

Anomaly detection with ExtraHop Addy

Published: 2018-02-08

The ExtraHop Addy™ service is a cloud-based service that applies machine learning techniques to automatically determine what is normal versus unusual behavior in your IT environment. Unlike other machine learning solutions that rely on logs or agent data, the Addy service applies machine learning technology to your wire data without requiring you to configure anything.

Addy learns about normal network behavior by analyzing the data stored on your Discover appliance. After Addy is activated, you can then browse detected performance and security anomalies in the ExtraHop Web UI and investigate root causes for issues on your network.

Overall, Addy offers 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 anomalies
- Find unknown performance issues, security issues, or infrastructure quirks
- Gain deeper insight into your network behavior

 **Important:** Addy does not analyze sensitive information and data types. For more information, download [ExtraHop Addy: Security Overview](#).

Here are important considerations about anomaly detection with Addy:


- You must have an ExtraHop Addy service license.
- You must have full system privileges, access to the Admin UI, and access through any firewalls to connect a Discover appliance to the Addy service through ExtraHop Cloud Services. For more information, see [Connect to the ExtraHop Addy service](#).
- You must have at least four weeks of wire data metrics stored on your Discover appliance before Addy can detect anomalies.
- On a Command appliance, you can access anomalies on a connected Discover appliance if that Discover appliance is connected to Addy.

Anomaly categories

Addy detects security or IT operations anomalies depending on the type of Addy license you have.

Security anomalies

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. The ExtraHop Addy service can detect anomalies associated with different phases of an attack. Addy tells you when a security risk occurred, which attack phase the risk is associated with, and which devices were affected by the risk.

 **Note:** This topic only applies to the ExtraHop Reveal(x) system edition.

Here are some important considerations about Addy security anomalies:

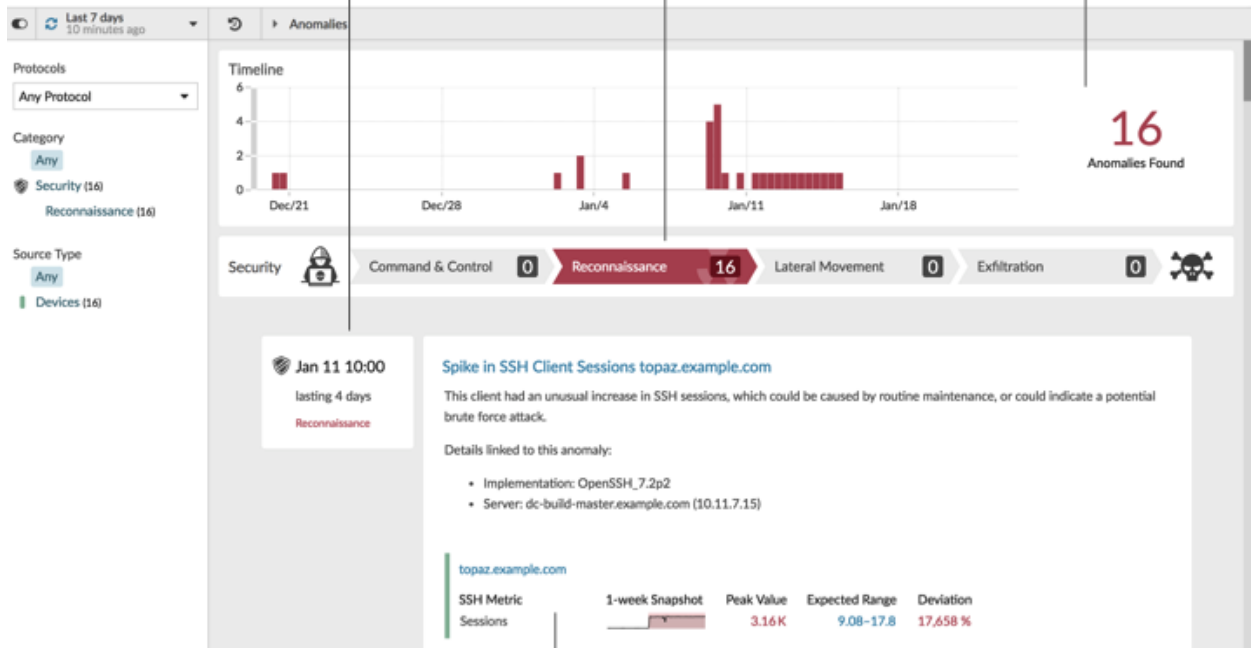
- You must have an ExtraHop Reveal(x) license to view security anomalies.
- Addy provides you with high-quality, actionable data about security risks. But these anomalies do not replace decision-making or expertise about your network. Always investigate anomalies to determine the root cause of the unusual behavior and when to take action.

When you log into the Web UI of your Security Edition and click **Anomalies**, an overview page appears with information about all the security anomalies detected during the selected [time interval](#).

See when the anomaly was detected and how long it occurred

See a timeline and summary of detected security anomalies

Click an attack indicator to see related anomalies



See details about anomalous protocols and metrics

Below the timeline chart, an attack chain highlights the number of anomalies that are associated with a different attack phase, as shown in the following figure.

Early attack phases represent attempts to infiltrate your network



Late attack phases represent attempts to steal data from your network

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

Addy detects the following security risk anomalies:

Command and control

An attacker has infiltrated a device on your network, and that device attempts to phone home to the attacker's command and control (C&C) server. The C&C server can then send malware or a payload to the device to gain control of that device. Addy detects when an internal device is frequently connecting to a suspicious server or client outside of your network.

Reconnaissance

An attacker has infiltrated a device on your network and is trying to learn about your network. Specifically, the attacker is looking for potential targets and associated vulnerabilities. Addy detects when an internal device is performing suspicious scans of devices, ports, services, applications, or files on your network.

Lateral movement

An attacker is gaining access to multiple devices within your network by determining valid user credentials. The attacker can then move between devices, or move data between devices on your network. Addy detects unusual movement in the network and unusual activity around the movement of data between devices.

Exfiltration

An attacker is attempting to transfer data from your network to a server outside of your network that an attacker controls. Addy detects when a device is sending an unusual amount of data to a suspicious server or client outside of your network.

IT operations anomalies

Addy is always on and always learning about network behavior across your IT infrastructure. Addy automatically surfaces network, application, and infrastructure problems and their root causes, so you can immediately focus on issues that matter.

Here are some important considerations about this type of anomaly:

- You cannot view IT operations anomalies if you have a ExtraHop Reveal(x) license.
- The Addy service provides you with high-quality, actionable data about potential performance and operation issues. But these anomalies do not replace decision-making or expertise about your network. Always investigate anomalies to determine the root cause of the unusual behavior and when to take action.

Addy detects anomalies in the following operational categories:

Authentication and authorization

Addy detects unsuccessful attempts by users, clients, and servers to log in or access resources.

Database

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

Desktop and app virtualization

Addy detects when there are long Citrix load times or poor quality sessions for end users. Addy also evaluates SSH (secure shell) activity.

Network infrastructure

Addy evaluates whether there are unusual events over the TCP, DNS, and DHCP protocols.

Service degradation

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

Storage

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

Web server

Addy analyzes web traffic to find unexpected spikes in HTTP errors and warning codes. Addy also detects poor web server performance.


Interpret anomalies

The Anomalies page displays the total number of anomalies for the selected time interval and details about each detected anomaly. The following sections show you what information you can learn from anomalies.


View total anomalies over time

The Timeline chart provides a summary of detected anomalies (y-axis) over time (x-axis) for the selected time interval. Each bar in the chart represents the total number of concurrent, active anomalies that were detected during a specific time period. Look for the tallest bar to determine when the most anomalies occurred in a time period.

Hover over a bar to view information, such as date, time, and the number of detected anomalies for a specific time period.



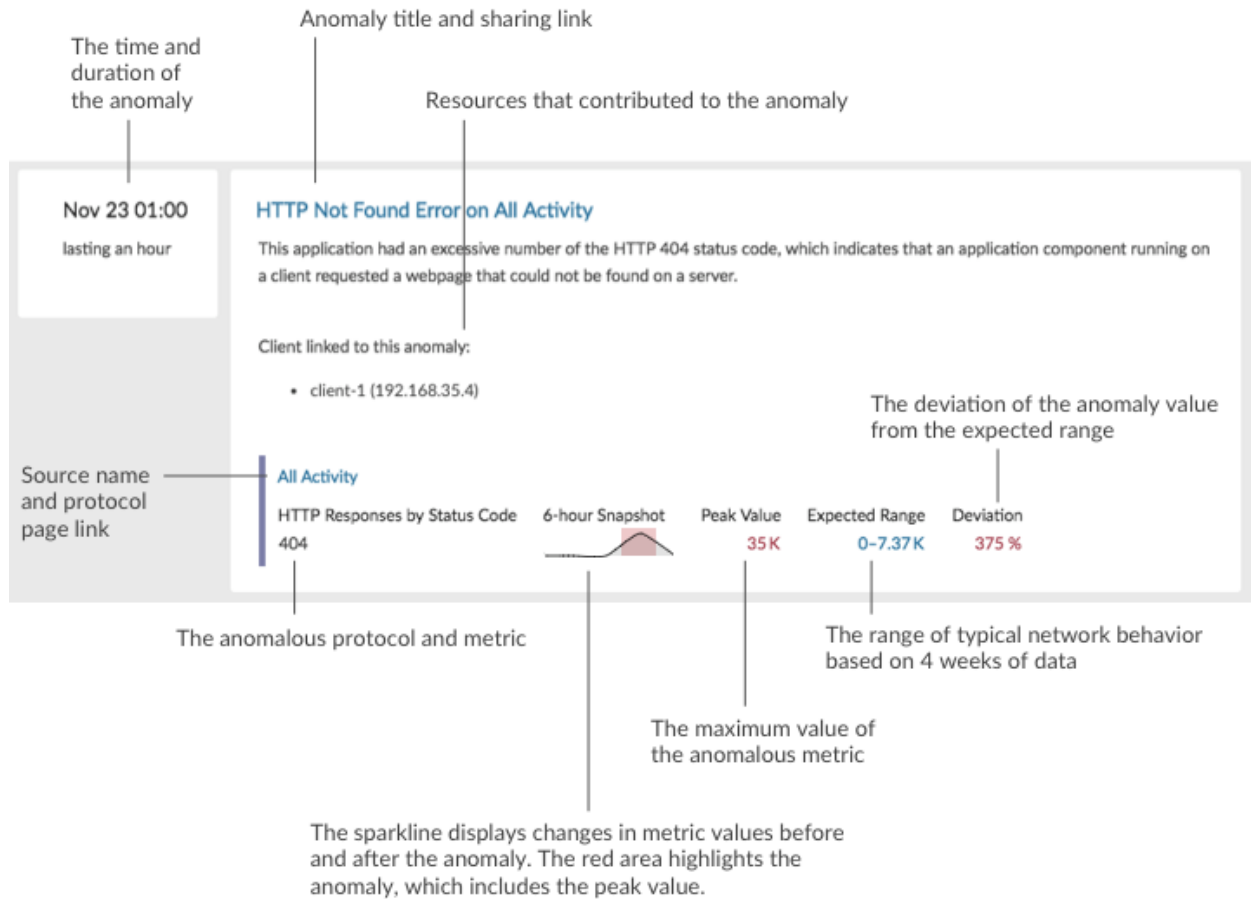
Click and drag across an area on the chart (which will become highlighted in green) to zoom in on a specific time range. The time interval in the Discover or Command appliance dynamically updates to match the new time range in the chart, and details about each anomaly that was detected in that time range are displayed below the chart.



View details for individual anomalies

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

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



Title

The title includes the anomalous metric and the device or application name that is the cause of the anomaly. Click the title to [share an anomaly](#).

Description

The description provides information about what the anomaly means. For most anomalies, Addy automatically surfaces detail metrics identified with Addy's machine learning capabilities, so you can immediately begin your investigation.

For more information, see [Investigate anomalies with Addy](#).

Duration

The duration of the anomaly indicates how long the anomalous value was detected by Addy.

The minimum duration of an anomaly is one hour, because Addy detects anomalies by analyzing metric data with 1-hour granularity. If the duration value is displayed as ONGOING, the anomalous metric is in the process of being detected.

Sparkline

Sparklines are simple line charts that show you the metric behavior that led up to the anomaly. The sparkline charts display a snapshot of metric data from the time frame around the duration of the detected anomaly (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 value from observed data that deviated from expected ranged for the duration of the anomaly.

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 by the Addy machine learning engine to indicate the extent of change from an expected range.

How the ExtraHop Addy service works

This section provides some background information on how the ExtraHop Addy service identifies anomalies.

Anomalies are unexpected deviations from normal patterns in device or application behavior. Addy detects anomalies from stored Discover appliance data with a proprietary algorithm that combines time series decomposition, unsupervised learning, heuristics, and ExtraHop's unique domain expertise. This combination helps to ensure that detected anomalies are both accurate and actionable. By detecting an anomaly as soon as it happens, you can identify and resolve a potential issue before it becomes a larger problem. You can also review historical anomaly data to investigate issues related to known security or network outage events.

In most network monitoring tools, anomalies are 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. Addy automatically delivers consistent and accurate results about anomalous metrics and protocols without requiring manual configuration for individual devices. The Addy machine learning engine analyzes the historical behavior of individual devices, and automatically adapts to each device across time when there are changes to the expected range of data in your network.

Here is how Addy anomaly detection generally works: the metrics that the Addy machine learning engine analyzes come from wire data that is collected by your Discover appliance. The Discover appliance processes this data, generates metrics, and associates the metric data with protocols, devices, and applications. Addy retrieves a subset of protocol metrics from the Discover appliance to analyze and report results about detected anomalies.

The algorithm that drives the machine learning engine in Addy evaluates unique information about your environment to calculate the expected range of normal network behavior and then adapts to changing variations in protocols and metric data. Outliers, or anomalies, are then detected based on three variables:

- Observed data, collected in real-time by the Discover appliance
- Expected range data, calculated from four weeks of historical data collected by the Discover 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





Note: If you need to define a specific threshold value for an anomaly, which might be associated with a service level agreement (SLA) for example, we recommend manually configuring an alert in the Discover appliance.

Essentially, an anomaly is detected when observed data deviates from the expected range of data by a significant amount. You can then view analysis results about anomalies on the Anomalies page in the Web UI of the Discover appliance. For each anomaly, Addy provides the measured deviation (which is the difference between the observed value and the expected range), the anomaly value, and the expected range of normal metric values at the time of the anomaly.

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

Related topics

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

- [Connect to the ExtraHop Addy service](#) 
- [Find and filter anomalies](#) 
- [Investigate anomalies with Addy](#) 