



ExtraHop Addy User Guide

© 2018 ExtraHop Networks, Inc. All rights reserved.

This manual in whole or in part, may not be reproduced, translated, or reduced to any machine-readable form without prior written approval from ExtraHop Networks, Inc.

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

Published: 2018-07-07

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

Contents

Anomaly detection with ExtraHop Addy	4
Connect to the ExtraHop Addy service	4
Anomaly categories	5
Security anomalies	5
IT operations anomalies	7
Find and filter anomalies	8
Interpret anomalies	9
Investigate anomalies	11
Open an activity map from an anomaly	12
Navigate to a protocol page	14
Best practices for investigating anomalies	15
Share an anomaly	16
Appendix	16
How the ExtraHop Addy service works	16

Anomaly detection with ExtraHop Addy

The ExtraHop Addy™ service is a cloud-based service that applies machine learning techniques to wire data to automatically determine what is expected versus unusual behavior in your IT environment. Unlike other machine learning solutions that rely on logs or agent data, Addy applies machine learning technology 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 or an ExtraHop Reveal(x) license.
- You must have full system privileges, access to the Admin UI, and access through any firewalls to connect a Discover or ExtraHop Reveal(x) 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 or ExtraHop Reveal(x) appliance before Addy can detect anomalies.
- When you create a user account with restricted read-only privileges, those users can only view the metrics in the dashboards that you share with them. Those users will be unable to view anomalies. For more information, see [Share a dashboard with a restricted user](#).
- On a Command appliance, you can access anomalies on a connected Discover appliance or ExtraHop Reveal(x) if that appliance is connected to Addy.

 **Note:** A Command appliance can only connect to either Discover appliances or ExtraHop Reveal(x) appliances.

Connect to the ExtraHop Addy service

The ExtraHop Addy service is a cloud-based service that applies machine learning techniques to automatically determine what is normal behavior and what is unusual behavior in your IT environment. After you acquire a license for the Addy service, the license status and ExtraHop Cloud Services settings are automatically updated on your Discover appliance.

Before you begin

- You must apply the Addy license before you can connect to the service. See the [License FAQ](#) for more information.
 - You must have [full system privileges](#) to access the ExtraHop Admin UI and to connect to ExtraHop Cloud Services.
1. Log into the ExtraHop Admin UI on the Discover appliance.
 2. In the Network Settings section, click **ExtraHop Cloud Services**.
 3. Click **Terms and Conditions** to read the content.
 4. After becoming familiar with the Addy service terms and conditions, select the checkbox.

5. Click **Connect to ExtraHop Cloud Services**.

After you are connected, the page updates to show status and connection information. Learn how to navigate and interpret anomalies in [Anomaly detection with ExtraHop Addy](#).

If the connection fails, there might be an issue with your firewall rules. See [Troubleshoot your connection to the Addy service](#) to identify and resolve the issue. If connection problems persist, contact ExtraHop Support for help by [creating a case on the Customer Portal](#) (requires login).

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 ExtraHop Reveal(x) 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

Click an attack phase to see related anomalies

See a timeline and summary of detected security 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.

Find and filter anomalies

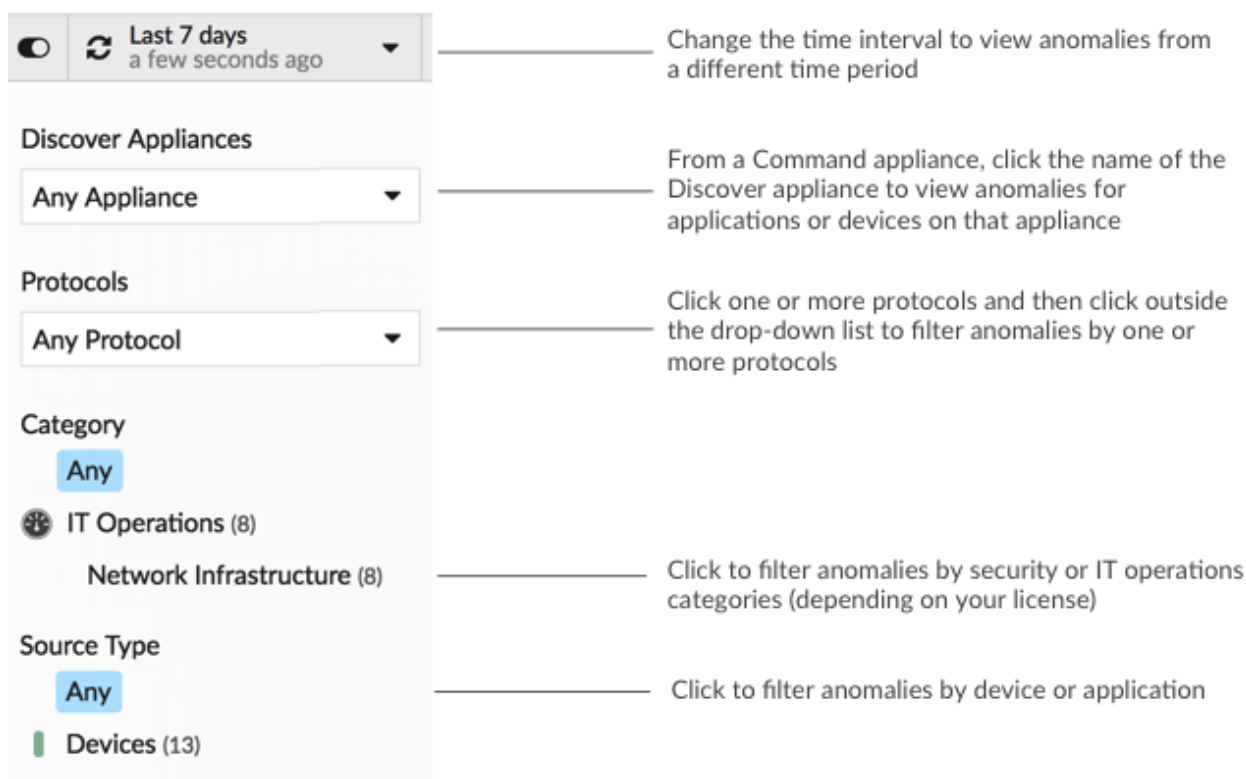
After activating Addy, a top menu item appears for anomalies. To browse anomalies detected by Addy, log into the Web UI and click **Anomalies** at the top of the page. You can then filter anomalies by time interval, protocol, category, applications, or devices. Anomalies are sorted by their start time and the most recent anomaly is listed first.

 **Note:** Configuring an anomaly alert from the Alerts page lets you monitor alerts or receive email notifications when a specific anomaly is detected. For more information, see the following topics:

- [Configure Addy anomaly alert settings](#)
- [Add a notification to an alert configuration](#) to receive emails when an anomaly is generated
- Monitor alerts from the [Alert History](#) page

The following steps show you how to find and filter anomalies:

1. Log into the Web UI on the Discover appliance, Command appliance, or ExtraHop Reveal(x) and then click **Anomalies** at the top of the page.
A list of anomalies for the current time interval appears. If the list is empty, then Addy has not detected anomalies for the selected time interval.
2. In the left pane, filter anomalies by selecting the options as shown in the following figure:



The screenshot shows a sidebar with the following filters and their corresponding callouts:

- Time Interval:** A dropdown menu currently set to "Last 7 days a few seconds ago". Callout: "Change the time interval to view anomalies from a different time period".
- Discover Appliances:** A dropdown menu currently set to "Any Appliance". Callout: "From a Command appliance, click the name of the Discover appliance to view anomalies for applications or devices on that appliance".
- Protocols:** A dropdown menu currently set to "Any Protocol". Callout: "Click one or more protocols and then click outside the drop-down list to filter anomalies by one or more protocols".
- Category:** A list of buttons: "Any" (selected), "IT Operations (8)", and "Network Infrastructure (8)". Callout: "Click to filter anomalies by security or IT operations categories (depending on your license)".
- Source Type:** A list of buttons: "Any" (selected) and "Devices (13)". Callout: "Click to filter anomalies by device or application".

Next steps

- [Investigate anomalies](#)

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 displays the total number of detected anomalies over time for the selected time interval. Each horizontal bar in the chart represents a single anomaly, so you can view the duration of each anomaly. Look for the tallest stack of bars to determine when the most anomalies occurred in the time interval. The total number of anomalies dynamically updates when you [filter anomalies](#).



Tip: Hover over a bar to view the anomaly title, or click the bar to navigate directly to the anomaly 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 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:

Click to open this anomaly in a separate page where you can copy and share the URL

Click to leave feedback about the anomaly

Today 08:00
lasting 2 hours
Database

Database Transaction Failures on mysql1

This server sent an excessive number of database response errors. Investigate all errors. "Login failure" errors could indicate a brute force attack.

Client linked to this anomaly:

- web2.nycdmz.example.com (172.22.1.81) - 99%
- web1.nycdmz.example.com (172.22.1.80) - 1%

Users linked to this anomaly:

- Anonymous - 83%
- eh - 17%

Errors linked to this anomaly:

- Host 'web2.nycdmz.example.com' is blocked because of many connection errors; unblock with 'mysqladmin flush-hosts' - 74%
- Table 'ecomapp.FAQ' doesn't exist - 17%

mysql1

Database Metric	6-hour Snapshot	Peak Value	Expected Range	Deviation
Errors		188 K	0-1	18,899,900 %

Activity Map

Click the application or device name to open a protocol page for that source

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](#).

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.

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 anomaly. For more information, see [Activity maps concepts](#).


Feedback


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

 **Note:** The option to provide feedback is determined by user privileges, which are assigned by the ExtraHop administrator. For more information, see the [User privileges](#) section in the ExtraHop Admin UI Guide.

Investigate anomalies

When you find an interesting anomaly, you want to quickly understand the root cause. Addy performs an automated investigation for most anomalies, which means that you can view detail metrics in the anomaly description to immediately learn what contributed to an issue. When multiple factors contribute to an anomaly, you can also see the percentage of their contribution to the anomaly. For example, the following figure show which client, server, and URI are linked to an HTTP 404 anomaly.



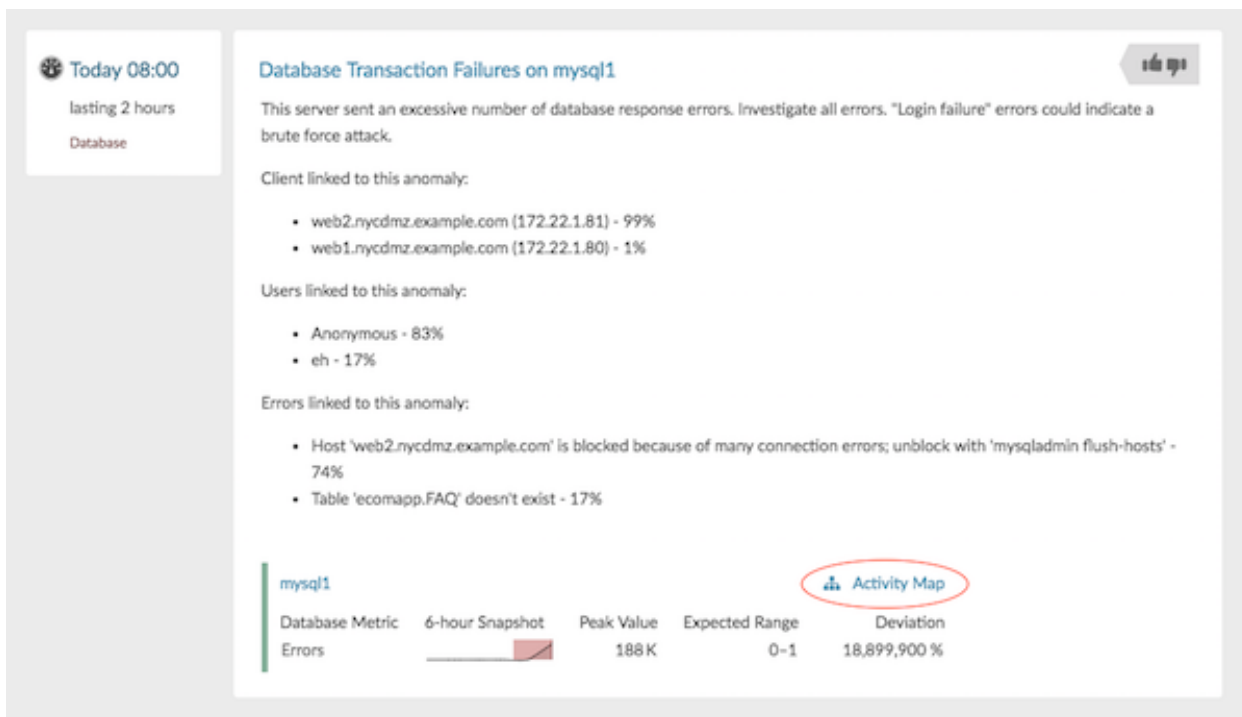
 **Note:** Automated investigation is not available for server processing time anomalies. For these anomalies, you can [investigate anomalies from a protocol page](#) in the Discover or Command appliance.

To learn more about the scope of an anomaly on your network, you can continue your investigation by opening an activity map or visiting a protocol page.

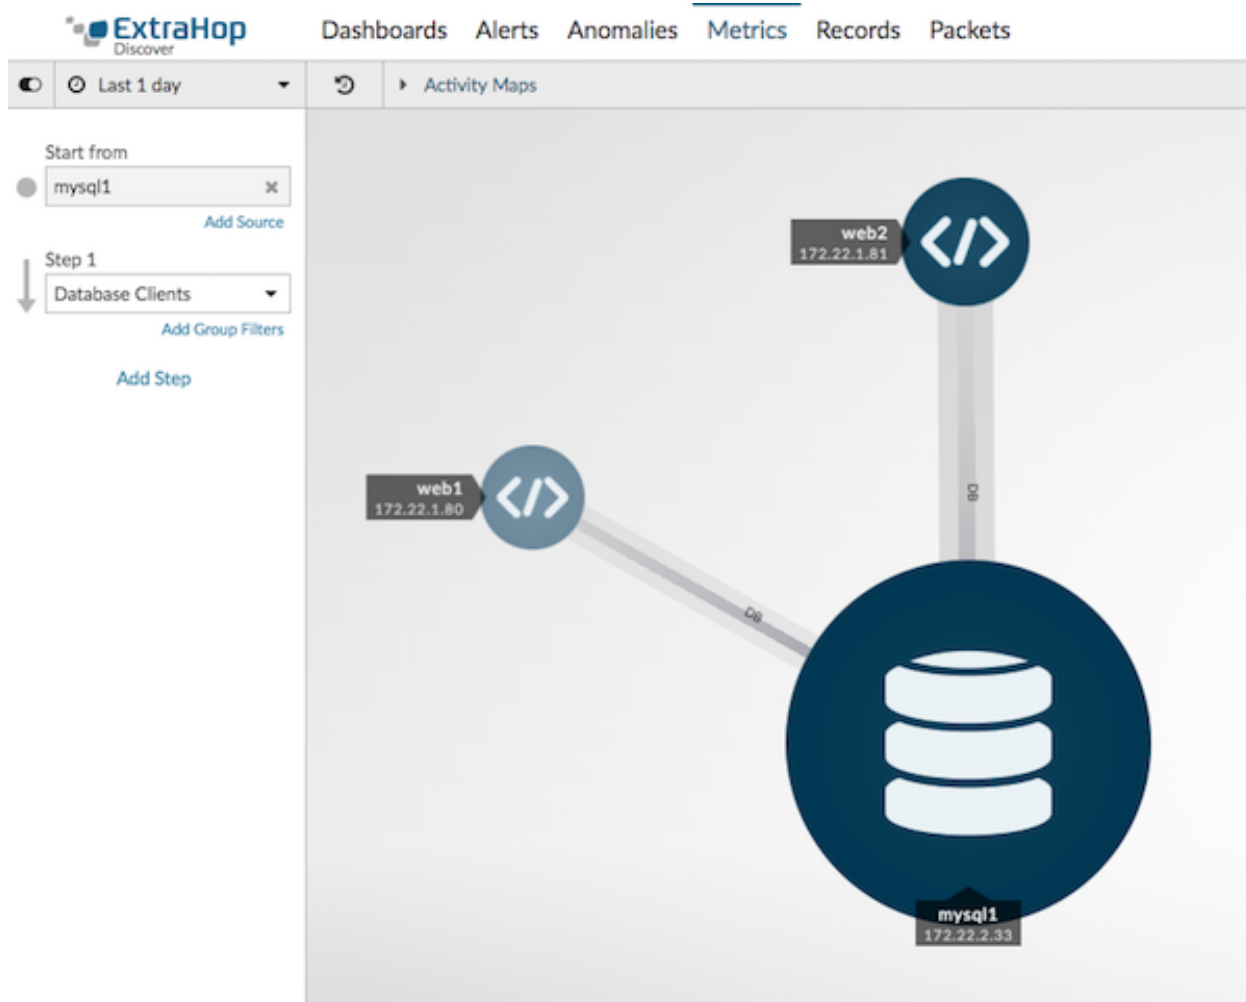
Open an activity map from an anomaly

When a single client or server is associated with unusual L7 protocol activity, such as a high number of HTTP errors or DNS request timeouts, an activity map link appears.

1. Log into the Web UI on the Discover or Command appliance and then click **Anomalies** at the top of the page.
2. Find the anomaly that you want to investigate. The following figure shows an example of the **Activity Map** link for a database server that sent an unusual number of errors.

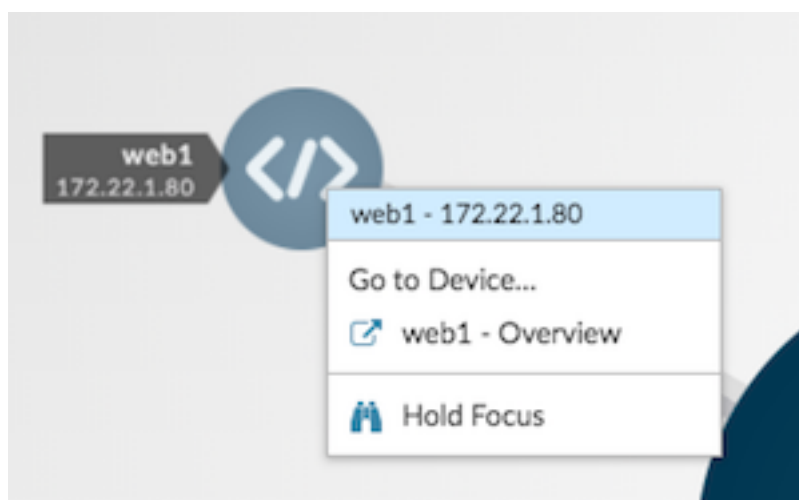


3. Click **Activity Map**.
An activity map appears for the database server. The activity map in the following figure shows the two database clients that were connected to the server during the anomaly time frame.

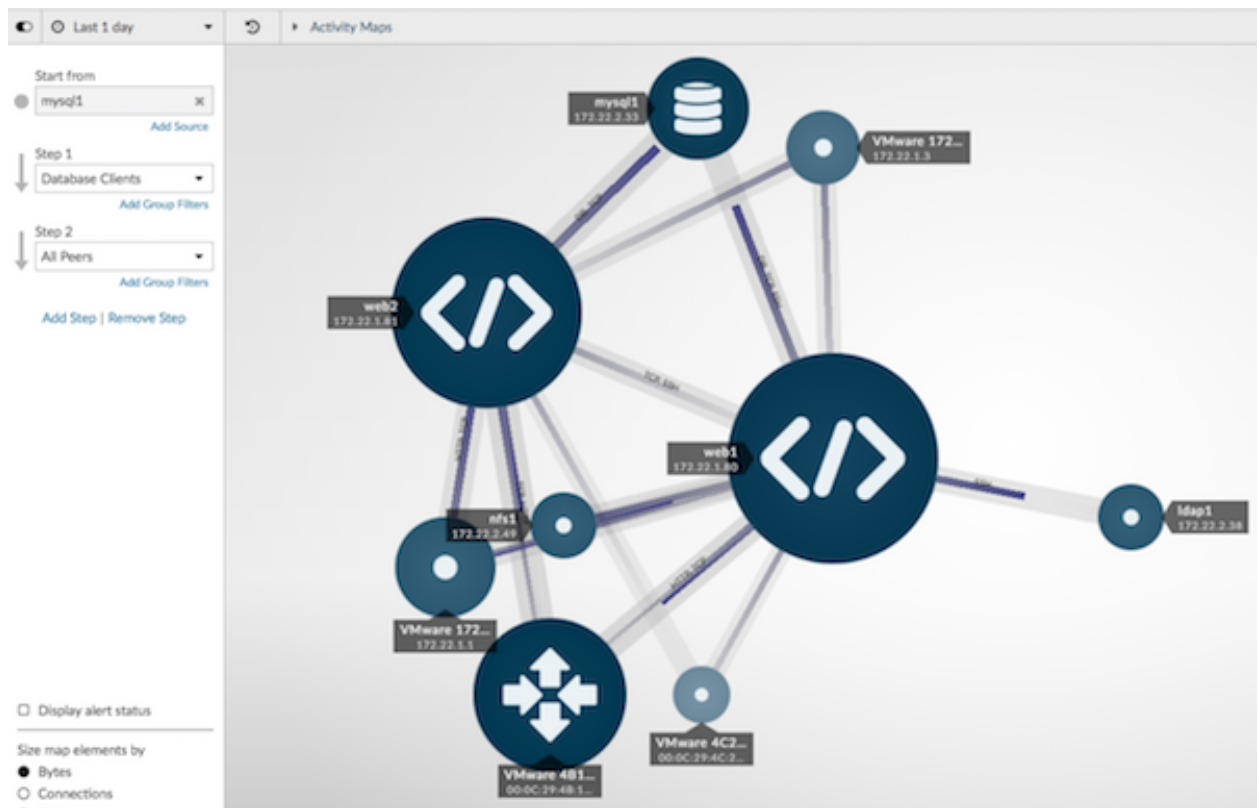


You can now interact with the activity map to learn more about the effect of the database errors across the network:

- Click any client in the map to access a menu that contains a Go to Device... link. Click the link to open a protocol page with client metrics, such as requests and responses.



- In the left pane below Step 1, click **Add Step** and then click **All Peers** in the drop-down list. The map updates to show you which downstream devices are connected to the database clients, as shown in the following figure.



- Save and then share your activity map with other ExtraHop users.

For more information about activity maps, see [Activity maps concepts](#).

Navigate to a protocol page

If you want to further investigate anomalous metrics, you can navigate to a protocol page where you have access to additional charts, metrics, and tools.

- Log into the Web UI on the Discover or Command appliance and then click **Anomalies** at the top of the page.
- Find the anomaly that you want to investigate.
- Click the source name, as shown in the following figure.

Dec 15 12:00

lasting an hour


CIFS Client Access Denied Errors on VMware 192.168.6.183

This client received an excessive number of errors with the SMB status code, STATUS_LOGON_FAILURE. This anomaly indicates that a user is trying to log in with an incorrect username or password. Investigate for a potential brute force attack.

VMware 192.168.6.183 Activity Map

CIFS Errors by Error	6-hour Snapshot	Peak Value	Expected Range	Deviation
STATUS_LOGON_FAILURE		9	0-1	800 %

The anomalous protocol page for the device or application appears, which displays all of the metric data associated with that specific device or application during the anomaly time interval, as shown in the figure below.



Next steps

From a protocol page, you can then choose one of the following options to further investigate metric data:

- [Create an activity map](#)
- [Drill down on metrics](#)

Best practices for investigating anomalies

Addy provides you with high-quality, actionable data about anomalies—but does not replace decision-making or expertise about your network. The following best practices explain how to determine which anomalies are worth further investigation and when to take action.

Change the time interval to see when anomalies occurred

Learn if anomalies occurred before, after, or during a reported problem. For example, does the time frame of the anomaly coincide with a reported issue, such as slow load times or login times? You can also compare anomalies from the past month to the current date, which gives you a sense of whether the occurrence or severity of anomalies is changing over time.

For more information, see [Find and filter anomalies](#).

Create an anomaly alert

You can configure an alert to receive email notifications when an anomaly occurs. Anomaly alerts also help you quickly find anomalies for a specific device or application on the [Alert History](#) page.

For more information, see [Configure Addy anomaly alert settings](#).

Filter anomalies by protocol

Filter by protocol to quickly monitor critical protocols with a role in security, commerce, or communication processes.

For example, an FTP 530 error anomaly might indicate that someone is trying to gain unauthorized access to information on your network. Or Citrix server and client latency anomalies might indicate that users are experiencing long load times for their roaming desktop profiles.

Selecting different protocols can also show you how anomalies correlate to each other. An anomalous HTTP response time followed immediately by an anomalous CIFS server processing time might suggest that web servers are dependent on how quickly your file storage servers can send and receive file data.

For more information, see [Find and filter anomalies](#).

Share an anomaly


If you find an anomaly that you want to share, you can send the URL from the anomaly detail page to other ExtraHop users.

Here are some important considerations about sharing anomalies:

- You must copy and share the entire URL.
- To view the anomaly, an ExtraHop user must have access to the Discover appliance or ExtraHop Reveal(x) where the anomaly was detected.

The following steps show you how to select and share an anomaly:

1. Log into the Web UI on the Discover or Command appliance where Addy is licensed and then click **Anomalies** at the top of the page.
2. Find the anomaly that you want to share.
3. Click the anomaly title, as shown in the following figure.



4. At the top of the browser, copy the entire URL. You can now send this URL to other ExtraHop users who have access to the Discover appliance where the anomaly was detected.

Next steps

- [Configure an anomaly alert](#) to receive email notifications about an anomaly.

Appendix

The following section contains reference information about the Addy service.

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.