Detections

Published: 2020-02-23

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).

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.

Here are important considerations about detections:

- 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.
- While some security detections can be identified immediately after connecting to the ExtraHop Machine Learning Service, the service continues to identify security detections as more historical data becomes available.
- You must have at least two weeks of wire data metrics stored on the ExtraHop system before performance 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.
- 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.

Detection categories

Depending on whether your license is activated on a Discover appliance or ExtraHop Reveal(x), your ExtraHop system will show detections about potential security or IT operations issues.

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 suspicious network behavior and security risks that are 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 response time and prevent disruptions from potential attacks.

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.

Risk score

Each security detection is assigned a risk score that can help you quickly identify urgent or critical detections in your environment. The risk score is displayed at the top of a detection card, similar to the following figure:

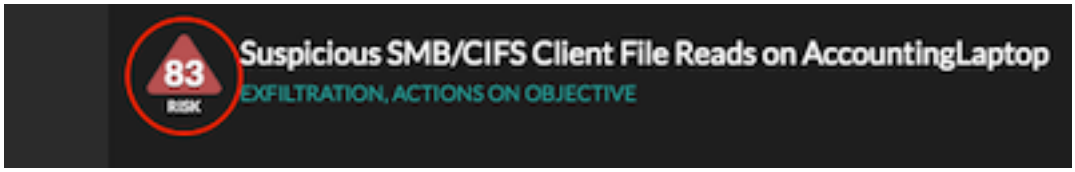

Each risk score is color coded by severity:

- 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.

Complexity

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

Business Impact

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

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 a category to all security detections. When the Group by filter is set to **None** 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.

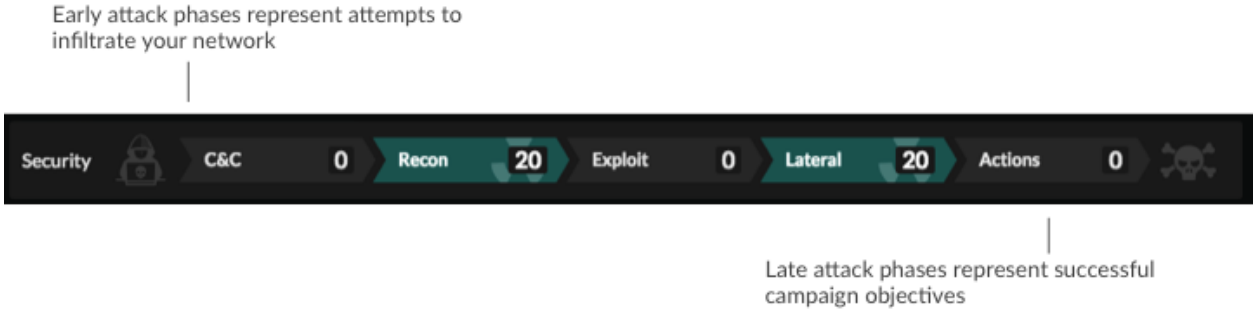

Important: Multiple detections in the attack chain can be associated with an attack. Detections associated with attack phases can be detected in any order.

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 external command and control (C&C) server that is controlled by the attacker. After the connection is established, the C&C server can send additional malware, commands, and payloads to support the attack. Reveal(x) detects when an internal device is communicating with a remote system that appears to be acting as a C&C channel for an attacker.

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 enumerate, or 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 known vulnerabilities on your network to actively exploit assets and vulnerabilities. For example, an attacker discovers a weak password after a brute force attack and then logs into an important file server or database. Or an attacker tries to cover their tracks by evading an intrusion detection system (IDS). Reveal(x) detects unusual and suspicious 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, denial of service attacks, or cryptomining. Reveal(x) detects when an attacker is close to completing a campaign objective.

Caution

Organizations should ensure that security policies and standards are enforced to mitigate risk. For example, proper access controls on file servers should be implemented and known software vulnerabilities should be addressed. Reveal(x) detects areas of risk on your network that might require attention. While this category is not located within the attack chain, this type of detection helps you find ways to mitigate risks associated with C&C, reconnaissance, exploit, lateral movement, and actions on objective attacks.

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.

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 authentication, authorization, and audit (AAA) protocols, 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 degradation associated with Voice over IP (VoIP), file transfer, 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 SMB/CIFS issues, or that network-attached storage (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.

Navigate the Detections page

The Detections page displays a list of all of the detections that occurred on your system and provides filtering and management options that help you find the detections that are most important to you. By default, detections are grouped by type and sorted by risk.

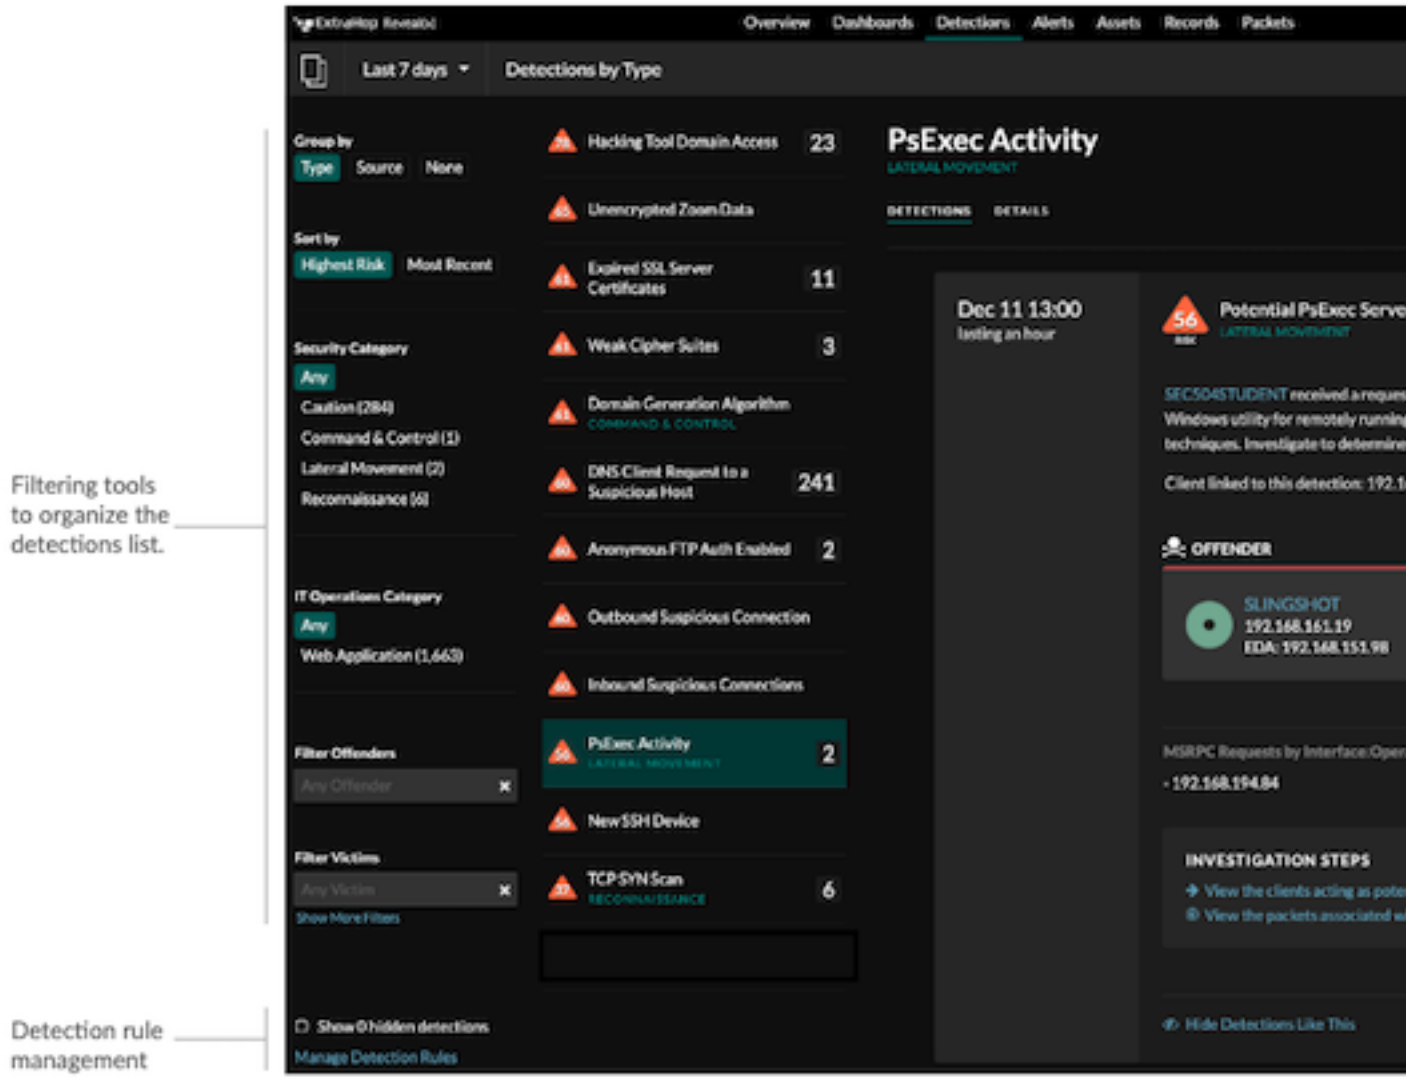

Filtering tools

You can filter the detections list to show only the detections that you want to see. For example, you might only be interested in exfiltration detections that occur over HTTP, or detections associated with specific devices. Here are some common ways that you can filter the detections page:

- Select the time interval
- Filter by source, protocol, and category
- Filter by offender or victim
- Group by source or title
- Sort by severity of the risk score (Reveal(x) only)

If you click a detection from a filtered list to view details, you can return to the same filtered list by clicking **Detections** in the navigation path at the top of the page.

Detection details

Certain detections offer additional details such as risk factors, attack backgrounds and diagrams, mitigation options, and reference links to industry security organizations such as MITRE.

You can find these details in the following locations:

- At the bottom of the detection card.
- By clicking **DETAILS** on the Detections page when Group by **Type** is selected.

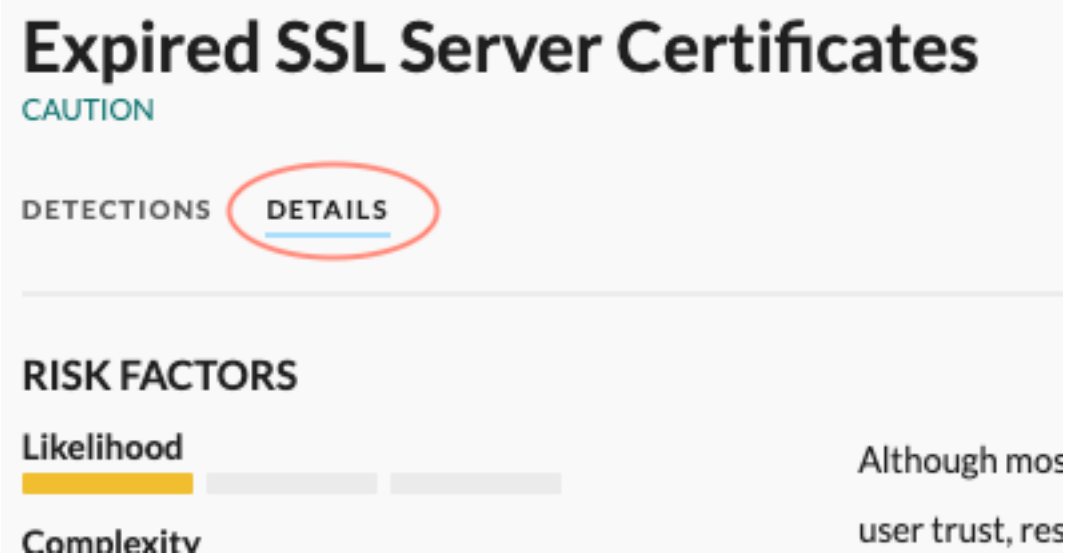

- Displayed next to the detection card on wide screens (1900 pixel or greater).

Detection rules

Detection rules enable you to hide detections that you don't want to see. For example, you might hide vulnerability scan detections over certain protocols because the scans are scheduled and running as planned.

Timeline chart

When the Group by filter is set to **None** on the Detections page, the Timeline chart displays the total number of detections identified within the selected time interval or applied filter. 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. Here are some ways that you can interact with the Timeline chart to get more information:

- Click and drag to highlight an area on the chart to zoom in on a specific time range. The detections list is filtered to the new time interval.
- Hover over a bar to view the detection title.
- Click a bar to navigate directly to the detection detail to view the timestamp.

Attack chain (Reveal(x) only)

Reveal(x) assigns an attack chain category to all security detections. When the Group by filter is set to **None** on the Detections page, the attack chain flow chart displays the number of detections that are associated with each attack phase. Click a phase in the attack chain to filter the detections list by that phase.

View a detection card

Each detection card identifies the cause of the detection, the detection category, and when the detection occurred.

The following figure shows an example of a detection card.

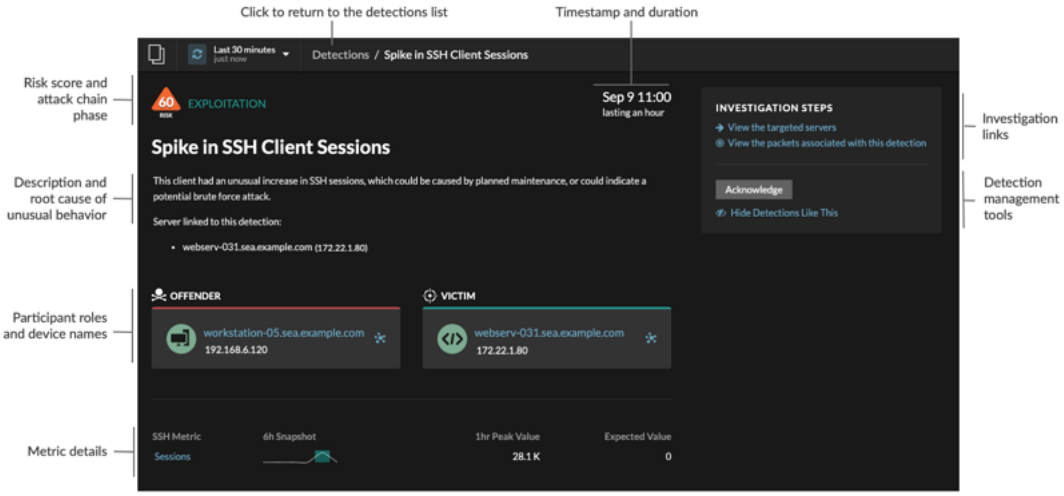

Metric data

When the unusual behavior is associated with a specific metric or key, the following metric data appears. Click the metric name or key (if available) to drill down from a detection.

- **Peak value:** 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.
- **Sparkline:** A simple line chart of metric activity before and during the detection.
- **Expected range or value:** Values that represent a normal level of activity, which is calculated based on two weeks of data. The expected range is the basis for comparison with observed values to detect abnormal changes in metric activity.

If metrics details are unavailable for the detection, the type of anomalous protocol activity is displayed.

Investigation Steps (Reveal(x) only)

A detection can include links to protocol pages, transaction-level records, packets, and activity maps. These links help you investigate detection data within the ExtraHop system. For more information, see Investigate security detections.

Participants

Endpoints associated with a detection are known as participants. A participant can be identified as a victim or offender, depending on the role the endpoint played in the detection; a role might not be identified for every participant. For example, ICMP-related detections always include a victim, but not an offender.

Note: For performance (IT operations) detections, a device identified as an offender is the likely source of an issue, such as a database server sending an excessive number of response errors. A device identified as a victim is usually negatively affected by the issue, such as clients experiencing slow or failed database transactions.

Each participant displays a device role icon, IP address, and hostname, if available. Click the device name to navigate to the device's Overview page. Click the activity map icon to view peer devices communicating with the participant.

Participants that are not on your network are identified as external endpoints; Overview page links, activity maps, and records are not available. You can hover over participants that are not on your network to see the geolocation of the IP address and a link to the ARIN Whois website.

Related detections timeline

Detection details can include a timeline of detections related to the current detection. A detection is considered related it meets one of the following criteria:

- An offender in the detection is also an offender in the current detection.
- A victim in an earlier detection became an offender in the current detection.

- An offender in a later detection was a victim in the current detection.

Each related detection participant included in the timeline displays the role and the role change relative to the current detection. For example, if the victim of the current detection becomes the offender in a later related detection, the timeline displays "Victim became offender" under the related detection.

The timeline displays the timestamp and duration of each detection, as well as when a related detection occurred before or after the current detection, in relative chronological order. The risk score and attack chain category are displayed for Reveal(x) only.

Click a related detection to go to the detail page for that detection. If the related detection occurs outside of the current time interval, the interval will not change unless you click an investigative link from the related detection details.

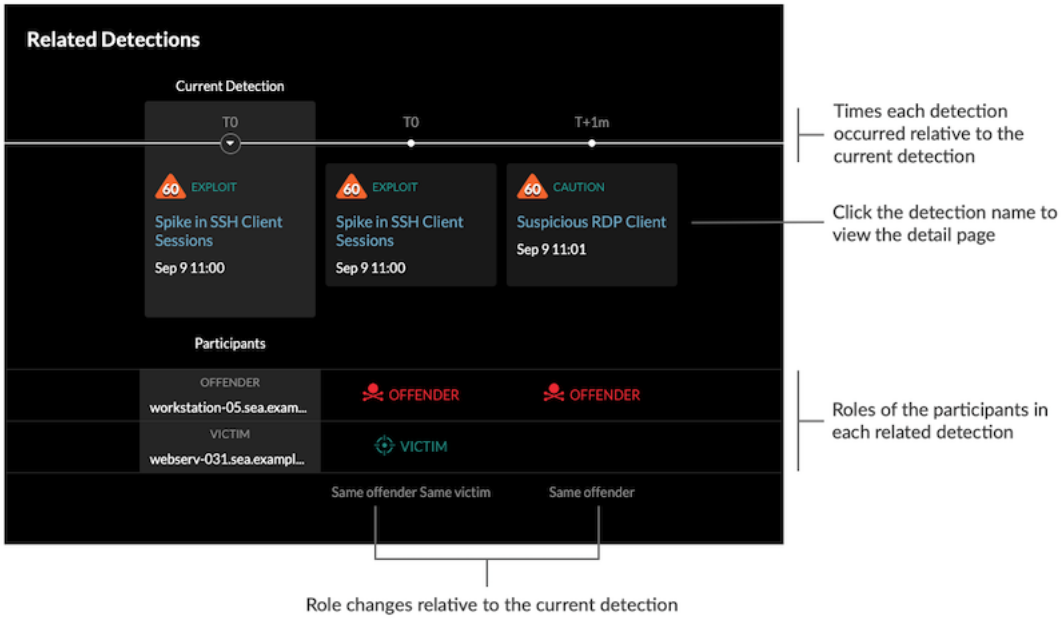

Duration

The duration of the detection indicates how long the unusual behavior was detected by the 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 behavior has not returned to a normal value or an anomalous event has not finished.

Detection management

Tools to manage and triage detections for investigation. You can acknowledge detections that you have reviewed or hide detections from view.

Find detections in the Web UI

While the Detections page provides quick access to all detections, there are indicators and links to detections throughout the Web UI.

- From the Activity page, click a detections link to go to the Detections page. The detections list is filtered to the associated protocol.
- From a device Overview page, click Detections to view a list of associated detections. Click the link for an individual detection to view detection details.
- From a device group Overview page, click the Detections link to go to the Detections page. The detections list is filtered to the device group as the source.
- From a device or device group protocol page, click the Detections link to go to the Detections page. The detections list is filtered to the source and protocol.

- On an activity map, click a device that displays animated pulses around the circle label to view a list of associated detections. Click the link for an individual detection to view detection details.
- From a chart on a dashboard or protocol page, hover over a detection marker to display the title of the associated detection or click the marker to view detection details.

Related topics

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

- Manage detections
- Investigate security detections
- Investigate performance detections
- Share a detection