


Detections

Published: 2018-12-05

The ExtraHop system applies machine learning techniques to your wire data to identify unusual behaviors and potential risks to the security or performance of your network. When notable behavior is identified, the ExtraHop system generates a detection that contains information about the behavior and the source on which it occurred.

 **Note:** This topic applies to all ExtraHop systems, including ExtraHop Reveal(x).

 **Note:** Detections require a [connection to the cloud-based ExtraHop Machine Learning Service](#).

Unlike other machine learning solutions that rely on logs or agent data or monitoring tools such as manually-configured alerts, detections do not require additional configuration or maintenance as your network infrastructure changes.

Detections offer the following types of help:

- Uncover hidden issues before they create problems for your users.
- Collect high-quality, actionable data to identify the root causes of network issues.
- Gain deeper insight into your network behavior.
- Find unknown performance issues, security issues, or infrastructure quirks.

After you [connect to the ExtraHop Machine Learning Service](#), the ExtraHop system begins to analyze your stored data to identify performance or security detections, and the Detections page is available from the top menu.

You can view detections from the Detections page of the ExtraHop Web UI. Each detection provides details and links to help you investigate the issue. You can learn which factors contributed to an issue, view protocol activity on the source application or device, and filter detections to help you prioritize which issues to investigate first.

Here are important considerations about detections:

- You must have at least four weeks of wire data metrics stored on the ExtraHop system before detections can be identified.
- Users with restricted read-only privileges cannot view the Detections page, and dashboards shared with these users do not display detection markers.
- Depending on your ExtraHop subscription, your detections highlight either potential performance issues or security risks. Security detections are available only in ExtraHop Reveal(x) and require a license for the Machine Learning Service.
- You can access detections on a Command appliance for any connected ExtraHop Discover appliances that are also licensed for the Machine Learning Service. Command appliances can only connect to Discover appliances that are on the same subscription, such as ExtraHop Reveal(x).
- Although detections provide you with high-quality, actionable data about performance issues and security risks, detections do not replace decision-making or expertise about your network. Always investigate detections to determine the root cause of unusual behavior and when to take action.

Security detection categories

The best way to stop attackers from stealing data or wreaking havoc on your network is to detect attacks before they cause harm. Even though attackers regularly develop new methods for evading detection, most attacks tend to follow familiar patterns or phases. ExtraHop Reveal(x) can detect suspicious network behavior associated with different phases of an attack chain, such as reconnaissance or lateral movement. Detections that are identified at one or more of these phases can help reduce dwell time and prevent disruptions from potential attacks.

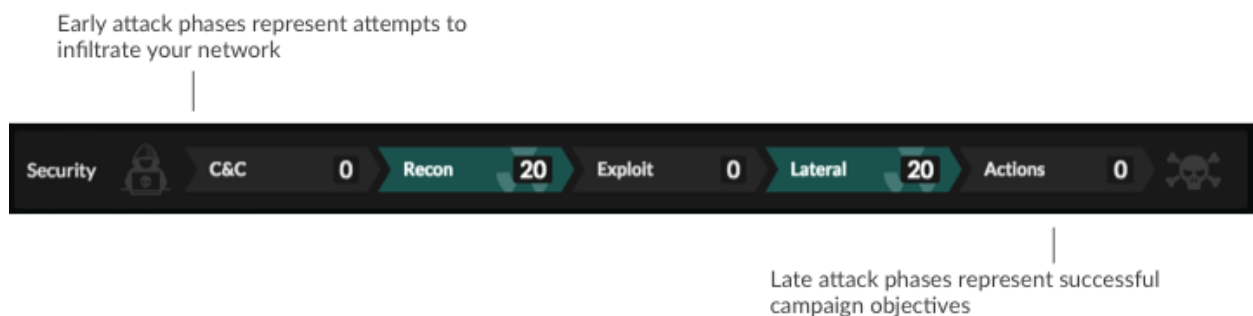
Note: This topic applies only to ExtraHop Reveal(x).

Note: Security detections provide you with high-quality, actionable data about security risks. But these detections do not replace decision-making or expertise about your network. Always investigate detections to determine the root cause of unusual behavior and when to take action.

Attack chain

Most network attacks tend to follow familiar patterns or phases. These phases can be assembled into an attack chain that characterizes the progression of steps an attacker takes to ultimately achieve their objective, such as stealing sensitive data.

Reveal(x) assigns an attack chain category to all security detections. On the Detections page, the attack chain flow chart highlights the number of detections that are associated with each attack phase, as shown in the following figure.



Important: Multiple detections in the attack chain can be associated with an attack. Detections associated with attack phases can be detected in any order.

The following types of security risks are associated with each phase of the attack chain.

Command and control

A compromised device on your network is attempting to contact an attacker’s external Command and Control (C&C) server. After the connection is established, the C&C server can send additional malware, instructions for remote execution, and payloads to support the attack. Reveal(x) detects when an internal device is communicating with a suspicious system outside of your network in support of an attack.

Reconnaissance

An attacker is looking for information about your network to find potential targets (such as critical assets) and weaknesses that can be exploited. Reveal(x) detects scans and various other techniques that map out devices and services on the network.

Note: Scans can be detected for known vulnerability scanners such as Nessus and Qualys. Click the device name to confirm if the device is already assigned a Vulnerability Scanner role in the ExtraHop system. To learn how to assign this role to a device, see [Change a device role](#).

Exploitation

An attacker is taking advantage of information about your network to actively exploit assets and vulnerabilities. For example, if an attacker logs into an important file server or database after finding valid credentials. Or if an attacker tries to cover their tracks by evading an intrusion detection system (IDS). Reveal(x) detects unusual behavior associated with various exploitation techniques such as brute force attacks and IP fragmentation.

Lateral movement

After the attacker infiltrates your network, they can start to progressively move from device to device in search of data, which might be the ultimate target of their attack campaign. Reveal(x) detects unusual device behavior associated with east-west corridor data transfers and connections.

Actions on objective

The ultimate objective of an attack can vary, from stealing sensitive data to encrypting files for a ransom. Another objective might include misappropriating network resources for botnet activity or cryptocurrency mining. Reveal(x) detects when an attacker is close to completing a campaign objective.

Performance (IT operation) detection categories

Detections automatically surface network, application, and infrastructure problems and identify their root causes, so that you can direct your investigation to any trouble areas.

Detections identify potential issues in the following performance and IT operation categories:

Authentication & Access Control

Unsuccessful attempts by users, clients, and servers to log in or access resources. For example, an authentication detection might reveal WiFi issues over the AAA protocol, excessive LDAP errors, or uncover resource-constrained devices.

Database

Database access problems for applications or users based on analysis of database protocols. For example, a database detection might show that the database server is sending an excessive number of response errors causing slow or failed transactions. A database detection might also reveal that an application cannot be reached due to Memcache issues.

Desktop & App Virtualization

Long Citrix load times or poor quality sessions for end users. For example, a virtualization detection might reveal an excessive number of Zero Windows, which indicates that the Citrix server is overwhelmed or experiencing issues.

Network Infrastructure

Unusual events over the TCP, DNS, and DHCP protocols. For example, a network detection might show DHCP issues that are preventing clients from obtaining a configured IP address from the server, or reveal that services were unable to resolve hostnames due to excessive DNS response errors.

Service Degradation

Service issues or performance problems identified during analysis of key Voice over IP (VoIP) and email communications protocols. For example, a service degradation detection might reveal that VoIP calls have failed and provide the related SIP status code, or show that unauthorized callers have attempted to make several call requests.

Storage

Problems with user access to specific files and shares detected when evaluating network file system traffic. For example, a storage detection might show that users were prevented from accessing files on Windows servers due to CIFS/SMB issues, or that NAS servers could not be reached due to NFS errors.

Web Application

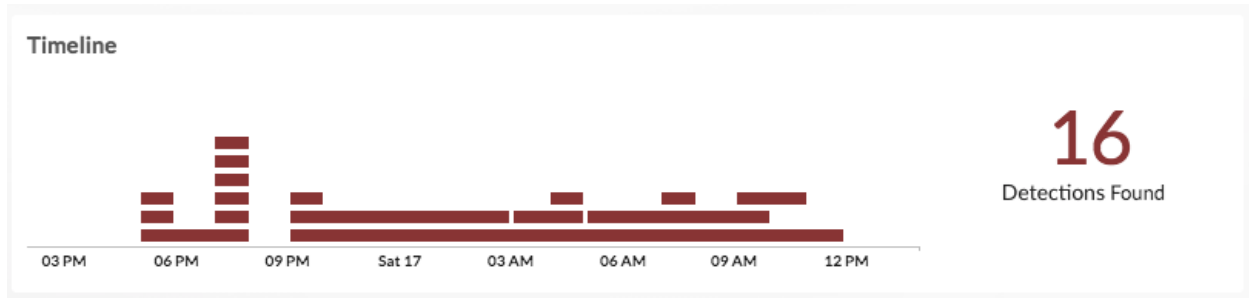
Poor web server performance or issues observed during traffic analysis over the HTTP protocol. For example, a web application detection might reveal that internal server issues are causing an excessive number of 500-level errors, preventing users from reaching the applications and services they need.

Interpret detections


The Detections page displays detections identified on your system and provides filtering options to help you find detections that are most important to you. The following sections show you what information you can learn from detections.

View total detections over time

The Timeline chart displays the total number of detections identified within the selected time interval, which might change as filters are applied. Each horizontal bar in the chart represents a single detection and the duration of the detection.



For Reveal(x) only, the color of each horizontal bar correlates to the risk score of the detection, as shown in the following figure:



Here are some ways to interact with the Timeline chart:

- Look for the tallest stack of bars to determine when the most detections occurred.
- Hover over a bar to view the detection title. For Reveal(x) only, the risk score is displayed.
- Click the bar to navigate directly to the detection detail page.
- Click and drag to highlight an area on the chart to zoom in on a specific time range. The time interval dynamically updates to match the new time range in the chart.

View details for individual detections

Each detection provides information about the type of issue that occurred, when it occurred, and the source of the issue. Individual detections are listed below the Timeline chart, sorted by the most recent start time.

The following figure shows the details that are provided for each individual detection:

The screenshot shows a detection card for 'Overwhelmed Data Transfer on glue-compute-xl-02.sea.i.extrahop.com'. Callout boxes point to various parts of the card:

- Start time and duration of the detection:** Sep 11 16:00 lasting an hour.
- Title of the detection:** Overwhelmed Data Transfer on glue-compute-xl-02.sea.i.extrahop.com (applebloom.sea.i.extrahop.com).
- Description and additional details:** This device sent an excessive number of Zero Windows to peer devices, which indicates that this device became resource-constrained during data transfer and could not accept any more data. Includes a link to the ExtraHop blog.
- Category of the detection:** NETWORK INFRASTRUCTURE.
- Details linked to this detection:**
 - IP: vc.sea.i.extrahop.com (10.10.11.4)
 - L7 Protocol: SSL:443
- Name of the anomalous metric and details about the behavior of the metric:** A table showing 'Zero Windows Out' with a 6-hour snapshot, 1-hour peak value of 22 K, expected range of 0-1, and a deviation of 2,198,700%.

Duration

The duration of the detection indicates how long the anomalous value was detected by Machine Learning Service. The minimum duration of a detection is 30 seconds. Detection data is analyzed every 30 seconds or every hour, depending on the metric. If the duration value is displayed as ONGOING, the anomalous metric has not returned to a normal value or an anomalous event has not finished.

Sparkline

Sparklines are simple line charts that show you the metric behavior around the time of the detection. The sparkline charts display a snapshot of metric data from the time frame around the duration of the detection (such as 6 hours), and not the overall time interval from the top of the page (such as the last 7 days).

Peak Value

The peak value is the maximum roll-up value of the metric observed over the duration of the detection. Metric values are rolled up, or aggregated, into either 1-hour, 5-minute, or 30-second periods.

Expected Range


The expected range includes values that represent a normal background level of activity, which is calculated based on 4 weeks of data. The expected range is the basis for comparison with observed values to detect abnormal changes in metric activity.

Deviation

A deviation is the percentage by which the metric value differs from the expected range.

Here are some ways you can interact with detection details:

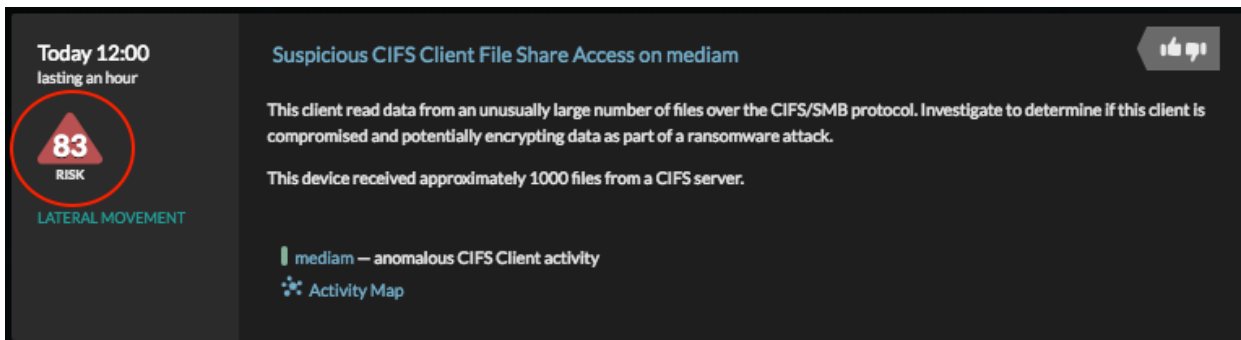
- Click the source device or application name to navigate to the protocol page associated with the metric and to access additional charts, metrics, and tools.
- Click the sparkline to open the Metric Explorer for the metric. Metric characteristics, such as the source, time interval, and drill-down details are preserved so that you can quickly create a chart from the metric or add additional sources and metrics for comparison.
- Click the activity map link to display all of the L7 protocol activity and device connections to the client or server in the detection. For more information, see [Activity maps](#).

Click the feedback icon  to let us know if the detection was helpful. Your feedback is valuable and helps us improve our identification process. All feedback is anonymous and will not have an immediate effect on your detections. You can submit feedback for a detection more than once.

See [Investigate detections](#) for tips and best practices.

Risk score (Reveal(x) only)

Each security detection has an associated risk score that can help you quickly identify urgent or critical detections in your environment. The risk score is displayed in the left pane of the detection details, similar to the following figure:



Each risk score is color coded by severity:

- Red = 80-99
- Orange = 31-79
- Yellow = 1-30

The risk score is calculated based on the following criteria:

Likelihood

An estimate of how likely it is that an attacker might discover and exploit the detection.

Skill level

The technical skill level required by an attacker to exploit the detection.

Impact

An estimate of the technical and business impact to company operations and value should an attacker exploit the detection.

Ticket tracking

Ticket tracking enables you to connect tickets, alarms, or cases in your work-tracking system to ExtraHop detections. Without leaving the ExtraHop Web UI, you can see who is working on a specific detection, as well as the status and outcome of that investigation.

When ticket tracking is enabled and configured, ticket details are displayed in the left pane of the detection details, similar to the following figure:

Today 14:00
lasting an hour

83
RISK

LATERAL MOVEMENT

Suspicious CIFS Client File Share Access on AccountingLaptop

This device sent an excessive number of read requests over the Common Internet File System (CIFS) protocol. This anomaly indicates that the device might be compromised and is preparing files for data exfiltration.

Server linked to this anomaly:

- corpshare.example.com (192.168.6.179)

AccountingLaptop Activity Map

CIFS Metric	6-hour Snapshot	Peak Value	Expected Range	Deviation
Reads		1.13 K	0-1	112,500%

Status — **CLOSED**

Ticket ID — ✓ EX-4437

Assignee — hopuser

Status

The status of the ticket associated with the detection. Ticket tracking supports the following statuses:

- New
- In Progress
- Closed
- Closed with Action Taken
- Closed with No Action Taken

Ticket ID

The ID of the ticket in your work-tracking system that is associated with the detection. If you have configured a template URL, you can click the ticket ID to navigate to the ticket in your work-tracking system.

Assignee

The username assigned to work on the ticket associated with the detection. Usernames in gray indicate a non-ExtraHop account.

See [Configure ticket tracking for detections](#) for more information.

Filter the detections list

The Detections page lists all detections in the specified time range by the detection start time; the most recent detection is listed first. You can further organize the detections list through filters, group by, and sort by options available in the left pane of the Detections page.

Filters

Filters enable you to view only those detections that match the criteria you select, as described in the following figure:

Change the time interval to view detections from a different time range.

Click to filter by a specific application or device.

From a Command appliance, click one or more Discover appliances.

Click to filter by one or more protocols.

Click to filter detections by application or device.

Additional filters are available, as shown in the following figure, if you have [ticket tracking](#) enabled:

Click to filter by New, In Progress, or Closed tickets. Select No Ticket Status to filter for tickets that have no status value configured yet.

Click to filter tickets by No Action Taken or Action Taken. Select No Resolution to filter for tickets that have no resolution value configured yet.

Click to filter tickets that are unassigned, assigned to you, or assigned to one or more other users. Usernames in gray indicate a non-ExtraHop account.

Type to filter by a ticket ID.

Filters applied to the detections list persist even when you apply grouping or sorting. For example, if you filter by ticket status and sort by highest risk, the detections list only displays detections that match all conditions.

Group by

The Group by option enables you to filter the detections list by source or title.

Grouping by source retrieves all of the source devices and applications on which detections occurred. Grouping by title retrieves the titles of all identified detections. The sources or titles display in a new pane to the left of the detections list. Click a source or title to display only those detections on that source or with that title in the detections list. In the example below, detections are grouped by title.

Title of detections identified in the selected time interval. | Start time of the most recent detection with this title. | Number of detections with this title. No number indicates a single detection.

Title of detections identified in the selected time interval.	Start time of the most recent detection with this title.	Number of detections with this title. No number indicates a single detection.
Overwhelmed Data Transfer	A DAY AGO	8
DNS Server Errors	5 DAYS AGO	
Delayed LDAP Server Authentication	5 DAYS AGO	
Delayed LDAP Client Authentication	5 DAYS AGO	
Suspicious CIFS Client Transactions	7 DAYS AGO	2
Suspicious CIFS Server Transactions	7 DAYS AGO	2

Overwhelmed Data Transfer
 Sep 18 10:00 lasting 2 hours
 NETWORK

Overwhelmed Data Transfer on VMWare 4C2693
 This device sent an excessive number of Zero Windows to peer devices, which indicates that this device became resource-constrained during data transfer and could not accept any more data. [Learn more about TCP analysis for applications on the ExtraHop blog.](#)
 Details linked to this detection:
 • L7 Protocol: SSL:443

TCP Metric	6-hour Snapshot	5-minute Peak Value	Expected Range	Deviation
Zero Windows Out		20 K	0-1	2,000,800%


Sort by (Reveal(x) only)

The sort by option enables you to sort the detections list by the most recent start time or by the highest risk score.



Category

The Category section enables you to filter the detections list by the number of detections associated with each category. If you apply filters that change the detections list or time interval, the number of detections displayed for a category might change.



Detection categories depend on your ExtraHop subscription. See [Reveal\(x\) detections](#) and [Performance \(IT Operations\) detections](#) for more information.

Find detections in the Web UI

Although you view detection details in the Detections page, indicators and links to detections are available throughout the Web UI.

Activity page

The Activity page provides links to detections found on each active protocol, as shown in the following figure:

Map	Protocol			Detections	Activity
✱	AAA	9 servers	26 clients	—	
✱	CIFS	74 servers	321 clients	10 detections	
✱	Database	5 servers	12 clients	7 detections	
✱	DHCP	19 servers	2,397 clients	—	
✱	DNS	29 servers	1,583 clients	9 detections	
✱	FTP	5 servers	12 clients	—	
✱	HTTP	293 servers	1,097 clients	3 detections	

Overview pages

The Overview page from any application, device, or device group provides a detections link, as shown in the figure below. Click the link on the right to view any associated detections on the Detections page.

Overview page for Dell B8B2E6. The page includes a sidebar with navigation options like Overview, Child Devices, Network, Client Activity, and DHCP. The main content area shows device details (Name, MAC, Groups, Tags, Role, Description) and performance metrics (Throughput, Throughput Summary, Total Traffic). A 'Detections' link is highlighted in the top right corner.

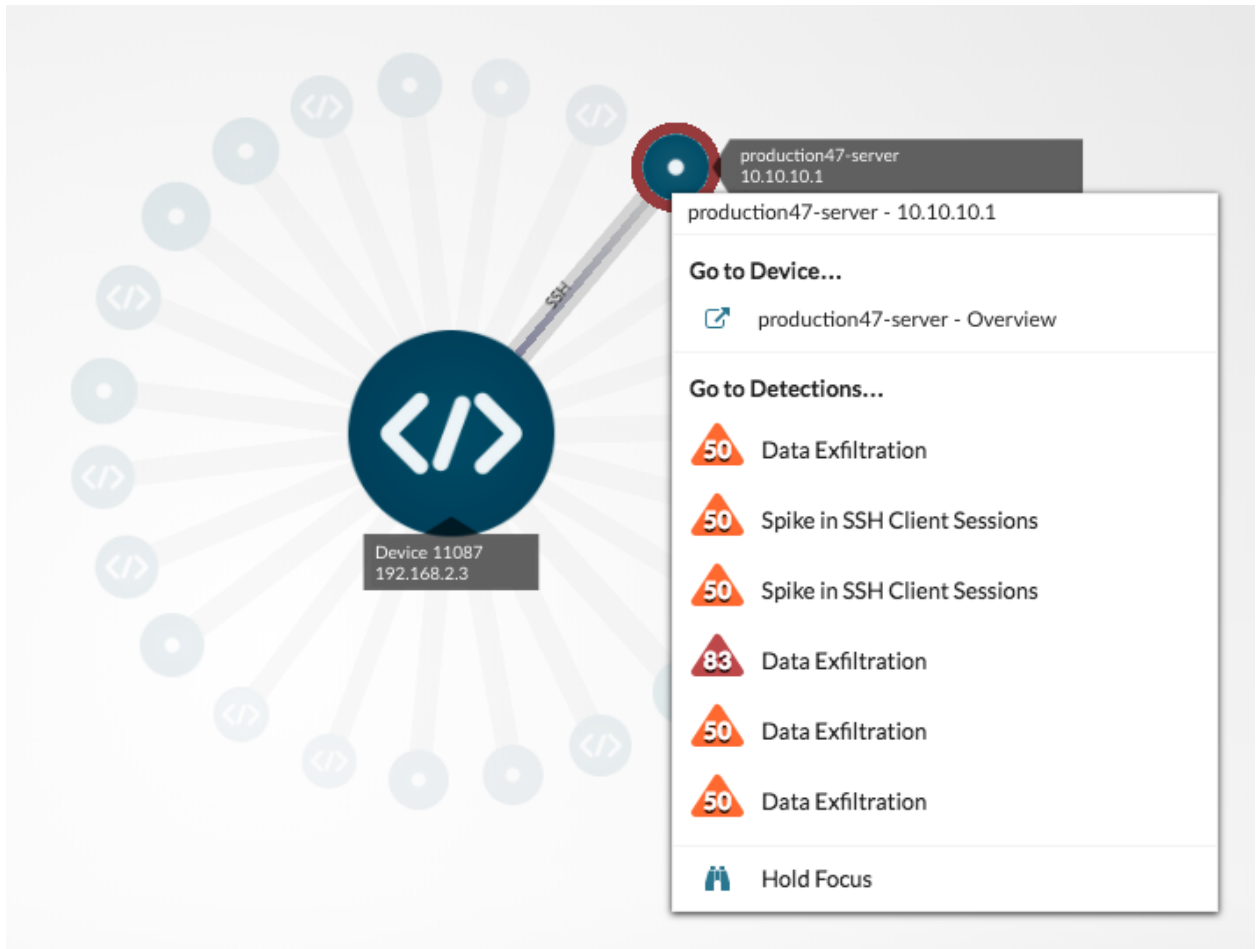
Protocol pages

Protocol pages for sources and groups provide a detections link, as shown in the figure below. Click the link in the top right corner to view any associated detections on the Detections page.


HTTP Summary for Group page. The page displays a line graph of Transactions (Responses and Errors) over time. Summary statistics show 279,388 Responses and 22 Errors. A 'Detections' link is circled in red in the top right corner.

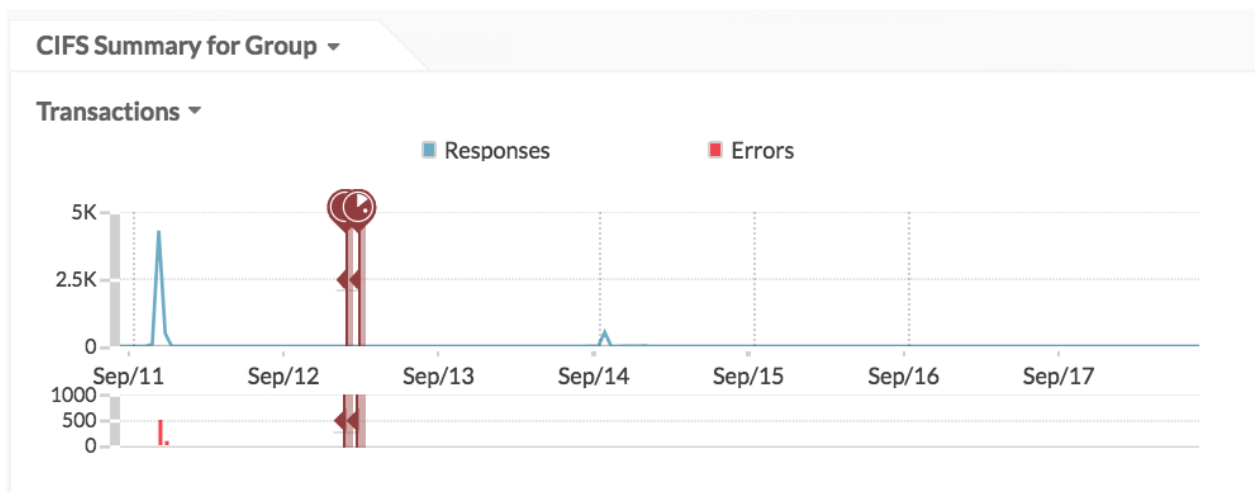
Activity maps

On activity maps, [devices with associated detections](#) display animated pulses around the circle label. Click a device to display detection links, as shown in the following figure:



Charts

Detection markers  appear on charts to indicate when a detection has occurred on a source device or application, as shown in the figure below. Hover over the marker to display the title of the detection or click the marker to view detection details



Alerts

If [detection alerts](#) are configured, the Alert History includes the title of the detection that generated the alert, as shown in the figure below. Click the title link to view detection details.

The screenshot shows the ExtraHop Alerts interface. At the top, there are navigation tabs: Dashboards, Alerts, Detections, Metrics, Records, and Packets. Below the navigation is a search bar and a filter section for 'Alerts'. The main area displays an 'Alert History' table with the following columns: Severity, Alert, Source, Time, and Alert Type. The table contains several rows of alerts, with the row 'Detection Alert Web Server Issues' circled in red.

Severity	Alert	Source	Time ↓	Alert Type
NOTICE	Storage Error Ratio - Yellow	All Activity	2018-09-20 11:26:30	Threshold
ERROR	Storage Error Ratio - Orange	eComApp	2018-09-20 11:22:30	Threshold
ERROR	Storage Error Ratio - Orange	All Activity	2018-09-20 11:22:30	Threshold
NOTICE	Storage Error Ratio - Yellow	eComApp	2018-09-20 11:22:30	Threshold
ALERT	Storage Error Ratio - Red	eComApp	2018-09-20 07:23:00	Threshold
ALERT	Storage Error Ratio - Red	All Activity	2018-09-20 07:23:00	Threshold
ERROR	Detection Alert Web Server Issues	VMware 4C261	2018-09-20 06:45:44.753	Detection
ERROR	Detection Alert Database Transaction Failures	mysql1	2018-09-20 06:45:44.753	Detection
ERROR	Detection Alert LDAP Server Auth Errors	ldap1	2018-09-20 06:45:44.753	Detection

How ExtraHop detections work

This section provides some background information on how the cloud-based ExtraHop Machine Learning Service identifies detections.

Essentially, a detection is identified when observed data deviates from the expected range of data by a significant amount.

The ExtraHop system observes metrics from wire data for the protocols, devices, and applications discovered on your network and then analyzes those metrics to identify detections. A subset of these metrics is delivered over an encrypted connection from the ExtraHop system to the Machine Learning Service in the cloud. The proprietary algorithm that drives the Machine Learning Service combines time series decomposition, unsupervised learning, heuristics, and unique domain expertise from ExtraHop. This combination helps to ensure that detections are both accurate and actionable. The ExtraHop system calculates the expected range of normal network behavior and then adapts to changing variations in protocols and metric data.


The ExtraHop system identifies detections based on three variables:

- Observed data, collected in real-time on your ExtraHop appliance.
- Expected range data, calculated from four weeks of historical data on your ExtraHop appliance.
- Threshold values, which are automatically adjusted by the algorithm based on historical metric data and heuristics defined by the IT networking and security experts at ExtraHop.

In most network monitoring tools, unusual activity is detected through manually-configured alerts and trend models for individual devices. However, as your network changes—because of hardware reconfigurations, organization mergers, business growth, or the addition of applications to your network—these types of alerts and models can become quickly outdated and potentially inaccurate. Detections automatically deliver consistent and accurate results about anomalous metrics and protocols without requiring manual configuration for individual devices.

Related topics

Check out the following resources that are designed to familiarize new users with Detections.

- [Investigate detections](#) 
- [Share a detection](#) 