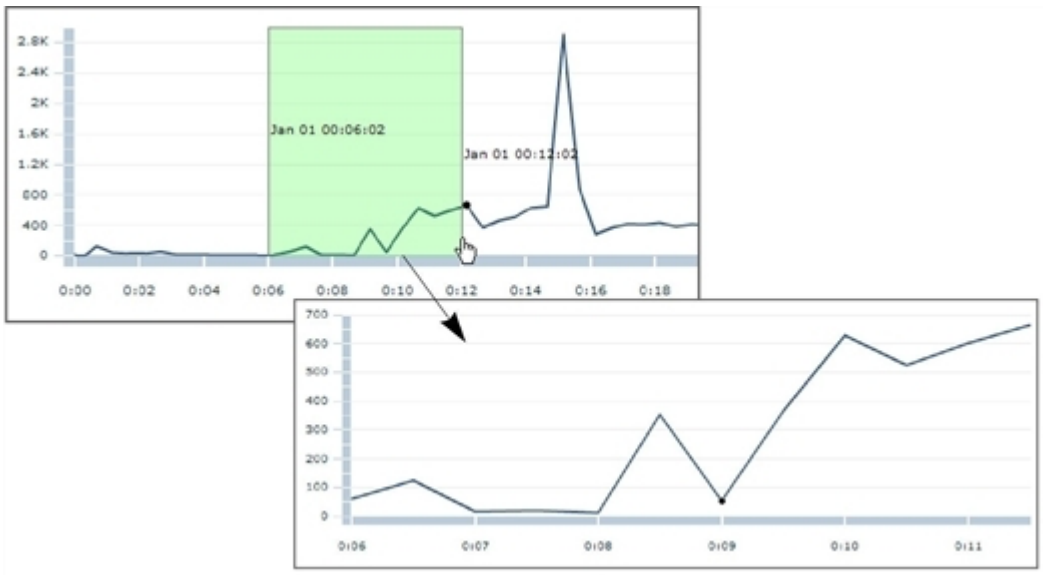

The scales on the chart's axes update to reflect the range of values in the zoomed time range. In addition, the **Custom Time Range** value in the Time Selector adjusts to reflect the time range in the chart.

If you want to revert from the zoomed time range back to your original time interval, click the undo icon—a magnifying glass with a minus sign—in the Time Selector. For example, if you originally specified "`Last 30 minutes`" as your time interval, and then perform a series of zoom operations on a chart, you can revert back to your original 30-minute time interval with one click on the undo icon.

ExtraHop modules

The Discover appliance provides metrics through the following types of modules:

Module Type	Protocols
L2-L3 Metrics	• Multicast • IP • IPv6 • ICMP • ICMPv6
L4 Metrics	• TCP • UDP
Naming	DNS
Directory Services	LDAP
Web	• HTTP/HTTPS • AMF • SSL
Middleware	• MS-RPC • Memcache • IBMMQ
Database	• IBM DB2 • IBM Informix • Microsoft SQL Server • MongoDB • MySQL • Oracle • PostgreSQL • Sybase ASE • Sybase IQ
Storage	• iSCSI • CIFS • NFS
File Transfer	FTP
Mail	SMTP
Citrix VDI	• ICA

Module Type	Protocols
	• CGP
Industry-Specific Protocols	• Diameter • FIX • HL7 • RADIUS • SMPP • Telnet
Decryption	Any protocol encrypted over end-to-end SSL channel, can be decrypted using the SSL decryption module.

For more information about ExtraHop modules, visit extrahop.com.

Browser compatibility

The following browsers are compatible with all ExtraHop appliances.

- Chrome 45
- Firefox 41
- Internet Explorer 10 and 11
- Safari 9

Common acronyms

The following common computing and networking protocol acronyms are used in this guide.

Acronym	Full Name
AAA	Authentication, authorization, and accounting
AMF	Action Message Format
CIFS	Common Internet File System
CLI	Command Line Interface
CPU	Central Processing Unit
DB	Database
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
ERSPAN	Encapsulated Remote Switched Port Analyzer
FIX	Financial Information Exchange
FTP	File Transfer Protocol
HTTP	Hyper Text Transfer Protocol
IBMMQ	IBM Message Oriented Middleware
ICA	Independent Computing Architecture
IP	Internet Protocol

Acronym	Full Name
IRL	Index record log
iSCSI	Internet Small Computer System Interface
L2	Layer 2
L3	Layer 3
L7	Layer 7
LDAP	Lightweight Directory Access Protocol
MAC	Media Access Control
MIB	Management Information Base
NFS	Network File System
NVRAM	Non-Volatile Random Access Memory
RADIUS	Remote Authentication Dial-In User Service
RPC	Remote Procedure Call
RPCAP	Remote Packet Capture
RSS	Resident Set Size
SMPP	Short Message Peer-to-Peer Protocol
SMTP	Simple Message Transport Protocol
SNMP	Simple Network Management Protocol
SPAN	Switched Port Analyzer
SSD	Solid-State Drive
SSH	Secure Shell
SSL	Secure Socket Layer
TACACS+	Terminal Access Controller Access-Control System Plus
TCP	Transmission Control Protocol
UI	User Interface
VLAN	Virtual Local Area Network
VM	Virtual Machine

Keyboard shortcuts

Keyboard shortcuts enable you to quickly navigate across the Discover appliance or perform specific actions with a few keystrokes.

The following keyboard shortcuts apply across the entire Discover appliance.

Key	Action
?	Show or hide a hot key help menu
G then S	Go to Dashboard

Key	Action
G then A	Go to Alerts
G then P	Go to Application Metrics
G then N	Go to Network Metrics
G then D	Go to Device Metrics
G then G	Go to Group Metrics
/	Global Search
O then M	Open Metric Explorer
G then E	Go to Settings
G then T	Go to Trigger Editor
G then H	Open Help
O then Q	View system information
Ctrl+S	Save widget configuration

The following keyboard shortcuts only apply to dashboards.

Key	Action
O then L	Toggle edit layout mode
O then P	Show dashboard properties
C then D	Copy the current dashboard
D then D	Delete the current dashboard
O then S	Toggle descriptions
Ctrl+Shift+F	Toggle presentation mode
N then D	Create a new dashboard
N then F	Create a new folder
O then D	Toggle dock edit mode

Built-in pages

The following pages are built-in to ExtraHop Command and Discover Appliances.

Applications page

ExtraHop provides a set of default applications based on all traffic. You can modify the default application template to suit the needs of your organization, and you can add your own applications.

Applications do not always adhere to device boundaries. Some applications use multiple devices, and some devices host multiple applications. You can use Application Inspection Triggers to define application boundaries based on criteria other than a list of devices (for example, URIs or database table names). Defining an application allows you to report on an application based on the subset of network traffic that comprises it, regardless of the devices associated with it. For information about using triggers to define applications, refer to Triggers.

The Applications page includes a table that lists all devices discovered on your networks. The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information. The counter at the bottom of the table identifies the number of applications currently displayed in the table. The table can show up to 100 applications per page.

The Applications page contains the following information:

Name

Specifies the name of the application.

Capture

Specifies the capture point for which the application was defined.

Description

Provides a space for an optional, user-defined description.

Custom application page

If a custom page has been assigned to an application, the name of the custom page appears in the left pane.

For information about creating custom pages, see Custom Pages.

Application overview page

The Application Overview sub-page includes interactive charts that provide an overview of a selected application.

Each chart shows an overview of activity for all active protocols. You can also view details for only certain protocols as well as a summary of a specific time or date.

To show overall details for only certain protocols, select those protocols in the chart legend.

To show a summary of activity for a specific time or date, mouse over the time period of interest.

- For statistical charts, a pop-up dialog showing a five-number summary appears, including the minimum, lower quartile, median, upper quartile, and maximum values.
- For area charts, a pop-up dialog showing total count and time appears.

Note: Because area charts are stacked, the total count represented by the number on the left side of the chart is a sum of the count for each individual protocol.

To show a particular region of the chart, click and drag across that region.

To show only a specific protocol in the chart, mouse over the protocol in the chart legend.

To view details for a specific protocol, click it in the chart. The protocol's application page appears.

For more information about working with the charts, refer to Drill-Down Functionality. For information about a specific protocol, refer to that protocol's application topic.

Transactions

Shows the total number of transactions (requests and their responses) for the active protocols excluding SSL and ICA, which are not transactional protocols.

Errors

Shows the total number of errors for the active protocols excluding SSL and ICA.

Processing Time

Shows the total server processing time for the active protocols.

L2 Bytes

Shows the total count of request bytes and response bytes transferred for the active protocols.

Packets

Shows the total count of request packets and response packets transferred for the active protocols.

Application geomaps page

The Geomaps sub-page lists the geomaps associated with the application. Geomaps display worldwide activity based on the metrics defined in that geomap.

The Geomaps sub-page displays the following information:

Geomap

Displays the name of the geomap.

Metric

Displays the metric displayed in the geomap.

Description

Displays a description of the geomap.

Application alert history page

The Alert History sub-page provides an alert summary for application-level alerts. The Discover appliance can be configured to generate both threshold and trend-based alerts for any metric in the system. Alerts can be configured to send email notifications or SNMP traps as proactive early warnings for potential performance problems.

The application Alert History page displays all alerts, including alerts that have been acknowledged previously, and the corresponding time for each alert for the current application. The Alert History page also includes additional information about trend alerts that have fired.

To use the Alert History page, you must first create alerts. For more information, refer to Configuring Alerts.

The Alert History page includes the following information:

Alert History

Displays alerts that have been generated.

Time

Displays the time that the alert was generated.

Alert

Displays the name of the alert.

For threshold-based alerts, clicking the name of the alert displays the following information:

Name

The name of the alert.

Expression

The metric, time interval, operator, and sensitivity that were defined when the alert was created.

Value

The value of the metric at the time the alert fired. This is used for comparison against the alert expression.

Description

The optional user-defined description of the alert.

For trend alerts, clicking the name of the alert displays the following information:

Name

The name of the alert.

Alert Conditions

The type of alert, time interval, operator, and/or percentage of the trend that were defined when the alert was created.

View at Time of Alert

The alert graph from when the alert was fired.

View Current State

The alert graph of the current trend state of the alert.

Current Trend State

Displays a list of trend alerts assigned to the application.

Trend

Displays the name of the trend alert.

Stat

Displays the metric that the trend alert is based on.

Description

Displays a description of the trend alert.

You can sort the table by the following parameters:

All Alerts

Displays alerts created on the Command appliance and the node.

Command appliance Alerts

Displays alerts created on the Command appliance only.

Local Alerts

Displays alerts created on the node only.

Click the name of the trend alert to view the following information about the alert:

Alert Graphs

Displays the trend alert over time and whether or not it has fired.

Alert Condition Nominal

Indicates the metrics being gathered have not reached an alert state.

Alert Firing

Indicates the metrics being gathered have met the alert criteria.

Alert Rules

Displays the rules of the trend alert and whether or not it has fired.

Alert Condition Nominal

Displays the alert rules in green.

Alert Firing

Displays the alert rules in red.

Devices page

This section provides information about viewing device metrics to troubleshoot network issues at the device level.

The Devices page includes a table that lists all devices discovered on your networks. The counter at the bottom of the table identifies the number of devices currently displayed in the table. The table can show up to 1,000 devices per page.

The Devices table contains the following columns:

Name

The primary name the device on the network. For more information, see the Device Search section. Hover over the device name shows a description of the device type, such as:

- WWW server

- DB (database) server
- File server
- Load balancer
- Gateway
- Custom device

MAC Address

The MAC address is a unique identifier of the device network interface.

VLAN

The Virtual Local Area Network (VLAN) of the device. VLAN information is extracted from VLAN tags, if the traffic mirroring process preserves them on the mirror port.

IP Address

The last IP address the device used to communicate on the network. For more information, see the Device Search section.

Discovery Time

The time when the device was first discovered. The day of the week, the calendar date, and time is displayed in the following format:Wed Feb 23 09:01.

Description

Provides a space for an optional, user-defined description.

The Devices page also includes a search feature that locates devices with plain text or regular expressions. For more information, see the Device Search section.

Custom device page

If a custom page has been assigned to an device, the name of the custom page appears in the left panel.

For information about creating custom pages, see Custom Pages.

Device overview page

The Device Overview sub-page includes interactive charts that provide an overview of a selected device.

Each chart shows an overview of activity for all active protocols. You can also view details for only certain protocols as well as a summary of a specific time or date.

To show overall details for only certain protocols, select those protocols in the chart legend.

To show a summary of activity for a specific time or date, mouse over the time period of interest.

- For statistical charts, a pop-up dialog showing a five-number summary appears, including the minimum, lower quartile, median, upper quartile, and maximum values.
- For area charts, a pop-up dialog showing total count and time appears.

Note: Because area charts are stacked, the total count represented by the number on the left side of the chart is a sum of the count for each individual protocol.

To show a particular region of the chart, click and drag across that region.

To show only a specific protocol in the chart, mouse over the protocol in the chart legend.

To view details for a specific protocol, click it in the chart. The protocol's application page appears.

For more information about working with the charts, refer to Drill-down Functionality. For information about a specific protocol, refer to that protocol's application topic.

Transactions

Shows the total number of transactions (requests and their responses) for the active protocols excluding SSL and ICA, which are not transactional protocols.

Errors

Shows the total number of errors for the active protocols excluding SSL and ICA.

Processing Time

Shows the total server processing time for the active protocols.

Device geomaps page

The Geomaps sub-page lists the geomaps associated with the device. Geomaps display worldwide activity based on the metrics defined in that geomap.

The Geomaps sub-page displays the following information:

Geomaps

Displays the name of the geomap.

Metric

Displays the metric displayed in the geomap.

Description

Displays a description of the geomap.

For more information about geomap settings, refer to Geomaps.

Device alert history page

The Alert History sub-page provides an alert summary for network-level alerts. The Discover appliance can be configured to generate both threshold and trend-based alerts for any metric in the system. Alerts can be configured to send email notifications or SNMP traps as proactive early warnings for potential performance problems.

The device Alert History page displays all alerts, including alerts that have been acknowledged previously, and the corresponding time for each alert for the current device. The Alert History page also includes additional information about trend alerts that have fired.

To use the Alert History page, you must first create alerts. For more information, refer to Configuring Alerts.

The Alert History page includes the following information:

Alert History

Displays alerts that have been generated.

Time

Displays the time that the alert was generated.

Alert

Displays the name of the alert.

For threshold-based alerts, clicking the name of the alert displays the following information:

Name

The name of the alert.

Expression

The metric, time interval, operator, and sensitivity that were defined when the alert was created.

Value

The value of the metric at the time the alert fired. This is used for comparison against the alert expression.

Description

The optional user-defined description of the alert.

For trend alerts, clicking the name of the alert displays the following information:

Name

The name of the alert.

Alert Conditions

The type of alert, time interval, operator, and/or percentage of the trend that were defined when the alert was created.

View at Time of Alert

The alert graph from when the alert was fired.

View Current State

The alert graph of the current trend state of the alert.

Current Trend State

Displays a list of trend alerts assigned to the application.

Trend

Displays the name of the trend alert.

Stat

Displays the metric that the trend alert is based on.

Description

Displays a description of the trend alert.

You can sort the table by the following parameters:

All Alerts

Displays alerts created on the Command appliance and the node.

Command appliance Alerts

Displays alerts created on the Command appliance only.

Local Alerts

Displays alerts created on the node only.

Click the name of the trend alert to view the following information about the alert:

Alert Graphs

Displays the trend alert over time and whether or not it has fired.

Alert Condition Nominal

Indicates the metrics being gathered have not reached an alert state.

Alert Firing

Indicates the metrics being gathered have met the alert criteria.

Alert Rules

Displays the rules of the trend alert and whether or not it has fired.

Alert Condition Nominal

Displays the alert rules in green.

Alert Firing

Displays the alert rules in red.

Activity groups page

Discover appliance automatically generates activity groups based on network traffic. A member might appear in more than one activity group if it has multiple types of traffic.

Click **All** in the navigation bar drop-down list to display all activity groups. Select **Client** or **Server** to filter activity groups by devices acting as a client or server, respectively.

The table includes the following group information:

Name
Specifies the name of the activity group.

Count
Identifies the number of devices that belong to this activity group.

To view details about the members in an activity group:

1. In the Name column, click the activity group name to view the group metrics.
2. On the group metrics page, click any of the metrics in blue to view device-level statistics.
3. In the table at the bottom of the page, click a name to view metrics about the group member.

   When a name is clicked from this page, the Discover appliance Web UI redirects to the Devices functional area and opens the device statistics page for the protocol specified by the activity group. For example, when the TCP activity group is open and you click a device name, the Web UI opens the L4 TCP protocol metrics page for that device.

Custom groups page

The Custom Groups page lists all user-defined device groups in the Discover appliance. There are two types of custom groups:

Static
Add devices to the group manually and modify the list of devices associated with the group. For more information, see the Static Custom Groups section.

Dynamic
Specify a rule that automatically adds and removes devices from the group. You can modify the criteria that defines the group, but you cannot manually add or remove devices from the group. For more information, see the Dynamic Custom Groups section.

The Filter text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

The Custom Groups table includes the following device group information:

Name
Specifies the name of the device group. The icon next to the name indicates whether the device group is a static or dynamic group.

Count
Identifies the number of devices that belong to the device group.

Description
Provides a space for an optional, user-defined description.

Group devices page

The Devices sub-page lists the devices in the group. You can filter the list of devices and manage the assignments for a device or group of devices. You can click a device to open a detailed metrics page for that device. To return to the list of devices, click the back button in your browser.

For information about searching for a device, refer to Device Search.

Group geomaps page

The Geomaps sub-page lists the geomaps associated with the group. Geomaps display worldwide activity based on the metrics defined in that geomap.

The Geomaps sub-page displays the following information:

Geomap

Displays the name of the geomap.

Metric

Displays the metric displayed in the geomap.

Description

displays a description of the geomap.

For more information about geomap settings, refer to Geomaps.

Networks page

This section describes the network capture attributes, network alerts, and network traffic details. The Network page is the entry point into the network capture. The metrics that are collected and displayed here provide a summary of all network activity retrieved in the capture.

Note: When using the Network page as the starting point for data analysis, remember that the information collected on network devices is determined by the port mirror configuration. The device is only aware of the traffic passed to it.

In addition, if your organization uses the Command appliance to manage multiple network capture points, the Networks page displays a table of all capture points for your entire networking environment. You can click a specific network listed in the table to open the detailed Network page with metrics for that network. Otherwise, clicking the **Networks** button leads directly to the capture point on the local system.

The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

The network capture provides the following information about the capture itself as well as the Discover appliance that initiated the capture:

Name

The name of the network capture. The name attribute includes an icon that opens a text box to edit the name of the network capture. This text area can be used to provide a more user friendly name for the capture.

Devices

The number of devices in the network capture.

MAC Address

The MAC address of the Discover appliance responsible for the network capture.

IP Address

The IP address of the Discover appliance responsible for the network capture.

Description

An optional detailed description of the network. This attribute includes an icon that opens a text box for a user-entered description of the network capture. This text area can be used to provide additional information about this particular network capture.

Alerts

A list of alerts assigned to the network. This list includes controls to add or remove network-level alerts from the network capture.

Pages

A list of all custom pages assigned to the network. This list includes controls to add or remove network-level custom pages from the network capture.

The Network page is the starting point to review the capture-level metrics collected by the Discover appliance.

To view capture-level metrics:

- On the Networks page, click a capture node in the list to view the capture details. The network capture details page appears.

Custom network page

If a custom page has been assigned to a network, the name of the custom page appears in the left pane.

For information about creating custom pages, see Custom Pages.

Network devices page

The Devices sub-page within the **Networks** functional area lists the devices discovered on the network in the current network capture.

The table contains the following columns:

Name

The primary name the device uses to communicate on the network. Names are discovered by passively monitoring a variety of naming protocols, including DNS, DHCP, NETBIOS, and Cisco Discovery Protocol. If a device name is not discovered, a NIC manufacturer-based identifier is assigned to the device by looking at the MAC address. If the MAC address range is not registered, or if it belongs to a private MAC address space, the name includes the last six characters of the MAC address (for example, Device 00000c0789b1).

The device-type icon to the left of the device name identifies the activity primarily associated with this device.

The device name and type can be edited by clicking on the name and using the edit tool on the Device page.

MAC Address

The MAC address is a unique identifier of the device network interface. For physical devices that have multiple interfaces, one entry per interface is maintained. The vendor icon displays to the left of MAC Address as determined by the MAC OID lookup.

VLAN

The ID of the VLAN the device is connected to.

IP Address

The Primary IP address the device uses to communicate on the network. By default, Address Resolution Protocol (ARP) traffic is used to determine the mapping from MAC addresses to IP addresses. In the absence of such traffic, IP packet header information is used. If there is no ARP traffic, the IP address 0.0.0.0 is assigned to routing devices, such as gateways, firewalls, and load balancers, to indicate that it handles packets from many sources.

Discovery Time

The time when the device was first discovered. The day of the week, the calendar date, and time are displayed in the following format:Wed Aug 06 09:01.

Description

A user-defined description of the device. To edit the device description, click the device name and use the edit tool on the Device page.

Network alert history page

The Alert History sub-page provides an alert summary for network-level alerts. The Discover appliance can be configured to generate both threshold and trend-based alerts for any metric in the system. Alerts can be configured to send email notifications or SNMP traps as proactive early warnings for potential performance problems.

The network capture Alert History page displays all alerts, including alerts that have been acknowledged previously, and the corresponding time for each alert for the current network capture. The Alert History page also includes additional information about trend alerts that have fired.

To use the Alert History page, you must first create alerts. For more information, refer to Configuring Alerts.

To check the network capture alert history:

1. In the left pane in the Networks functional area, click the **Alert History**.
2. Find a specific alert in the table.

   To sort the table by time, click the **Time** column heading, and then click the arrow in the right corner of the column to sort in ascending or descending order.

   To sort the table by alert entry, click the **Alerts** column heading, and then click the arrow in the right corner of the column to sort in ascending or descending order.
3. Click the alert to view more information. The Alert Details window includes the following:

   Name

   The name of the alert.

   Expression

   The metric, time interval, operator, and sensitivity that were defined when the alert was created.

   Value

   The value of the metric at the time the alert fired. This is used for comparison against the alert expression.

   Description

   The optional user-defined description of the alert.

   For trend alerts, the Trend Alert Details window includes the following:

   Name

   The name of the alert.

   Alert Conditions

   The type of alert, time interval, operator, and/or percentage of the trend that were defined when the alert was created.

   View at Time of Alert

   The alert graph from when the alert was fired.

   View Current State

   The alert graph of the current trend state of the alert.

To view trend alerts:

1. On the Alert History page, click the **Current Trend State** tab to view a list of trend-based alerts assigned to the network.
2. Find a specific trend in the table.

   (Command appliance Only) Click the **Show** drop-down list and select one of the following options:

   All Alerts

   Displays alerts created on the Command appliance and the node.

   Command appliance Alerts

   Displays alerts created on the Command appliance only.

   Local Alerts

   Displays alerts created on the node only.

   To sort the table by trend, click the Trend column heading, and then click the arrow in the right corner of the column to sort in ascending or descending order.

   To sort the table by metric, click the Stat column heading, and then click the arrow in the right corner of the column to sort in ascending or descending order.
3. Click the trend name to view more information about the trend alert.

4. Click the **Alert Graphs** tab to view the trend alert over time and whether or not it has fired.

   Alert Condition Nominal

   Indicates the metrics being gathered have not reached an alert state.

   Alert Firing

   Indicates the metrics being gathered have met the alert criteria.
5. Click the Alert Rules tab to view the rules of the trend alert and whether or not it has fired.

   Alert Condition Nominal

   Displays the alert rules in green.

   Alert Firing

   Displays the alert rules in red.
6. Click **Back to Trend Alerts** to return to the Current Trend State table.

Alert History page

The Alert History page contains a list of triggered alerts for a specified time interval. This page provides an overview of the most recent application, device, and network alerts that have fired during the capture period.

To use the Alert History page, you must first create alerts. Click the **Configure Alerts** button, and the Alerts page opens in the **System Settings** pop-up window. For more information about configuring alerts, refer to the Configuring Alerts section.

The **Filter** text box above the table uses ActionScript regular expressions. Refer to ActionScript documentation for more information.

The Alert History table provides the following information about the alerts:

Source Type

The type of object that triggered the alert, either an application or a device.

Source

The name of the application or device that is the source of the alert.

Node

The Discover node on which the alert exists.

Alert

The name of the alert.

Most Recent

The time of the most recently fired alert.

Settings page

The Settings pop-up window provides controls to define and manage the alerting, tagging, reporting, and custom scripting features within the Web UI. It also provides a link to the Discover appliance Administration UI, which is a separate web application used to configure the Discover appliance settings.

The Settings window includes the following pages:

Administration

Manage the Discover appliance in the Administration UI. For more information about the features and pages in the Administration UI, see the Discover appliance Administration UI Users Guide.

Alerts

Define alerts and apply them to devices and device groups.

Bundles

Upload a set of related objects into the Discover appliance.

Custom Devices
Configure custom devices.

Device Limits
Select which devices receive full analysis of metrics.

Device Tags
Manage all defined device tags.

Flex Grids
Create and manage flex grids.

Geomaps
Configure and manage information displayed on geomaps.

Metric Catalog
View and edit built-in and custom metrics.

Pages
Configure and manage user-defined pages within the Discover appliance Web UI.

Record Formats
Configure and manage record formats.

Reports
Define ExtraHop reports and add device metrics to specified reports.

System Health
View usage and other metrics about the Discover appliance.

Triggers
Define and manage the triggers that execute user-defined scripts.