Detections

Published: 2018-10-10

The ExtraHop system applies machine learning techniques to your wire data to identify unusual behaviors and potential risks to your network security or performance. Unlike other machine learning solutions that rely on logs or agent data, detections do not require additional configuration or maintenance as your network infrastructure changes.

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

After you have connected to the ExtraHop Machine Learning Service, the Detections page is enabled, and the ExtraHop system begins to analyze your stored data to identify performance and security detections.

Detections offer the following types of help:

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

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 can only view metrics included in the dashboards that you share with them. Those users will be unable to view detections. For more information, see Share a dashboard with a restricted user.
- If you are managing multiple ExtraHop Discover appliances through a Command appliance, you can access detections for any connected Discover appliance that is enabled for detections.

Depending on your ExtraHop system edition, your detections can highlight potential performance issues or security risks. Security detections are available only in ExtraHop Reveal(x).

Security detections

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 unusual network behaviour associated with different phases of an attack. Security detections help you learn about security risks, what type of attack is associated with the risk, and which devices are affected by the risk.

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.

When you log into the Web UI of your ExtraHop system, click **Detections**. The Detections page appears with all of the security detections identified during the selected time interval.

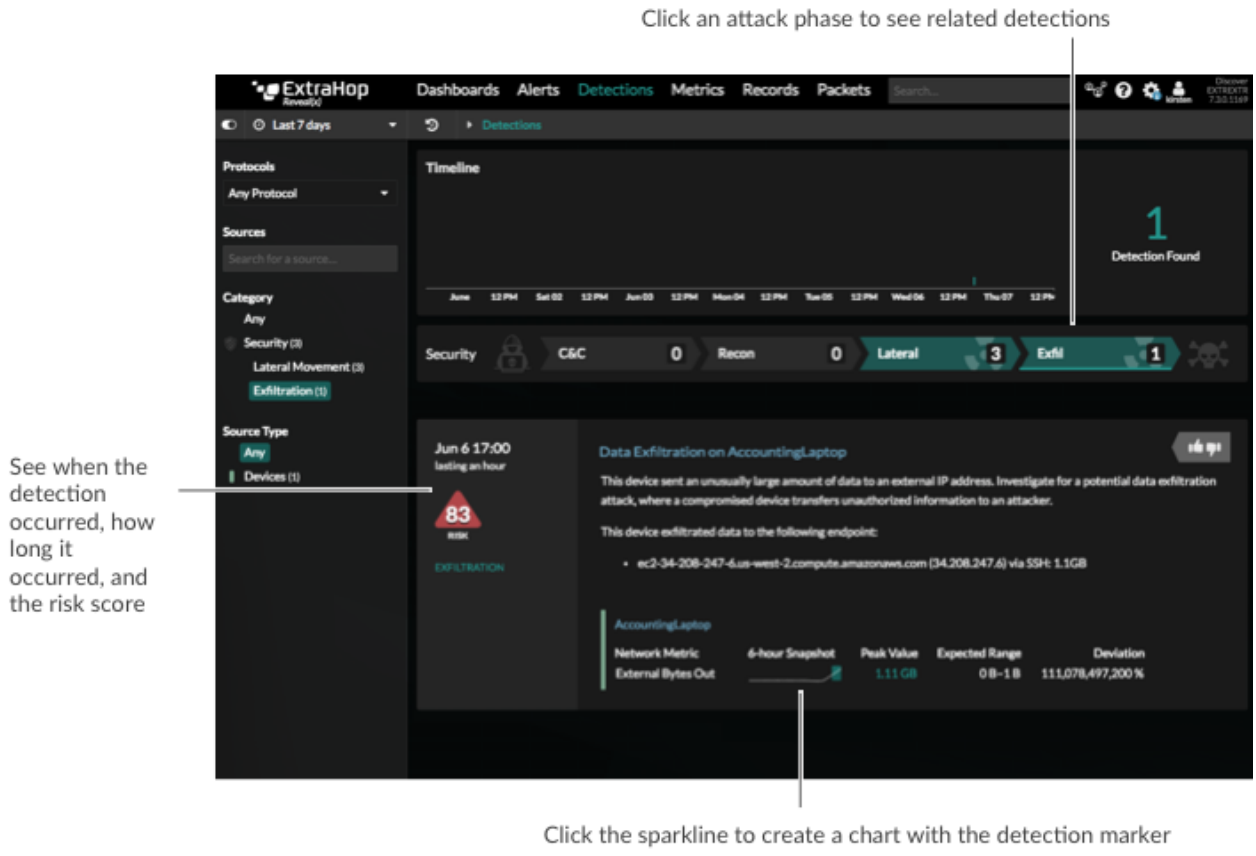

Attack chain

Most network attacks tend to follow familiar patterns or phases. These phases can be assembled into an attack chain to characterize the progression of an attack. Below the timeline chart on the Detections page, the attack chain highlights the number of detections that are associated with each attack phase, as shown in the following figure.

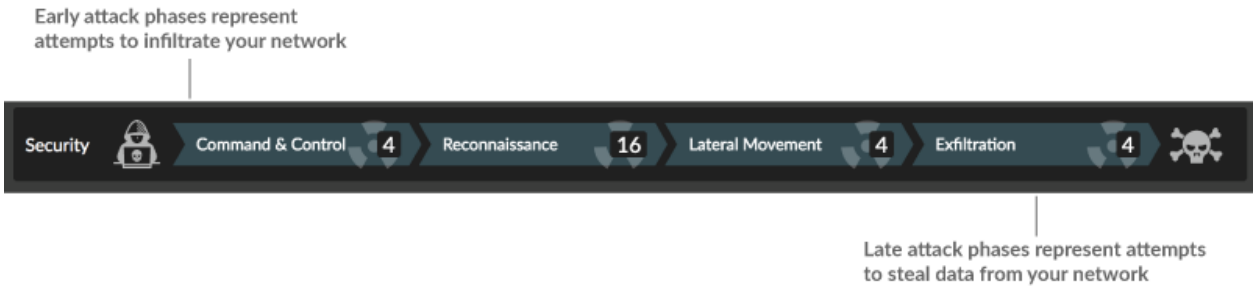

Important: Multiple detections in the attack chain can be associated with an attack. Detections associated with attack phases can be detected in any order.

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. Detections identify

when an internal device is communicating to a suspicious system outside of your network in support of an attack.

Reconnaissance

An attacker has compromised a device and initiates suspicious scans from that device to learn about your network. The attacker is looking for potential targets (critical assets) as well as attempting to gain direct control of resources. Detections identify when an internal device is performing suspicious scans of devices, ports, services, applications, or files on your network.

Lateral movement

An attacker is progressively moving through your network from device to device in search of data and critical assets that are ultimately the target of their attack campaign. Detections identify the unusual movement of users or data within your network.

Exfiltration

An attacker is attempting an unauthorized transfer of data from your network to a system that the attacker controls. Detections identify unusual transfers of data from devices within your network to external systems.

Performance (IT operation) detections

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

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

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

Authentication & Access Control

Detections identify unsuccessful attempts by users, clients, and servers to log in or access resources.

Database

Detections evaluate a suite of database protocols to determine whether your applications or users might be experiencing database access problems.

Desktop & App Virtualization

Detections identify when there are long Citrix load times or poor quality sessions for end users. SSH (secure shell) activity is also evaluated.

Network Infrastructure

Detections evaluate whether there are unusual events over the TCP, DNS, and DHCP protocols.

Service Degradation

Detections analyzes key protocols for Voice over IP (VoIP) and email communications within a network to identify service issues or performance problems.

Storage

Detections evaluate network file system traffic to determine whether users are having issues accessing specific files and shares.

Web Application

Detections analyze web traffic to find unexpected spikes in HTTP errors and warning codes. Poor web server performance is also analyzed.

Interpret detections

The Detections page displays the total number of detections for the selected time interval and details about each detection. 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 over time for the selected time interval. Each horizontal bar in the chart represents a single detection, so you can view the duration of each detection. Look for the tallest stack of bars to determine when the most detections occurred in the time interval. The total number of detections dynamically updates when you filter detections.

Tip: Hover over a bar to view the detection title, or click the bar to navigate directly to the detection detail page.

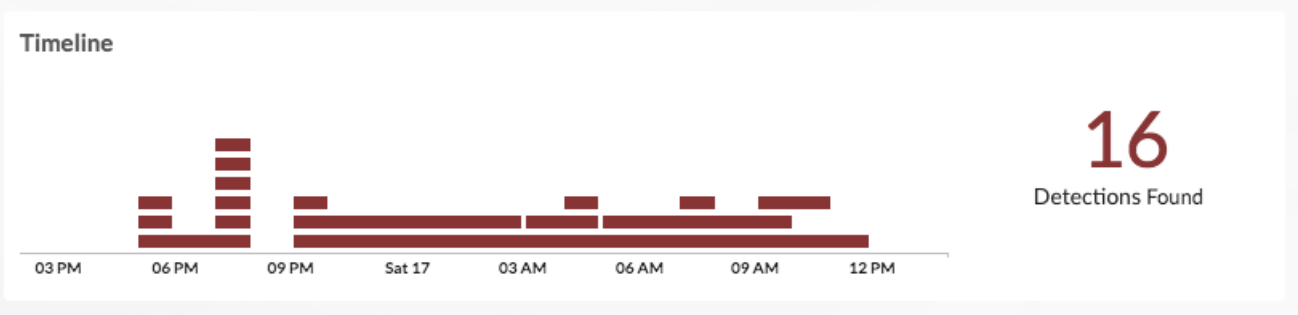

Click and drag across an area on the chart (which will become highlighted in green) to zoom in on a specific time range. The time interval dynamically updates to match the new time range in the chart, and details about each detection is displayed below the chart.

View details for individual detections

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

The following figure shows you what type of information is provided within an individual detection:

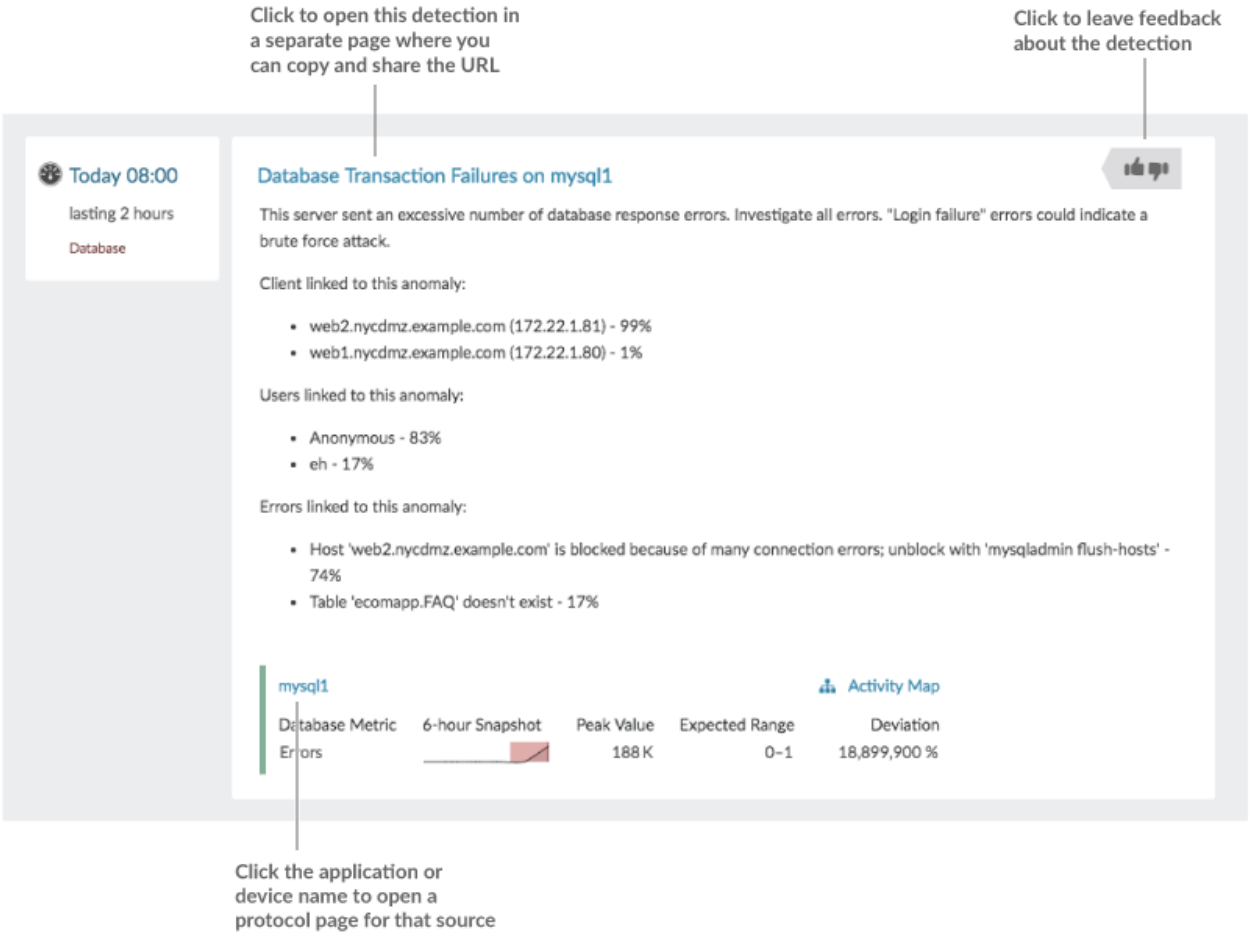

Title

The title includes the anomalous metric and the device or application name that is the cause of the detection. Click the title to share detection.

Description

The description provides information about what the detection means. For most detections, detail metrics are provided so you can immediately begin your investigation.

For more information, see Investigate detections.

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 displayed is ONGOING, the metric is being analyzed.

Risk score (ExtraHop Reveal(x) only)

Each detection has an associated risk score that can help you quickly identify urgent or critical detections in your environment. A risk score is displayed for each security detection and 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.

Sparkline

Sparklines are simple line charts that show you the metric behavior that led up to 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).

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.

Peak Value

The peak value is the maximum value from observed data that deviated from expected ranged for the duration of the detection.

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 changes in metric activity.

Deviation

A deviation is the quantity calculated to indicate the extent of change from an expected range.

Activity Maps

Click **Activity Map** to open an activity map that displays all of the L7 protocol activity and device connections to the client or server in the detection. For more information, see Activity maps.

Feedback

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 an detection more than once.

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 exhibits anomalous behavior such as deviating from the expected range of data by a significant amount. You can view analysis results about anomalies on the Detections page in the ExtraHop Web UI. If available, the following information is provided for each detection: the measured deviation (which is the difference between the observed value and the expected range), the detection value, and the expected range of normal metric values at the time of the detection.

Here is how detections are generally identified: the ExtraHop system generates metrics from wire data for the protocols, devices, and applications discovered on your network. 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 ExtraHop's unique domain expertise. 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 domain experts at ExtraHop

Detections also provide anomalous 50th percentile or 75th percentile values for a subset of metrics that account for server processing time.

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.

Because unusual behavior is detected in real time, you can identify and resolve a potential issue before it becomes a larger problem. You can also review historical detection data to investigate issues related to known security or network outage events that previously occurred.

Note: If you need to define a specific threshold value for an anomaly, such as a service level agreement (SLA), we recommend manually configuring an alert.

Related topics

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

- Find and filter detections
- Investigate detections
- Share a detection