



Investigator documentation

Updated Jun 06, 2024

CONTENTS

1	Release notes	2
1.1	June 2024	2
1.2	May 2024	2
1.3	April 2024	3
1.4	March 2024	3
1.5	February 2024	4
1.6	January 2024	4
1.7	December 2023	4
1.8	November 2023	5
1.9	October 2023	5
1.10	September 2023	5
1.11	August 2023	6
1.12	July 2023	6
1.13	June 2023	7
1.14	May 2023	7
1.15	April 2023	7
1.16	March 2023	8
1.17	February 2023	8
1.18	January 2023	8
1.19	December 2022	9
1.20	November 2022	9
1.21	October 2022	9
1.22	September 2022	10
1.23	August 2022	10
2	Quickstart	12
2.1	Before you start	12
2.2	Log in for the first time	12
2.3	Configure your sensors	13
3	Investigator overview	15
3.1	Alerts	16
3.2	Detections	16
3.3	AI-generated content	16
3.4	Analytics engine	17
3.5	Roles	17
4	Security overview	18
5	Work with alerts & detections	20

5.1	Identify and manage threats through the Detections page	20
5.2	Machine learning detections	29
5.3	Search-based alerts	35
5.4	Normalized severity scores	43
5.5	CrowdStrike data in detections	43
6	Explore data through Dashboards	47
7	Find details in the logs	50
7.1	Search the logs	50
8	Federated tenants	53
9	Account settings	56
9.1	Account alias	56
9.2	Cookies	56
9.3	Password	57
9.4	Two-factor authentication	57
9.5	Theme settings	59
10	System settings	61
10.1	Licensing	61
10.2	Autoclose detections	62
10.3	Alert Catalog	63
10.4	Audit activities through logs	67
10.5	Sensor monitoring and management	69
10.6	User management	69
10.7	GPT settings	73
10.8	ServiceNow integration	73
10.9	CrowdStrike EDR integration	75
10.10	Alert Exports	76

Step through these documentation topics to learn about the Corelight Investigator features and functionality.

Investigator Overview Get to know the core concepts and definitions behind Corelight Investigator.

Overview

Quickstart guide Connect your sensors to Investigator and log in for the first time.

Get started

Security overview Learn about the Investigator landing page and how it provides a high-level summary of what's going on in your network.

Security overview

Alerts & Detections Explore and manage alerts and detections.

Alerts and detections overview

Dashboards See how dashboards can help you monitor events and explore data.

Dashboards

Log search Jump in to the deep end and see what you can do with searchable access to the data ingested from logs.

Log search

Account settings Learn how to change your password, reset 2FA, and update your alias.

Account settings

System settings Learn how an admin user can manage users, manage sensors, and configure alert exports.

System settings



PDF Version of Documentation

Contents

RELEASE NOTES

Here's what's new with Investigator.

1.1 June 2024

1.1.1 Fixed bugs

- Fixed an issue where Investigator incorrectly reported the number of sensors sending data. (INVEST-9649)
- Removed the character limit for the Category filter on the Related Detections page. (INVEST-9518)
- Fixed an issue that prevented the ability to copy the MFA code during initial setup. (INVEST-10013)
- Updated the behavior of buttons in the UI so you can right-click a button and open a new tab. (INVEST-10016)
- Corrected an issue that prevented the tenant display name from appearing in the page header. (INVEST-10017)

1.2 May 2024

1.2.1 Features and enhancements

- Introduced support for *federated tenants* to provide administrators with a collective view of data from configured sub-tenants (child tenants). The aggregated data provides comprehensive insights into network security within a unified interface and streamlines the management of tenants.

1.2.2 Fixed bugs

- Resolved an issue where LogScale dashboards returned a 404 error when opened in a new tab. (INVEST-9772)
- Fixed a content refresh issue related to LogScale tokens. (INVEST-9757)
- Fixed an issue where reactivated users encountered an invalid OTP error despite entering the correct OTP. (INVEST-9624)
- Corrected the icon representation for domain type entities in the excluded entities data table. (INVEST-9571)
- Resolved an issue where the toast message indicating an unsuccessful save did not disappear automatically while saving alert exporters. (INVEST-9569)
- Fixed issues related to the quick view functionality in the detection page. (INVEST-9548)

- Corrected UI distortion in the timeframe field and ensured proper detection listing when navigating from the Alert Catalog to the Detections page, particularly when a detection shares an alert name. (INVEST-9543)
- Improved the loading time of the details pane on the Detections page, ensuring faster data display. (INVEST-9602)
- Fixed the incorrect display of the local authentication state for some users. (INVEST-9474)
- Fixed an issue where a role change for a user was not reflected at the top of the Account Settings menu in the upper-right corner. (INVEST-9873)
- Corrected the mapping of MITRE techniques to their respective tactics on the Security Overview page. (INVEST-9589)
- Removed extra spacing in the Alert Overview card on the Security Overview page. (INVEST-7746)
- Fixed an issue where the Related Detections tab showed all detections instead of filtering to only related detections. (INVEST-10010)

1.3 April 2024

1.3.1 Features and enhancements

- Added a view only user role for individuals who need to view detections and their associated data without making modifications to the system or taking action.
- Added the ability to *isolate entities* from network activity to address potential compromises or ongoing attacks based on Corelight evidence. This feature requires a properly configured *CrowdStrike integration*.
- Streamlined the license model to include only an advanced license and evaluation version.
- Added these new search-based alerts:
 - *JetBrains TeamCity Authentication Bypass CVE-2024-27198 and CVE-2024-27199*
 - *AsyncRAT remote access trojan*
 - *Fortra FileCatalyst remote code execution CVE-2024-251538*
 - *Onion domain*
 - *Fortigate RCE CVE-2024-21762*

1.4 March 2024

1.4.1 Features and enhancements

- Added an integration for CrowdStrike EDR (Endpoint Detection and Response) that provides additional context to detections that helps analyze threats and helps analysts make informed decisions during the triage process. When configured, detections show valuable entity context and the expanded data. For details, see *CrowdStrike data in detections*.
- Enhanced log search queries from detections with pre-populated queries to provide more contextually relevant results.
- Added a new ML detection: *C2 HTTP Frameworks*.
- Added these new search-based alerts:

- *Citrix Netscaler CVE-2023-4966 (CitrixBleed)*
- *ScreenConnect Authentication Bypass CVE-2024-1709*
- *Sonicwall RCE/DOS CVE-2022-22274 / CVE-2023-0656*

1.5 February 2024

1.5.1 Features and enhancements

- Added a menu to the *Dashboard* item in the left navigation to provide quick access to dashboards.
- Enhanced the Alert Category page so it preserves the search query and filters if you navigate from the page. (INVEST-8943)

1.5.2 Fixed bugs

- Fixed an issue where the Exfiltration via DNS alert displayed inconsistent severity scores. (INVEST-8935)
- Fixed an issue where you could not filter detections to exclude those with a score of 10 on the Detections page. (INVEST-8935)
- Fixed an issue where the Investigate Logs query did not work if the Suricata alert name changes. (INVEST-8921)
- Fixed an issue where not all Tor Connection detections were displayed if you navigate to the Detections page from the Alert Catalog details. (INVEST-8815)
- Improved security audit logging to eliminate extraneous entries, improve sorting, and track SSO users. (INVEST-8805)
- Fixed an issue where bulk updates for alert status in the Alert Catalog resulted in an invalid status on the back end and caused issues with alert filtering. (INVEST-9110)

1.6 January 2024

- Enhanced the *search feature* for the Alert Catalog to use AND/OR operators, wildcards, and regular expressions.
- Enabled view-only access for Analyst users to the Integrations in *System Settings*.
- Merged the **Related Detections** and **Related Entities** tabs on the detailed view of detections and added filters to easily find relevant alerts and entities.

1.7 December 2023

- Reorganized System Settings and created a dedicated page for integrations.

1.8 November 2023

1.8.1 Features and enhancements

- Added a table view to the Detections page that provides a more structured and concise representation of detections and lets you quickly scan rows and columns.
- Added the ability to *configure the autoclose time period* for detections in General Settings.
- Added the ability to use local authentication for all admin users and for analyst users not in the SSO domain when SSO is enabled.
- Added new ML model for *malicious certificates*.

1.8.2 Fixed bugs

- Fixed an issue where you could not open a link in a new tab if you signed in using SSO. (INVEST-8119)

1.9 October 2023

1.9.1 Features and enhancements

- Moved the System Settings to the left-side navigation. (Account Settings remain available in the menu that appears when you click your username.)
- Added support for search-based alerts, which are Corelight-defined log search queries. You can review and manage search-based alerts through the *Alert Catalog*.
- Repositioned and redesigned the sort and filter options on the Detections page.

1.10 September 2023

1.10.1 Features and enhancements

- Transitioned content from the Alerts page to the Security Overview and the Detections pages and removed the Alerts page.
- Added Highest Risk Detections section to the Security Overview page.
- Added the ability to *send a detection to ServiceNow* and create a security incident for further investigation.
- Added the ability to refresh the detections from the time window if Investigator identifies new detections.

1.11 August 2023

1.11.1 Features and enhancements

- Added the Detections page for better grouping, filtering, and visibility to alerts and the ability to triage detections and take action.
 - Updated sort terms for the Detections page to be more meaningful.
 - Moved the time window to be in the upper-left corner for consistency.
 - Added a header to the Detection page.
 - Removed redundant inline actions.
- Added GPT functionality to provide AI generated information to Suricata and machine learning detections.
- Reorganized system settings and created a general settings area that includes licensing and GPT.
- Added new ML detection: *domain combosquatting*.

Note: The Alerts page will be retired. The full functionality is available in the Security Overview and the Detections pages.

1.11.2 Fixed bugs

- Fixed an issue where authentication through SSO was incorrectly prompting for a one-time password. (INVEST-7379)

1.12 July 2023

1.12.1 Features and enhancements

- Enhanced exported alerts to include a URL that links to the associated detection details page within Investigator.

1.12.2 Fixed bugs

- Fixed an issue where the **Modified By** field would be empty for a changed severity score. (INVEST-7007)
- Fixed an issue with the CrowdStrike/IOC Overview dashboard where an error indicated a canceled query due to the regex backtrack limit. (INVEST-4085)

1.13 June 2023

1.13.1 Features and enhancements

- Updated the user interface for improved spacing, typography and readability, and accessibility.

1.13.2 Fixed bugs

- Fixed an issue that prevented SAML configuration with the error “Identify Provider configuration already exists”. (INVEST-6641)
- Fixed an issue for Notices and Suricata detections that didn’t show the normalized severity in the Alert Catalog. (INVEST-6727)

1.14 May 2023

1.14.1 Features and enhancements

- Added the ability to *customize the severity score* of an alert category.

1.15 April 2023

1.15.1 Features and enhancements

- Added the *Security Audit* log to provide a record of user and system activities.
- Upgraded the log search engine to *Falcon LogScale 1.76*.
- For IDN homograph machine learning alerts, added a more readable version of the domain in Unicode next to the Punycode value to demonstrate why the domain was flagged. (INVEST-5927)

1.15.2 Fixed bugs

- Fixed an issue to let search functionality recognize special characters in the Alert Catalog details page. (INVEST-6525)
- Fixed an issue where full search terms where not returning appropriate results in the Alert Catalog details page. (INVEST-5905)

1.16 March 2023

1.16.1 Features and enhancements

- Added the ability for admin users to reset 2FA by deactivating and reactivating users.

1.16.2 Fixed bugs

- Standardized the minimum normalized score for alerts to 1. (INVEST-5393)
- Fixed an issue where a low severity (benign) Machine Learning finding appeared in the Alert Catalog search results. (INVEST-5206)

1.17 February 2023

1.17.1 Features and enhancements

- Added support for SAML SSO user management.
- Added support for multiple alert exports and for exports through Elastic, CrowdStrike Falcon LogScale, and a generic HTTP exporter.
- Added the ability to select multiple entries in the Alert Catalog and change their status.

1.17.2 Fixed bugs

- Fixed an issue to allow deep pagination (results beyond 1000 pages) for the Alert Catalog. (INVEST-5892)
- Fixed an issue where machine learning alerts displayed an incorrect time interval. (INVEST-5890)
- Fixed an issue where some search results do not match the specified search term. (INVEST-5727)

1.18 January 2023

- Fixed an issue where alerts were not suppressed properly when an alert category had more than 10 excluded entities. (INVEST-5904)
- Fixed an issue where you could not configure Splunk Exporter due to permission issue. (INVEST-5858)
- Fixed an issue where the Domain Typosquatting machine learning model generated an alert for google.com. (INVEST-5523)

1.19 December 2022

- Migrated the existing user database to support enhanced user management. This change requires each user to reset their password and 2FA token. (INVEST-5725)
- Added more information for machine learning alerts including the associated connection's timestamp, source and destination IP, and domain. (The added information varies by the type of alert.) Added a **View ML Analysis** icon for easy access to details about how features influenced the model score. (INVEST-5078)
- Added a **Suppress Entity** button to the entity detail and alert detail pages that adds the entity to the Excluded Entities list. Excluded entities do not generate new alerts for the alert category. (INVEST-5505)

1.20 November 2022

1.20.1 Features and enhancements

- Added the ability to filter the alert categories and top entities on the Alerts dashboard based on a search term. (INVEST-4646)
- Suricata alert details include a **View Rule** button that shows the definition of the rule. (INVEST-4629)

1.20.2 Fixed bugs

- Fixed an issue where some top entities on the Alerts page did not show entity details when the pointer hovers over the entity. (INVEST-5460)
- Fixed an issue where you could not copy or download sensor export details. (INVEST-5486)
- Fixed an issue to properly sort search results in the Alert Catalog. (INVEST-5399)
- Fixed an issue where adding or deleting an excluded entity does not reflect in the list until a page refresh. (INVEST-5191)
- Corrected the entity type for the Social Engineering Domains and Domain Typosquatting models. (INVEST-5501)
- Fixed an issue where a previously configured sensor was not recognized in the interface and you were prompted to configure a new sensor. (INVEST-5096)
- Fixed an issue so the entity details page retrieves all alert types for the entity and not just a single alert type. (INVEST-5002)

1.21 October 2022

1.21.1 Features and enhancements

- The Alert Details page (available from the Alert Catalog) now shows the MITRE technique number as well as the technique description for relevant alerts. (INVEST-4471)

1.21.2 Fixed bugs

- Fixed an issue where alert details mistakenly showed zero analytics contributing to model score. (INVEST-5059)
- Fixed an issue where new machine learning alerts did not include machine learning feature data or the Top Model Advanced Analytics Summary data in the UI. (INVEST-5319)

1.22 September 2022

1.22.1 Features and enhancements

- Added the ability to change the status of alert categories and control if alert categories appear on the alert dashboard.
- Added the ability to *exclude alerts from specific entities* for an alert category in the Alert Catalog.

1.22.2 Fixed bugs

- Fixed an issue in the Users & Access page where selected users were not retained when navigating to a previous or next page in the list of users. (INVEST-4147)
- Fixed an issue where a low severity (benign) Machine Learning finding was generated as an alert. (INVEST-3889)

1.23 August 2022

1.23.1 Features and enhancements

- General Data Protection Regulation (GDPR) updates, including:
 - User-level consent to Corelight Privacy Policy during initial login.
 - Ability to manage *cookies* and tracking data during initial login and on the Account Settings page.
 - Investigator deletes raw log data at the end of the retention period. (Raw log data is kept for a minimum of 30 days and you can adjust the retention period with a license.)
- Product availability in the European Union (EU).
- The *Alert Catalog* provides a read-only list of all alerts in the system.

1.23.2 Fixed bugs

- Fixed an issue where the Investigate Logs button for a Suricata alert on the Entity page returns an error. (INVEST-4187)
- Fixed an issue where the number of alerts shown on the Alert or Entity page was one number higher than the number of alerts shown when you click the Investigator Logs button. (The logs would drop the last event.) (INVEST-4213)
- Fixed an issue where some log data was deleted before reaching the log retention limit due to a data retention storage setting. (INVEST-4444)

1.23.3 Known issues

- Investigator displays Falcon LogScale (Humio) content in an iFrame, which can result in display issues with the content. For example, data can extend beyond the visible frame of the Humio logs at smaller screen sizes. (INVEST-4083)
- In the Falcon LogScale (Humio) frame of the Log Search page, the **Save As | Export to File** option returns only a blank page for large downloads (downloads taking more than 15 minutes). Smaller downloads are not impacted. (INVEST-4354)

QUICKSTART

These topics help you get started with Corelight Investigator.

2.1 Before you start

Make sure that all the Corelight Sensors that you want to connect to Investigator meet these requirements and are configured to export to Investigator.

- You must use hardware, virtual (VMware or Hyper-V), or cloud (EC2 or Azure) sensors.
- Your sensors must be running software release v25 or later.
- Your sensors must have Corelight Cloud Services enabled.
 - Run `corelight-client configuration update --remote.enable 1`
 - *–or–*
 - Go to **Sensor | Updates | Enable Corelight Cloud Services**.
- Your sensors must have external network connectivity (north/south).
- We recommend Chrome for the best web experience.

2.2 Log in for the first time

Once Corelight creates your Investigator account, the system sends you a confirmation email welcoming you to the platform. You also receive an account registration email containing your username and a temporary password. Use these credentials to complete your account registration and log in for the first time.

To register your account

1. Click the **Go to Investigator** link in the account registration email to activate your account.

A welcome screen appears.
2. Click **Start**.
3. Enter the username and password provided in the account registration email and click **Register now**.
4. Review the Corelight privacy policy and click **Continue**.

When you click **Continue**, you acknowledge the Corelight privacy policy and agree to the processing of data in accordance with the policy.
5. Create a secure password that meets the requirements. Confirm your new password and click **Next**.

6. Open an authenticator app (such as Google Authenticator) on your phone and add a new account.
7. Scan the provided QR code to connect your app to Investigator.
8. Enter the one-time password (OTP) from your app in Investigator.
9. Click **Verify OTP** to log into Investigator.
10. Click **Login**.

Your account is set up and active. Enter your login details again to get started.

Next, enable and configure Corelight Investigator on your sensor.

2.3 Configure your sensors

If you are an admin user and your Investigator account is not connected to a sensor, you are prompted to add sensors and import logs after your initial login. Your Investigator license determines what logs are imported. For details, see [Find details in the logs](#).

Tip: You can also perform these steps from the Sensor Monitoring page in System Settings.

To configure your sensors

1. Click **Configure Sensor**.
2. Enter your Investigator credentials and click **Confirm**.
3. Confirm your OTP code and click **Verify**.
4. Click **View Sensor Information**.
5. Note the Access Key, Secret Key, Region, and Stream Name specified by Investigator.
You can use the copy button to easily get the values and you can download the values as a CSV file.
6. Outside of Investigator, enable export to Corelight Investigator for your sensors, either through the sensor interface or through Fleet Manager.

Web interface

- a. From your sensor or policy configuration, go to **Export** and turn on **Export to Investigator**.
- b. Enter the Access Key, Secret Key, Region, and Stream Name specified by Investigator.
- c. Optionally, select Zeek logs to exclude and use the Corelight filter language to define a log filter.
- d. Click **Apply Changes**.
- e. Repeat this step for all sensors that you want to connect to Investigator.

See [Investigator export](#) in the *Corelight Sensor User Guide* for more details.

corelight-client

```
corelight-client configuration update --bro.export.investigator.enable 1 \
--bro.export.investigator.access_key=<access_key> \
--bro.export.investigator.secret_access_key=<secret_key> \
--bro.export.investigator.region=<region> \
--bro.export.investigator.stream=<stream>
```

7. In Investigator, click **Close**.

The Sensor Monitoring dashboard shows the imported activity.

Note: If needed, update your firewall to allow external HTTPS traffic to the regional Kinesis service endpoint from the sensor. For regional endpoint details, see the AWS help topic [Amazon Kinesis Data Streams endpoints and quotas - AWS General Reference](#).

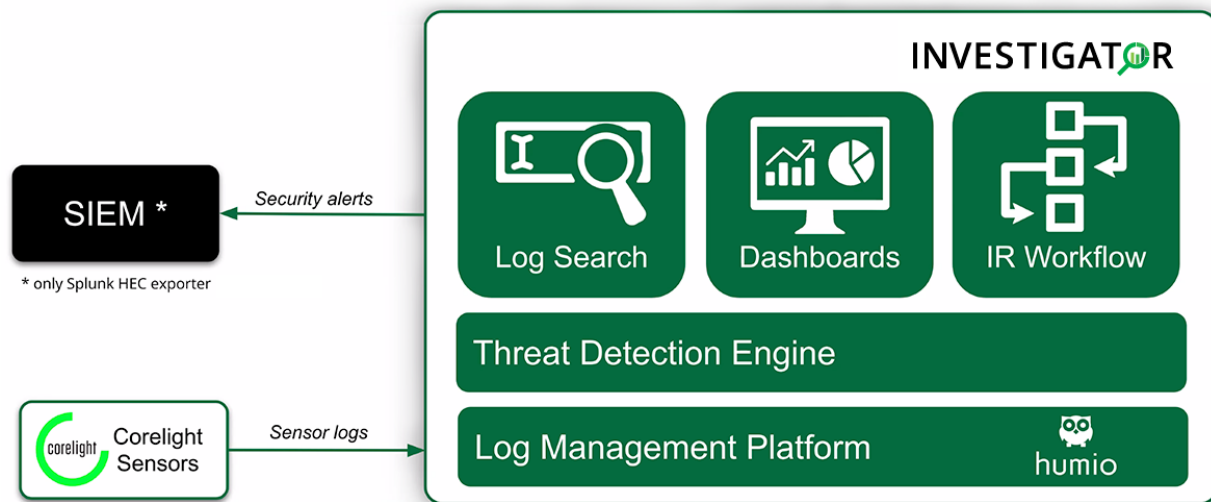
If you experience problems logging in or connecting your sensors, read through the topics in this help system. If you can't resolve your issue, contact Corelight Support.

INVESTIGATOR OVERVIEW

Corelight Investigator is a cloud-based security platform that provides threat detection and investigation, incident response, and log storage with scalable search. The platform combines Corelight’s rich logging framework with Investigator’s sophisticated analytics engine and gives security responders a clear path to identifying and understanding potentially malicious traffic.

Corelight Investigator is a software-as-a-service (SaaS) platform so you don’t need to manage installations, upgrades, or storage systems. Investigator is accessible directly from a browser and receives automatic updates.

Connect one or more Corelight Sensors to Investigator and get immediate access to its powerful analytics engine. See the [Quickstart](#) for details.



Watch a Corelight YouTube video with an overview of Investigator.

3.1 Alerts

As logs are exported from the sensor, Investigator uses a variety of detection mechanisms to raise specific categories of security alerts.

- **Notice** - Alerts related to notice logs generated by the sensor, Corelight collections, or any additional custom scripts.
- **Suricata** - Suricata alerts raised as a result of your configured Suricata ruleset on a sensor.
- **Machine Learning** - Alerts raised by the Investigator analytics engine after Corelight logs are run through machine learning models.
- **Search Based** - Alerts generated by Corelight-defined log search queries. The Corelight Labs team creates these alerts using LogScale searches. This type of alert provides expanded coverage and rapid response.

Each alert is associated with an **entity** – the host domain or IP address – that Investigator identifies as associated with a potential threat. Alerts are assigned a score based on the likelihood that the alert represents a real security threat. The score ranges from 1 to 10, with higher scores indicating more severe threats. Investigator *normalizes scores* across alert types. You can customize the severity score in the *Alert Catalog*.

This information lets you quickly identify, prioritize, and respond to security incidents on your network.

There can be anywhere from 10 to 10,000 alerts within a week for an alert category. The volume of alerts can be overwhelming and often analysts are unsure of what action is required. That's where detections can help.

3.2 Detections

A detection is an aggregation of alerts from the same alert category and entity combination. The detection can include Suricata, Notices, and ML alerts or a combination of these alerts.

Through the *Detections* page, Investigator helps you manage the alert triage process with simplified workflows and instructions to process and resolve the alerts. Investigator leverages Zeek data along with the transparency of machine learning models to provide context and helpful information to take the next steps.

3.3 AI-generated content

Investigator uses GPT by OpenAI to generate descriptions, next steps, and additional details for some alerts. For example, for Suricata rules, Investigator sends the rule definition to GPT to create a meaningful explanation of the alert and provide recommendations for next steps.

GPT generated content and features are identified by this icon:



Note: Content generated through AI might have errors or omissions. Use your best judgment for this content.

AI-generated content only applies to Corelight-provided data, rules, and alerts and is not available for unknown or customer-generated data.

The GPT implementation is based on the GPT 3.5-Turbo model and uses Corelight authentication.

You can turn off *GPT integration* in General Settings.

3.4 Analytics engine

The analytics engine is a core component of the Investigator platform. In addition to alerts raised by packages on the sensor, the analytics engine provides an additional layer of insight for security responders.

The analytics engine steps through four phases:

1. Normalization
2. Feature computation
3. Model scoring
4. Alert aggregation

During **normalization**, log data is processed and transformed so that it can be properly consumed by the algorithms.

Investigator uses the normalized data for **feature computation**. In this step, the analytics engine calculates values for characteristics that it deems important in determining whether the log is associated with an attack.

Next, the analytics engine uses the computed feature values for **model scoring**. Algorithms use the features to assign a numeric score to each enabled model. A higher score indicates a higher likelihood that the data represents the attack that the model specifies.

Finally, the model scoring output and alerts produced by Corelight Sensors are sent to **alert aggregation**. In this step, the analytics engine detects which entities represent threats and sends those threats to Investigator's detection dashboard.

3.5 Roles

Investigator provides three types of user roles: analyst, administrator, and viewer.

The analyst role has access to all the entity and alert data and can follow and identify threats. This is the role for a security analyst and threat hunter.

The admin role includes all the data visibility of the analyst, and can also configure system settings to perform tasks such as manage users, connect sensors, set up integrations, and export alerts.

The viewer role is for individuals who need to view detections and their associated data without making modifications to the system or taking action.

Note: Sensors have their own user roles that are managed separately. You will need an admin role on your sensors to configure them for Investigator exports.

SECURITY OVERVIEW

The Security Overview page appears when you log in. This page provides a high-level summary of what's going on in your network based on the data coming into Investigator from your sensors.

The Security Overview displays summary data about detections on your network, including the number of detections over time and the distribution of detections across MITRE ATT&CK categories as well as information about individual detections.

The screenshot displays the Security Overview dashboard with the following components:

- Highest Risk Detections:** A table with columns for Score, Entity, and Detections. The top entries are:

Score	Entity	Detections
8	10.2.128.198	1
5	fulnqcnquwacif.com	1
5	mahronis.com	1
5	jdhpqgajxxqwq.com	1
5	hoozffgordgv.com	1
5	edsexpressinc.com	1
5	btmqnqozo.com	1
5	psufqokase.com	1
5	ybuaupf.com	1
5	ryeavxsthkysv.com	1
- Entities With Detections:** Shows 5,068 entities with a 6315% increase.
- Alert Categories With Detections:** Shows 313 alert categories with a 502% increase.
- MITRE ATT&CK Map:** A heatmap showing the distribution of detections across various MITRE ATT&CK categories. Key categories include:
 - Active Scanning: 318
 - Phishing: 15
 - User Execution: 6
 - Masquerading: 6
 - Modify Authentication: 2
 - Network Services: 204
 - Remote Services: 9
 - Application Layer: 136
 - Exfiltration Over C2: 5
 - Resource Hijacking: 2
 - Exploit Public: 12
 - Scheduled Task/Job: 4
 - Exploitation for: 2
 - OS Credential: 2
 - System Services: 4
 - Windows Management: 4
 - Remote Service: 5
 - Exploitation of: 4
 - Lateral Tool: 4
 - Dynamic Resolution: 5
 - Automated Exfiltration: 1
 - Data Transfer: 1
 - Exfiltration Over: 1
 - Account Discovery: 17
 - File and Directory: 17
 - Network Share: 17
 - Non-Standard: 17
 - Permission Groups: 17
 - Remote System: 17
 - Encrypted Channel: 5
 - Fallback Channels: 5
 - Protocol Tunneling: 4
 - Data Encoding: 1
 - Non-Standard: 1
 - Proxy: 1
- Corelight Blogs:** A list of recent blog posts, including "Detections and Findings using Corelight in the Black Hat Asia NOC | Corelight" and "Detecting Storm-0558 Using Corelight Evidence | Corelight".

The Security Overview organizes information in these cards:

- **Highest Risk Detections** - highlights important alerts and alert categories. You can switch between these tabs:
 - **Entities** - Displays the entities with the highest severity scores for open detections. The score ranges from 1 to 10 with higher scores indicating more severe threats.

Mouse over an entity to see more information about its associated alert categories and click **View Detections** to see the detections for the entity in the Detection summary page.

Entities with high threat scores are a good place to start your analysis.

– **Alert Categories** - Displays the alert categories with open detections and their severity score.

- **Entities with Detections** - Displays the number of entities (IP addresses and domain names) that have been flagged with a security alert. The card also shows the percentage change from the previous time window. The number of entities includes both open and closed detections.
- **Alert Categories with Detections** - Displays the number of distinct detection types (both open and closed) that have been identified in the specified time interval. This section also shows the percentage change from the previous time window.
- **MITRE ATT&CK Map** - Provides a heat map that aligns security alerts with MITRE ATT&CK tactics, techniques, and procedures. The MITRE ATT&CK framework shows how attacks evolve through an enterprise.

This card shows the observed tactics (such as reconnaissance and initial access execution) and techniques in the color-coded blocks.

For active scanning alert categories, a number in the bottom-right corner of each card shows how many unique entities were found for the category. Mouse over the color coded technique blocks for details about the associated detections.

Tip: For more information about the MITRE ATT&CK framework, see <https://attack.mitre.org>.

- **Blogs** - Populated with the latest security blogs from Corelight on corelight.com.

You can focus the contents of the Security Overview page by specifying a time interval from the menu in the upper-right corner. The default time interval is 7 days. You can change the time period to range from one hour to three months, or you can specify a custom date range. The time range applies to all parts of Investigator and changing the time interval in one place changes it for all features with a time interval. For example, the Detections page uses the same time interval as the Security Overview page.

The Investigator pages do not automatically refresh. If new detections are available for your time window, the **Refresh** icon to the left of your time interval selection is blue. Click the **Refresh** icon to show the new detections in the results. If new detections are not available, the **Refresh** icon is gray. The time window shows the length of time since the last update.

To learn more about these entities and alerts, go to the [Detections](#) dashboard.

WORK WITH ALERTS & DETECTIONS

5.1 Identify and manage threats through the Detections page

Investigator helps you monitor and manage alerts by grouping them as detections and providing simplified workflows and instructions to process and resolve them.

5.1.1 What are detections?

A detection is an aggregation of alerts from the same alert category and entity combination. Investigator creates detections when an entity generates an alert for an alert category. Future alerts from the same alert category and entity pair are appended to the same detection.

Detections are accessible to Admin, Analyst, and Viewer users (although Viewers cannot see the action buttons associated with detections).

For more information, see *Overview of alerts and detections*.

5.1.2 How to investigate a detection

Here are the basic steps to investigate a detection:

1. From the *Detections page*, review the list of discovered detections.

The list summarizes basic information about each detection. You can drill down to learn more about detections and their associated entities and alerts.

2. *Sort and filter* detections.

You can sort detections by severity score and time of detection.

You can filter detections by severity score range, status, alert category, entity, or assignee.

3. View essential information about a detection in the *Quick View* panel.

4. View complete information in the *Detailed View*.

The Detailed View includes information provided by the Corelight content team and can include description, significance, validation, next steps, and associated MITRE ATT&CK techniques.

Some detections provide AI generated descriptions and let you Ask GPT for more information through pre-formed prompts.

The Detailed View also shows alerts for the detection, related detections, and related entities.

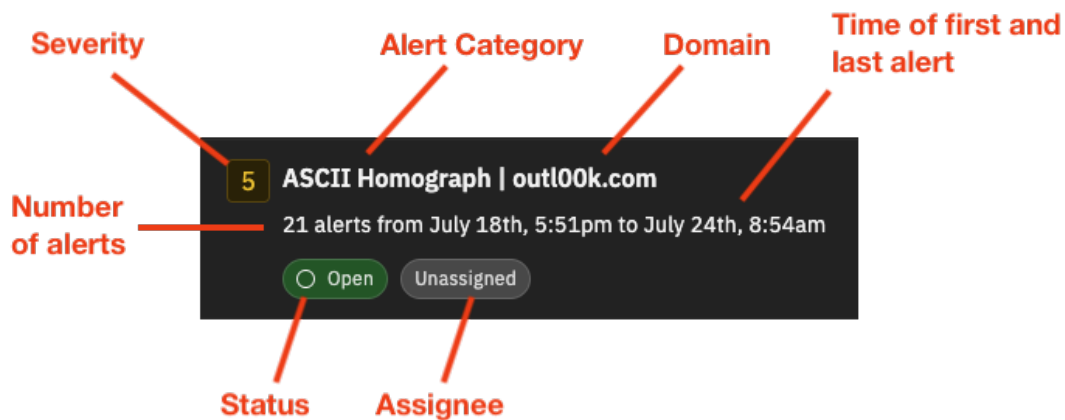
5. *Assign* a detection to any user in the system for further analysis.

6. *Exclude* an entity from a detection.
7. *Send to ServiceNow* as a security incident.
8. *Close* a detection.

5.1.3 Detections page and summary view

The Detections page shows a summary of each detection. Each detection includes these details:

- Severity – a number ranging from 1 to 10; more severe threats have a higher score
- Alert Category – the name of the security alert
- Entity – IP address or domain name
 - With a configured CrowdStrike integration, the detections show *additional host information* from CrowdStrike data and can let you isolate entities from all network activities in response to security breaches.
- Number of Alerts in the detection
- First Alert Time – time of the first alert in the detection
- Last Alert Time – time of the most recent occurrence of the alert in the detection until the detection is closed
- Status – open or closed
- Assignee – indicates a user has been assigned to the detection or the detection is unassigned.



By default, detections appear in the list view. If you prefer a more structured and concise representation of the detections, you can switch to the table view. The table view lets you quickly scan rows and columns to find relevant data. You can switch between views using the **List** and **Table** buttons under the filter options.

Detections in the list view

5 NXDOMAIN Beaconsing | fljxvuyd.com
1 alert from October 23rd, 11:24pm to October 23rd, 11:24pm
Closed Unassigned

5 NXDOMAIN Beaconsing | gkmvxsidfvau.com
2 alerts from October 19th, 6:21pm to October 23rd, 11:21pm
Closed Unassigned

5 NXDOMAIN Beaconsing | wjjyzhwwbnrl.com
1 alert from October 23rd, 11:18pm to October 23rd, 11:18pm
Closed Unassigned

Detections in the table view

Score	Alert Category	Entity	Alert Type	#Alerts	Status	Assignee	Latest Start Timestamp ↓
5	NXDOMAIN Beaconsing	fljxvuyd.com	Machine Learning	1	Closed	Unassigned	October 23rd, 11:24pm
5	NXDOMAIN Beaconsing	gkmvxsidfvau.com	Machine Learning	2	Closed	Unassigned	October 23rd, 11:21pm
5	NXDOMAIN Beaconsing	wjjyzhwwbnrl.com	Machine Learning	1	Closed	Unassigned	October 23rd, 11:18pm

The *sort and filter options* along the top of the results help you focus the list of detections and the *Quick View* panel provides more information about the detection.

The screenshot displays the 'Detections' interface. At the top, there's a 'Filters' section with a 'Severity Score' slider (set to 5), 'Status' (Open/Closed), and search boxes for 'Category', 'Entity', and 'Assignee'. Below this is a table view showing a list of detections. Each row includes a severity score (5), a title (e.g., 'NXDOMAIN Beaconing | rduozrvtwvt.com'), a time range, and status buttons ('Open', 'Unassigned'). A 'View Detection' button is visible on the right. The detailed view on the right shows the 'Description', 'Significance', and 'Detection Summary' for a selected detection.

Note: The Detections page does not automatically refresh. If new detections are available for your time window, the **Refresh** icon in the time window is blue. Click the **Refresh** icon to show the new detections in the results. If new detections are not available, the **Refresh** icon is gray. The time window shows the length of time since the last update.

5.1.4 Quick view of detection information

Each detection displays a Quick View panel to the right. This panel provides essential information and actions for the detection, including

- Description of the detection. Machine learning alerts and notices from Corelight collections include a description. Suricata alerts include an AI generated description. Custom generated alerts will not have a description.

Note: *AI-generated content* is identified by an icon that appears next to the description. Although our goal is a detailed and accurate description, use your judgment for any AI generated description.

- Significance – the potential impact of this detection for machine learning (ML) detections and some notices.
- Detection Summary
 - Status – if the detection is open or closed.
 - Assignee – indicates if the detection is assigned to a user or if it is unassigned.
 - Number of alerts
 - First alert time and last alert time
- Entity for the detection, including IP address or domain.

With a *CrowdStrike integration configured and enabled*, the detections show *additional host information* from CrowdStrike data to provide even more context and let you isolate entities from network activity in response to security events.
- Alert Category
 - Alert Category – click the name to see to the full entry in the Alert Catalog and customize the severity score
 - Severity
 - Type – Notice, Suricata, Search Based, or Machine Learning
- MITRE ATT&CK techniques (if available) – Notice and ML detections that Investigator can map to the MITRE ATT&CK framework include links to the relevant MITRE ATT&CK techniques.

Buttons at the top of the Quick View let you perform any of the following actions for a detection:

- *Close the detection*
- *Send to ServiceNow as a security incident*
- *Suppress Entity from a detection*
- *Assign a detection to a user*

Tip: In table view, the action icons appear when you hover over a table row and also in the Quick View panel.

5.1.5 Detailed view of the detection information

By default, the Detection page displays a Quick View panel with more information for the selected detection. Click the **View Detection** button in the Quick View panel to display a more complete view of the detection in a full page view.

The screenshot displays the 'Detailed View' of a detection. At the top, there's a navigation bar with a 'Back to Detections Page' link and the detection title 'Intel::Notice | 221.226.212.189'. Below the title are action buttons: 'Close Detection', 'Suppress Entity', 'Assign To', and 'Investigate Logs'. A breadcrumb trail shows 'Detection Details', '1 Alert', and '278 Related Detections'. The main content is divided into two columns. The left column contains sections for 'Description', 'Significance', 'Validation', 'Next Steps', and 'Ask GPT'. The right column contains sections for 'Detection Summary', 'Entity', 'Alert Category', and 'MITRE ATT&CK Techniques'. The 'Ask GPT' section includes two prompts: 'What type of attacks might be associated with this alert?' and 'What techniques specific to this alert might an adversary use?'. The 'Detection Summary' shows the status as 'Open', the assignee as 'Unassigned', and the number of alerts as 1. The 'Entity' section shows the IP address 221.226.212.189. The 'Alert Category' section shows the category as 'Intel::Notice' and the severity as 4. The 'MITRE ATT&CK Techniques' section lists 't1595 Active Scanning' and 't1071 Application Layer Protocol'.

The Detailed View includes all the information and actions available in the *Quick View*, plus this additional information (as applicable for each alert type):

- Description – explains the detection
- Significance – the potential impact of this detection for machine learning (ML) detections and some notices.
- Validation – how to assess the correctness of the detection
- Next Steps – recommendations on how to address the detection (for ML detections and some notices)
- Suricata detection details include a **Suricata Rule** section that shows the definition of the rule.
- Ask GPT – Corelight-provided alerts and detections include an Ask GPT section that lets you query GPT from OpenAI about the alert or detection through pre-formed chat prompts.

Click a prompt to ask GPT for additional details about an alert generated by a detection. Prompts might let you ask “what does this alert mean” or “what next steps should I take.” As you get an answer, related queries appear as appropriate and available.

If preferred, you can turn off *GPT integration* in General Settings.

You can use the **Request More** button to send feedback to the Investigator team and suggest prompts related to the detection. (Note that this feedback is not interactive.)

Buttons along the top let you perform these actions:

- *Close the detection*

- *Send to ServiceNow as a security incident*
- *Suppress Entity from a detection*
- *Assign a detection to a user*
- Investigate logs - pivots to the logs to review the evidence exported from your sensors. Investigator pre-populates the log search query to provide contextually relevant results.

You can click the **Copy Detection URL** icon in the upper-right corner to copy the Detailed View page location and share it as a unique URL for the detection.

Tabs under the action buttons provide access to:

- Alerts – shows the alerts that contributed to this detection and details about the alerts. *Machine learning detections* show details about the model score, including analytics that contributed to the calculation of that score.

You can pivot to the logs and view in LogScale by clicking the **Investigate Logs** icon.

- Related Detections – provides additional context and lets analysts review detections related to the entity and the alert category to see a more complete view. The tab shows detections (both open and closed) from other related alert categories and other entities. By default, Investigator shows the newest detections first, but you can also sort by severity and change sort options. Click a detection to display additional information.

You can use the filters to limit the results to a specific alert category or entity.

To filter by alert category, type an alert category in the **Category** search field. Investigator suggests matches as you type. You can also click the checkbox to show detections only for the current category. You can add multiple alert categories to the filters.

To filter by an entity, type an IP address or domain in the **Entity** search field. Investigator suggests matches as you type. You can also click the checkbox to show detections only for the current entity. You can add multiple entities to the filters.

You can switch the results list format and view in either table view or list view, and you can sort by severity score.

A number before the tab name indicates the quantity of current related alerts/detections.

5.1.6 Sort and filter detections

By default, detections appear in a list that shows the most recent activity first and the detections appear for the last 7 days. You can change the time period to range from one hour to three months. You can also specify a custom date range.

Note: When you change the time window for detections, the time setting applies to all parts of Investigator that use a time window.

You can sort detections based on:

- Newest Detection – show the most recent detections first
- Oldest Detection – show the oldest detections first for the specified time window
- Highest Severity – show the most severe detections first
- Lowest Severity – show the least severe detections first

You can filter detections based on:

- Severity Score – use the slider to specify a minimum and maximum score for the results. By default, all scores are included.
- Status – open or closed.
 - Open – the alerts/detection have been generated and are awaiting investigation and resolution.
 - Closed – the alerts/detections have been fully investigated and resolved, and no further action is required or the detection was open for more than a week.
- Category – Limit the results based on a single alert category. This search box suggests matches based on values found in the filtered results; start typing keywords to see available filtering options. You can only filter by one alert category at a time.
- Entity – Find detections by IP or domain. This search box suggests matches based on values found in the filtered results; start typing search terms to see available filtering options. You can only filter by one entity at a time.
- Assignee – A search box lets you find and select an assignee from the full list of users associated with the account and show only detections assigned to that user. You can also show only unassigned detections or detections assigned to you.

You can combine multiple filters to focus results. Active filters are shown at the top of the Filters pane.

You can reset custom filters. Click **Reset** to clear all filters, or click the **X** next to a specific filter label to remove it.

You can show or hide filters with the icon in the upper-right corner of the filter pane.

5.1.7 Exclude an entity

Analyst and admins can exclude entities from a detection. Excluded entities do not generate new alerts for the alert category. Typically, you exclude trusted entities so you can focus on other entities.

To exclude an entity from an alert category

1. From the Quick View panel or the Detailed View, click the **Suppress Entity** button.
A dialog box prompts for confirmation.
2. Click **Suppress Entity**.

If excluded, you can click the **Unsuppress** button to restart alerts for the entity.

You can see a list of suppressed entities for an alert category in the *Alert Catalog*.

5.1.8 Assign a user to a detection

You can assign a detection to any admin or analyst user, including yourself.

To assign a detection

1. From the Quick View panel or the Detailed View, click the **Assign To** button.
2. Search for an assignee and select their alias.
Each assignee is listed by their account alias.
3. Click **Apply**.

Once assigned, you can change the assignee or revert to **Unassigned**. Users do not get notified when assigned a detection, but they can sort based on their assigned detections.

5.1.9 Send a detection to ServiceNow

If an analyst or admin identifies a detection as a potential incident, they can send the detection to their ServiceNow instance. A detection sent to ServiceNow creates a security incident for further investigation and can initiate response workflows.

Important: Users send detections manually on a case-by-case basis.

To enable this functionality, an admin for your account needs to provide access settings for your ServiceNow instance in the *integration settings*.

To send a detection to ServiceNow

1. From the Quick View panel or the Detailed View, click the **Send to** button.
2. You are prompted to confirm the action.

When sent, the status for the detection changes to closed, the username of the person who sent to detection is added to the details, and an icon indicates the detection has been sent to ServiceNow.

Important: Once you send a Detection to ServiceNow, it cannot be reverted or reopened.

When you send a detection to ServiceNow, Investigator includes these fields:

- Description of the detection
- Alert category name
- Severity score
- Detection status (will always be closed since Investigator closes detections once sent to ServiceNow)
- Entity
- Entity type
- Assignee
- Number of alerts
- Unique URL for the detection
- Detection created time
- Last updated time

Tip: Corelight video on YouTube: [How Corelight's ServiceNow integration speeds response](#)

5.1.10 Close a detection

Without any user activity, Investigator automatically closes the detection one week after creation. (Admins can *configure the autoclose time period*.) Additionally, analysts and admins can close detections when they determine it is not a security issue or when they have addressed the issue.

If your active filters only show Open detections (the default), the closed detection no longer appears in the list. Once you close a detection, you cannot re-open it.

Note: The system automatically closes detections 7 days after creating the detection.

To close a detection

1. From the Quick View panel or the Detailed View, click the **Close Detection** button.
2. When prompted to confirm the action, click the **Close Detection** button.

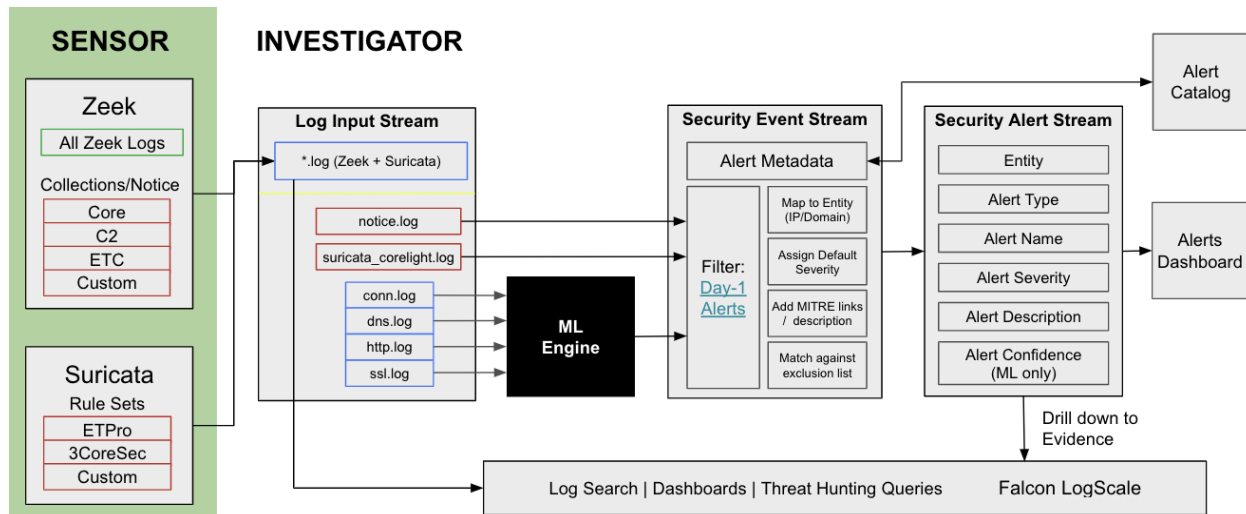
A message appears when the action completes successfully.

Note: If the detection is closed and alerts with the same attributes, Investigator creates a new detection for the same alert category and entity pair.

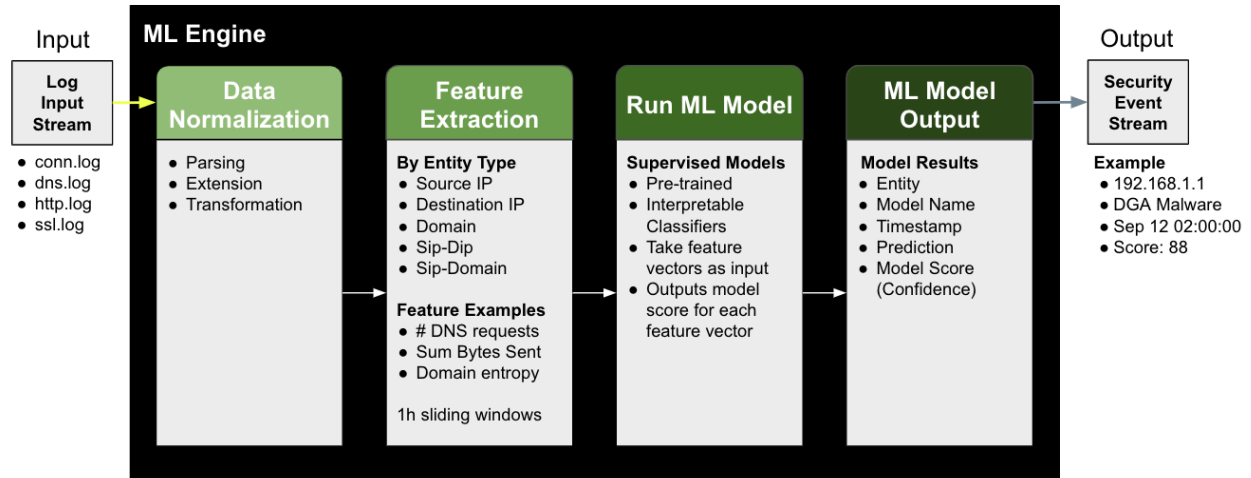
You can still perform actions on closed detections, such as assign a user or suppress an entity.

5.2 Machine learning detections

The Investigator analytics engine provides several machine learning models that analyze Corelight logs and generate alerts when they detect invariant patterns and potential threats.



The machine learning engine in Investigator ingests the log input, normalizes and organizes the data, extracts and analyzes the features, and applies the machine learning models. The result is a focused security event with a severity score.



Machine learning alerts complement notices, Suricata alerts, and Falcon LogScale queries.

5.2.1 Machine learning models

This section describes the machine learning model detections available in Investigator.

You can find more information for each machine learning model in Investigator by clicking the alert name in the Alert dashboard. The alert details provide a summary and suggest next steps for each detection to guide your investigation and troubleshooting.

5.2.1.1 ASCII homograph

The adversary registers a domain containing one or more ASCII homographs, making it visually similar to a trusted domain. This is possible because the ASCII table contains characters that look similar, for example the uppercase character “O” and the number “0”. Homograph domains are often used by adversaries to trick Internet users into visiting phishing sites.

- Detections are based on information from the HTTP log.
- Default severity: 5
- This model is silent by default and can be tuned per account.

5.2.1.2 Attempted connection to a DGA domain

One or more hosts attempted to connect to the domain. The domains generated by Domain Generation Algorithms (DGAs) exhibit random-looking patterns of letters, numbers, or words, and are generally long to minimize collisions with existing domains.

- Detections are based on information from the DNS log.
- Default severity: 5
- Enabled and set to alert by default.

5.2.1.3 DGA malware

A Domain Generation Algorithm (DGA) is a technique used by cyber attackers to generate new domain names and IP addresses for malware command and control servers.

Malware often relies on DGAs to obtain C2 (Command and Control) rendez-vous locations. DGAs result in an elevated number of DNS requests performed by the same internal IP, where the requests attempt and mostly fail to resolve random domains (NXDOMAIN response code).

This attack is considered stealthy because organizations often allow web traffic by default and new AGDs (algorithmically generated domains) bypass blacklists.

- Detections are based on information from the DNS log.
- Default severity: 8
- Enabled and set to alert by default.

5.2.1.4 Discovery via network service scanning

Adversaries perform port scans to discover the ports and services available in a network, and which could be used to establish connections with the targeted machines. Port scans result in an elevated number of connection attempts performed by the same source IP, targeting a wide range of destination ports in one or more destination IPs.

- Detections are based on information from the conn log.
- Default severity: 6
- Enabled and set to alert by default.

5.2.1.5 DNS reconnaissance

A DNS reconnaissance attack tries to get information about the network infrastructure of the company, in particular, about the DNS servers and their records. The two main DNS reconnaissance techniques are zone transfer attacks and brute force subdomain enumeration.

This attack is considered stealthy because most organizations do not monitor DNS traffic.

- Detections are based on information from the DNS log.
- Default severity: 3
- Enabled and set to alert by default.

5.2.1.6 Domain combosquatting

This detection involves the identification of domain names that are similar to legitimate domain names by combining multiple words or phrases.

For example, a cybercriminal might create a domain name like “facebooksocialnetwork.com” to trick users into thinking it is the official Facebook website. This technique is effective because it can be difficult for users to differentiate between legitimate domain names and similar, but fake ones.

To detect domain combosquatting, machine learning algorithms are trained on large datasets of domain names to learn patterns of legitimate and abusive behavior. The algorithms can then identify domain names that are likely to be used for combosquatting based on features such as the similarity of the name to a legitimate domain, the use of common keywords, and the registration date of the domain. By identifying

these malicious domain names, machine learning can help prevent users from falling victim to phishing attacks and other forms of online fraud.

- Detections are based on information from the HTTP log.
- Default severity: 5
- This model is silent by default and can be tuned per account.

5.2.1.7 Domain typosquatting

Adversaries rely on errors made by users when typing a website address to deliver malware, to redirect to a malicious site, to commit fraud, or to phish credentials.

- Detections are based on information from the HTTP log.
- Default severity: 5
- This model is silent by default and can be tuned per customer.

5.2.1.8 Exfiltration via DNS

Adversaries can exfiltrate data by encoding data in the subdomains of DNS queries. This technique results in a high volume of DNS queries generated by the same source IP, where the queries aim at resolving different random (and often long) subdomains of a domain owned by an adversary. This activity can be difficult to detect since continuous monitoring is rarely applied to DNS traffic.

- Detections are based on information from the DNS log.
- Default severity: 9
- This model is silent by default and can be tuned per account.

5.2.1.9 C2 HTTP Frameworks

This detection identifies Command and Control (C2) frameworks that use the HTTP (Hypertext Transfer Protocol) for communication between compromised systems (infected devices or computers) and a central command server controlled by an attacker.

This detection considers multiple aspects of the HTTP connections between the pair such as beaconing activity, payload, and header characteristics.

C2 frameworks might differ in implementation but can share similar characteristics such as connection duration and frequency, URI pattern, and payload consistency.

- Detections are based on information from the HTTP log.
- Default severity: 9
- Enabled and set to alert by default.

5.2.1.10 IDN homograph

The adversary registers a domain containing one or more homoglyphs, making it visually similar to a trusted domain. This is possible because different international alphabets contain letters that look the same and are coded as different Unicode characters, for example the Latin character “a” and the Cyrillic character “а”. Homograph domains are often used by adversaries to trick Internet users into visiting phishing sites.

- Detections are based on information from the HTTP log.
- Default severity: 5
- This model is silent by default and can be tuned per account.

5.2.1.11 Malicious file download

An adversary might rely on a user opening a malicious file to gain execution. Users may be subjected to social engineering to get them to open a file that will lead to code execution. Adversaries use several types of files that require a user to execute them, including .doc, .pdf, .xls, .rtf, .scr, .exe, .lnk, .pif, and .cpl.

- Detections are based on information from the HTTP log.
- Default severity: 7
- This model is silent by default and can be tuned per account.

5.2.1.12 Malicious SSL certificate

Malicious SSL certificates are digital certificates that have been obtained or issued fraudulently or through an unauthorized channel. Attackers can use these certificates to conduct malicious activities, including man-in-the-middle (MITM) attacks, phishing, and malware distribution.

This detection indicates one or more hosts attempted to connect to the domain with a potentially malicious SSL certificate. The certificate associated with the domain shares characteristics with certificates used by malware and adversaries. These certificates often contain random, unpopular domain names in the certificate subject common name (CN), use free Certificate Authorities, and do not populate optional subject and issuer information such as the CountryName (C), Locality (L), Organization (O), OrganizationalUnit (OU), or the StateOrProvinceName (ST).

- Detections are based on information in the x509 log.
- Default severity: 5
- Enabled and set to alert by default.

5.2.1.13 NXDOMAIN beaconing

Attackers and malware can rely on hard-coded domains for C2 and often these domains are unavailable. Malware and adversaries repeatedly attempt to connect to the predefined domains (beaconing) until a successful connection is established or a limit of attempts is reached. When the same IP tries to connect periodically to a domain that does not resolve, the result is an NXDOMAIN response code.

- Detections are based on information from the DNS log.
- Default severity: 4
- This model is silent by default and can be tuned per account.

5.2.1.14 Social engineering domains

Social engineering domains trick internet users into visiting malicious sites on the false pretext of prizes, free software updates, fake antivirus alerts, and such. These sites attempt to phish credentials, install malware, or redirect users to other malicious sites.

- Detections are based on information from the HTTP log.
- Default severity: 5
- Enabled and set to alert by default.

5.2.1.15 Tor connections

Adversaries, malware, and insider threats use Tor to bypass blacklist-based detection and prevention. Tor can be detected by analyzing the subject names of SSL requests. Tor connections often populate the subject name with random-looking domains that do not actually exist, where the domains are preceded by www., have a .com or .net top-level domain, range in size from 8 to 20 characters, and are base-32 encoded. Next-generation firewalls (NGFWs) often misclassify Tor traffic due to encryption.

- Detections are based on information from the SSL log.
- Default severity: 7
- Enabled and set to alert by default.

5.2.2 Machine learning analysis

The Alerts tab in the Detailed View for machine learning detections shows details about the ML analysis, including features that contributed to the calculation of that score.

5 NXDOMAIN Beaconing | onsiteingo.com

Close Detection Suppress Entity Assign To

Detection Details **7 Alerts** 1 Related Detections 605 Related Entities

7 NXDOMAIN Beaconing alerts have been generated between July 21st, 11:09am and July 26th, 9:44am for onsiteingo.com as part of this detection

1-7 of 7 items

- 2023-07-26T09:44:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100
- 2023-07-25T22:58:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100
- 2023-07-25T12:10:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100
- 2023-07-23T06:14:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100
- 2023-07-22T08:43:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100
- 2023-07-21T21:58:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100
- 2023-07-21T11:09:00-07:00
NXDOMAIN Beaconing | onsiteingo.com
Model Score: 100

Details

Category: NXDOMAIN Beaconing
Timestamp: 2023-07-26T09:44:00-07:00
Source IP: 10.2.128.198
Domain: onsiteingo.com
Model Score: 100

ML Analysis

Top Feature Contribution

- 27% Average number of NXDOMAIN r...
- 22% Number of NXDOMAIN requests ...
- 19% Number of NX subdomains
- 12% Maximum number of DNS reques...
- 12% Number of DNS requests
- 2% Length of the longest subdomain
- 1% Length of the longest subdomain ...
- 1% Average length of the subdomains
- 1% Number of subdomains

Entity vs. Population

Feature	Actual	Mean	Std. dev	25th to 75th percentile	Distance from mean	x.x
Average number of NXDOMAIN r...	230				21.64	
Number of NXDOMAIN requests ...	230				0.57	
Number of NX subdomains	1				1.89	
Maximum number of DNS reques...	230				18.54	
Number of DNS requests	230				0.5	
Length of the longest subdomain	0				Feature has constant value	
Length of the longest subdomain ...	0				Feature has constant value	
Average length of the subdomains	0				Feature has constant value	
Number of subdomains	1				-0.17	

5.3 Search-based alerts

Search-based alerts are generated by Corelight-defined log search queries. The Corelight Labs team creates these alerts using LogScale searches, correlating threat indicators from the logs. This type of alert provides expanded coverage and rapid response.

You can find more information for each search-based alert by clicking the alert name in the Alert Catalog. The alert details provide a summary and suggest next steps for each detection to guide your investigation and troubleshooting.

The search-based alerts described in the following sections are available in Investigator, grouped by their MITRE ATT&CK categories.

5.3.1 Initial access

Initial Access consists of techniques that use various entry vectors to gain their initial foothold within a network.

The next sections describe the initial access search-based alerts.

5.3.1.1 Atlassian Preauth RCE

Atlassian Bitbucket CVE-2022-0540 RCE attempt detected.

CVE-2022-0540 is a command injection vulnerability in multiple API endpoints. In this attack, an attacker with access to a public repository or with read permissions to a private Bitbucket repository would be able to execute arbitrary code by sending a malicious HTTP request.

Default severity: 10

5.3.1.2 Cisco IOS XE Command Execution Attempt

Cisco IOS XE CVE-2023-20273 command execution detected.

The previously unknown vulnerability, which is tracked as CVE-2023-20198, resides in the Web User Interface of Cisco IOS XE software when exposed to the Internet or untrusted networks.

Default severity: 10

5.3.1.3 Cisco IOS XE Software Backdoor

Cisco IOS XE CVE-2023-20198 backdoor detected.

This previously unknown vulnerability, tracked as CVE-2023-20198, resides in the Web User Interface of Cisco IOS XE software when exposed to the Internet or untrusted networks. This alert detects backdoor activity associated with in-the-wild-exploits observed by Talos.

Default severity: 10

5.3.1.4 Citrix Netscaler CVE-2023-4966 (CitrixBleed)

Detects successful CVE-2023-4966 Citrix Netscaler exploits over HTTP.

CVE-2023-4966 is a software vulnerability found in Citrix NetScaler ADC and NetScaler Gateway appliances with exploitation activity identified as early as August 2023. This vulnerability provides threat actors, including LockBit 3.0 ransomware affiliates, the capability to bypass MFA and hijack legitimate user sessions.

Default severity: 10

5.3.1.5 Confluence Authorization Vulnerability

CVE-2023-22518 - Improper Authorization Vulnerability In Confluence Data Center and Server detected.

All versions of Confluence Data Center and Server are affected by this vulnerability. This Improper Authorization vulnerability allows an unauthenticated attacker to reset Confluence and create a Confluence instance administrator account. Using this account, an attacker can perform all administrative actions that are available to Confluence instance administrator leading to a full loss of confidentiality, integrity, and availability.

Default severity: 10

5.3.1.6 Curl Package Vulnerability

Curl CVE-2023-38545 and CVE-2023-38546

CVE-2023-38545 is a SOCKS5 heap buffer overflow, which makes curl overflow a heap-based buffer in the SOCKS5 proxy handshake. CVE-2023-38546 is a cookie injection which allows an attacker to insert cookies at will into a running program, using libcurl, if a specific series of conditions are met.

Default severity: 10

5.3.1.7 DApp Injection

Possible Water Labbu cryptocurrency theft attempt detected.

Threat actor Water Labbu capitalizes on social engineering schemes of other scammers, injecting malicious JavaScript code into decentralized application websites of other scammers to steal cryptocurrency.

Default severity: 8

5.3.1.8 Fortigate RCE CVE-2024-21762

Detects a successful Fortigate Remote Code Execution (RCE) per CVE-2024-21762.

CVE-2024-21762 is an RCE vulnerability in Fortigate FortiOS and FortiProxy. An out-of-bounds write vulnerability in FortiOS and FortiProxy can let a remote unauthenticated attacker execute arbitrary code or command through specially crafted HTTP requests. A Proof of Concept exploit was publically released and Fortigate made a patch available.

Default severity: 10

5.3.1.9 Fortiguard Auth Bypass

Fortiguard CVE-2022-40684 auth bypass attempt detected.

This vulnerability allows adversaries to bypass authentication and login into target systems as an administrator in FortiOS / FortiProxy / FortiSwitchManager products. With these privileges, the adversary may create new users, update or download network/system configurations, reroute traffic, or listen to and capture sensitive data by running packet capturing programs.

Default severity: 10

5.3.1.10 Fortinet Key Upload

Fortinac CVE-2022-39952 exploitation attempt detected.

External control of a filename or path in Fortinet FortiNAC may allow an unauthenticated attacker to execute unauthorized code or commands through a specifically crafted HTTP request.

Default severity: 10

5.3.1.11 Fortra FileCatalyst remote code execution CVE-2024-251538

Detects Fortra FileCatalyst successful remote code execution (CVE-2024-25153).

CVE-2024-25153 is a Remote Code Execution vulnerability in Fortra FileCatalyst. A Proof of Concept exploit has been publically released and Fortra has made a patch available.

Default severity: 9

5.3.1.12 JetBrains TeamCity Authentication Bypass CVE-2024-27198

Detects systems vulnerable to JetBrains TeamCity Authentication Bypass CVE-2024-27198.

Vulnerability CVE-2024-27198 lets attackers bypass JetBrains authentication using an alternate path or channel. This vulnerability allows for a potential compromise of a vulnerable TeamCity server by a remote unauthenticated attacker. Compromising a TeamCity server gives an attacker full control over all TeamCity projects, builds, agents and artifacts, and could position an attacker to perform unauthenticated remote code execution and a supply chain attack. A patch is available through JetBrains.

Default severity: 9

5.3.1.13 JetBrains TeamCity Authentication Bypass CVE-2024-27199

Detects systems vulnerable to JetBrains TeamCity Authentication Bypass CVE-2024-27199.

Vulnerability CVE-2024-27199 lets attackers bypass JetBrains authentication using an alternate path or channel. This vulnerability allows for a limited amount of information disclosure and a limited amount of system modification, including the ability for an unauthenticated attacker to replace the HTTPS certificate in a vulnerable TeamCity server with a certificate of the attacker's choosing. A patch is available through JetBrains.

Default severity: 8

5.3.1.14 Kali Presence

A host using the Kali offensive security Linux distribution was detected performing software updates. Attackers may use these software updates to move laterally through the network and/or potentially exfiltrate data.

Kali Linux is an open-source Linux distribution geared toward information security tasks, including penetration testing, security research, computer forensics, and reverse engineering. A host using this distribution may be able to perform lateral movement or data exfiltration.

Default severity: 9

5.3.1.15 OpenSSL Punycode

OpenSSL CVE-2022-3602 exploitation attempt detected.

Exploit CVE-2022-3602 is an arbitrary 4-byte stack buffer overflow that has been assessed as critical by the OpenSSL project. Exploitation of this vulnerability may lead to remote code execution.

Default severity: 10

5.3.1.16 Proxy Not Shell

MS Exchange CVE-2022-41040 / CVE-2022-41082 ProxyNotShell exploitation attempt detected.

ProxyNotShell exploitation of Microsoft Exchange Server 2013, Exchange Server 2016, or Exchange Server 2019 may allow attackers to perform Server-Side Request Forgery or Remote Code Execution. Authenticated access to the vulnerable Exchange Server is necessary to successfully exploit either vulnerability.

Default severity: 10

5.3.1.17 ScreenConnect Authentication Bypass CVE-2024-1709

Detects systems vulnerable to ScreenConnect Authentication Bypass CVE-2024-1709.

Critical vulnerability CVE-2024-1709 allows anonymous attackers to exploit an authentication bypass flaw to create admin accounts on publicly exposed instances. Using the system admin role would allow the attacker to delete other users and take over the instance.

Default severity: 10

5.3.1.18 Sonicwall RCE/DOS CVE-2022-22274 / CVE-2023-0656

Detects systems vulnerable to Sonicwall CVE-2022-22274 and CVE-2023-0656 over HTTP. A stack-based buffer overflow vulnerability in the SonicOS through HTTP request allows a remote unauthenticated attacker to cause Denial of Service (DoS) or potentially results in code execution in the firewall. This vulnerability only impacts the web management interface, the SonicOS SSLVPN interface is not impacted.

Default severity: 10

5.3.1.19 Text4Shell

Apache Commons Text CVE-2022-42889 log4text exploitation attempt detected.

In Apache Commons Text, a default interpolator allows for string lookups that can lead to Remote Code Execution. This is due to a logic flaw that makes some lookup keys (script, dns, and url) interpolated by default. Those keys allow an attacker to execute arbitrary code through lookups and possibly perform remote code execution (RCE) to execute arbitrary code on the machine and compromise the entire host.

Default severity: 10

5.3.1.20 WS_FTP Remote Commands

WS_FTP CVE-2023-40044 RCE attempt detected.

In WS_FTP Server versions prior to 8.7.4 and 8.8.2, a pre-authenticated attacker could leverage a .NET deserialization vulnerability in the Ad Hoc Transfer module to execute remote commands on the underlying WS_FTP Server operating system.

Default severity: 10

5.3.2 Command and control

Attackers often use software that is controlled through a mechanism called a C2 (command and control) channel. To avoid detection, malicious software often attempts to obfuscate the C2 traffic to make it look benign. HTTP traffic is very commonly used to carry C2 communications since it easily traverses between most organizations and the internet.

The next sections describe the search-based alerts related to command and control.

5.3.2.1 AsyncRAT remote access trojan

AsyncRAT is a remote access trojan (RAT) that uses an encrypted C2 channel to allow remote monitoring and control of the infected machine, including keylogging, screen recording, controlling the desktop and webcam, running remote shells, and injecting new payloads.

Default severity: 10

5.3.2.2 BunnyLoader 1

BunnyLoader C2 communications using known uri and user_agents.

Malware loader BunnyLoader provides functionalities such as downloading and executing malware, stealing browser credentials and system information, keylogging, stealing credentials, and running remote commands on the infected machine.

Default severity: 10

5.3.2.3 BunnyLoader 2

BunnyLoader C2 communications using known uri parameters.

Malware loader BunnyLoader provides various functionalities such as downloading and executing malware, stealing browser credentials and system information, keylogging, stealing credentials, and running remote commands on the infected machine.

Default severity: 10

5.3.2.4 Hyperscrape

Possible APT 35 HYPERSCRAPE malware data exfiltration detected.

Default severity: 10

5.3.2.5 Kolobko

Malware C2 communications associated with Lapsu\$/UNC2447/Yanluowang detected.

Default severity: 10

5.3.2.6 Manjusaka C2 search 2

Manjusaka malware C2 communication detected.

Default severity: 10

5.3.2.7 Manjusaka C2 search 3

Manjusaka malware “keep alive” detected.

Default severity: 10

5.3.2.8 Nim Plant C2

NimPlant C2 communications detected.

Default severity: 10

5.3.2.9 Onion domain

Top level onion domain detected.

Anonymous websites on the Tor network utilize top-level domain .onion, which is accessible only from the Tor anonymity browser. Presence of .onion domains in the logs might indicate the originating system (a potential attacker) is establishing TOR connections and C2 communications to hide the destination of the connections and evade blacklist-based detection.

Default severity: 7

5.3.2.10 Scanbox

SCANBOX browser exploitation framework traffic detected.

The ScanBox browser exploit starts with phishing emails that include links to a website impersonating Australian media entities such as Australian morning news. Through the phishing attempt, victims visit an infected website, where they would be infected with SCANBOX malware. This detection alerts if visitors clicked the link to the phishing site or if communications occur between the implant and the C2 over HTTP.

Default severity: 10

5.3.2.11 Sliver GET URI

Sliver exploitation framework network traffic detected.

Default severity: 10

5.3.2.12 Sliver POST URI

SLIVER exploitation framework network traffic detected.

Default severity: 10

5.3.2.13 Vidar C2

Malware C2 communications associated with Vidar C2 detected.

Vidar C2 Malware has the ability to collect sensitive information from an infected computer and exfiltrate this data. Vidar may collect a variety of information from infected computers, browsers, and digital wallets including OS data, credentials, and browser history.

Default severity: 10

5.3.3 Credential access

Credential Access consists of techniques for stealing credentials like account names and passwords.

The next sections describe the search-based alerts related to credential access.

5.3.3.1 Confluence Hardcoded Password

Confluence CVE-2022-26138 hardcoded password usage detected.

Confluence user accounts with hardcoded credentials stored inside the plugin jar file may be created by the Questions for Confluence app. An attacker with knowledge of these credentials could log into the Confluence application and access all contents within the confluence-users group. Atlassian has rated the vulnerability critical and highlighted that the vulnerability is being exploited in the wild.

Default severity: 10

5.3.3.2 Kerberos Weak Ciphers

Kerberos tickets are used to obtain access to resources. If weak encryption such as RC4 is used, it is possible to obtain passwords using attacks such as Kerberoasting. If a user on the network were to attempt to use such a ticket, this search would detect this ticket and generate an alert.

Default severity: 6

5.3.4 Privilege escalation

Privilege escalation consists of techniques that adversaries use to gain higher-level permissions on a system or network.

The next section describes the search-based alerts related to privilege escalation.

5.3.4.1 Confluence Server Privilege Escalation

Potential CVE-2023-22515 attempt, Zero-Day Privilege Escalation in Confluence Server and Data Center.

Default severity: 10

5.4 Normalized severity scores

Investigator integrates Notices, Suricata, and ML alerts. Each of those sources calculates severity scores for alerts in their own way. Investigator normalizes scores from each source and maps the original scores into a comparable value ranging from 1 to 10, with 10 being the most critical.

- Notices – Imported Notices scores range from 0 to 7, with 0 being the most critical. Investigator normalizes the scores according to this table.

Original severity	Normalized severity
0	10
1	9
2	8
3	4
4	2
5	1
6	1
7	1
Unknown	2

- Suricata – Imported Suricata alert scores range from 1 to 4, with 1 being the most critical. Investigator normalizes Suricata alert scores according to this table.

Original severity	Normalized severity
1 (Critical)	8 (The alert review process can increase this score)
2 (Major)	6
3 (Minor)	3
4 (Informational)	2
Unknown	2 (Defaults to informational level)

- ML models, generated by Investigator, range from 1 to 10, with 10 being the most critical

5.5 CrowdStrike data in detections

The CrowdStrike integration blends CrowdStrike EDR with Investigator network detection capabilities and maps Corelight IP addresses from detections to CrowdStrike host information. The expanded data provides additional context to analyze threats and helps analysts make informed decisions during the triage process.

With a *CrowdStrike integration configured and enabled*, the detections show additional host information from CrowdStrike data to provide even more context. The detection summary highlights some of the most important host information, and the detection details show greatly expanded entity information with even more context.

The CrowdStrike information is seamlessly integrated with the Investigator content. To view the enhanced entity content, go to the Detections page, click a detection, and review the Entity section.

Entity
^

Entity 96.35.155.226

Entity Type IP

Data Source Crowdstrike

Entity Status 🔄

[Isolate Entity](#)

HOST INFORMATION

Timestamp March 7, 2024 1:05pm

Mac Address 00-15-5d-c4-db-74

Host Name JAMES-DESKTOP

Platform Windows

Platform Type Workstation

OS Windows 11

External IP 96.35.155.226

Device ID ecc6a481d55f40a684db15f7512103f2

Sensor Version 7.06.17807.0

Cloud Account ID ---

Cloud Instance ID ---

Cloud Provider ---

AD Domain Name ---

AD Organizational Unit Name ---

INTERFACE HISTORY

IP Address	Mac Address	Timestamp
96.35.155.226	00-15-5d-c4-db-74	March 12th, 1:44pm
172.30.176.1	00-15-5d-c4-db-74	March 11th, 12:27pm
192.168.10.175	04-d9-f5-82-72-c0	March 11th, 12:25pm
192.168.10.175	04-d9-f5-82-72-c0	March 11th, 12:25pm
100.97.0.124		February 11th, 6:58am

[View All Interface History](#)

USER HISTORY

User	Last Login
NT VIRTUAL MACHINE\54AFFE6F-3F30-4964-A355-C34CF...	March 12th, 2:10pm
NT VIRTUAL MACHINE\193EC090-D662-4252-9571-87C14...	March 12th, 2:10pm
NT VIRTUAL MACHINE\F3082CCA-B619-465D-8B22-0DB0...	March 12th, 2:10pm
NT AUTHORITY\LOCAL SERVICE	March 7th, 1:05pm
JAMES-DESKTOP\James	March 7th, 1:05pm

[View All User History](#)

Without this integration, the Entity section for a detection has two values: Entity and Entity Type. With this integration, you can see the following additional fields in the Entity section of the detection details page.

Field	Description
Timestamp	The timestamp of device's most recent connection to Falcon.
MAC Address	The MAC address of the device.
Hostname	The name of the machine.
Platform	The operating system platform, such as Linux, Mac, or Windows.
Platform Type	The name of the product type, such as Workstation, Server, or Domain Controller.
OS	The version of the operating system, such as Windows Server 2012 R2.
External IP	The external IP address of the device, as seen by CrowdStrike.
Device ID (Host ID)	The ID of the device, which is obtained from CrowdStrike. You can use this ID to look up information in the CrowdStrike console.
Sensor Version	The version of the CrowdStrike sensor, such as 7.06.16108.0.
Cloud Account ID	The cloud account ID, if applicable.
Cloud Instance ID	Cloud resource information, if applicable.
Cloud Provider	The cloud service provider, such as AWS_EC2_V2, if applicable.
AD Domain Name	Active Directory domain name.
AD Organizational Unit name	Active Directory organizational unit name.
Interface History	The Interface History displays the most recent IP addresses and MAC addresses used by a host and a timestamp indicating the time of access. The 5 most recent entries appear; click View All Interface History to see more entries.
User History	The User History displays the most recent user accounts to log in to the device and the time of their last login. The 5 most recent entries appear; click View All User History to see more entries.

5.5.1 Isolate entities

With a *CrowdStrike integration configured* and set up to isolate entities, admin users can quickly isolate an entity from all network activity to address potential compromises or ongoing attacks based on Corelight evidence. The entity isolation uses the [Network Contain](#) feature in CrowdStrike Falcon and CrowdStrike Falcon must recognize the entity.

As an admin user, you can isolate an entity and quickly respond to an attack by implementing quarantine measures. The entity isolation is integrated into the Investigator detection workflow.

The Entity Status field reflects the current state of network containment for the entity.

To isolate an entity

- From the Entity section of the Detections page, click **Isolate Entity**. You are prompted to confirm the action, and click **Isolate Entity** again.

Investigator submits the request to trigger the Network Contain feature in CrowdStrike Falcon and implement quarantine measures for the specified entity.

When submitted, the entity status changes to Pending Isolation and then to Isolated when the request succeeds. (An unsuccessful request displays an error message to help you understand and address the issue.)

When isolated, the entity loses the ability to make network connections to anything other than the CrowdStrike cloud infrastructure. Also, the **Isolate Entity** buttons changes to a **Lift Isolation** button.

To remove a device from network isolation

- Click **Lift Isolation**. You are prompted to confirm the action, and click **Lift Isolation** again.

Investigator submits the request and the entity status changes to Pending Lift and then to Normal, and the device resumes normal network activity.

Note: Analyst users can see the **Isolate Entity** and **Lift Isolation** buttons, but cannot perform the related actions.

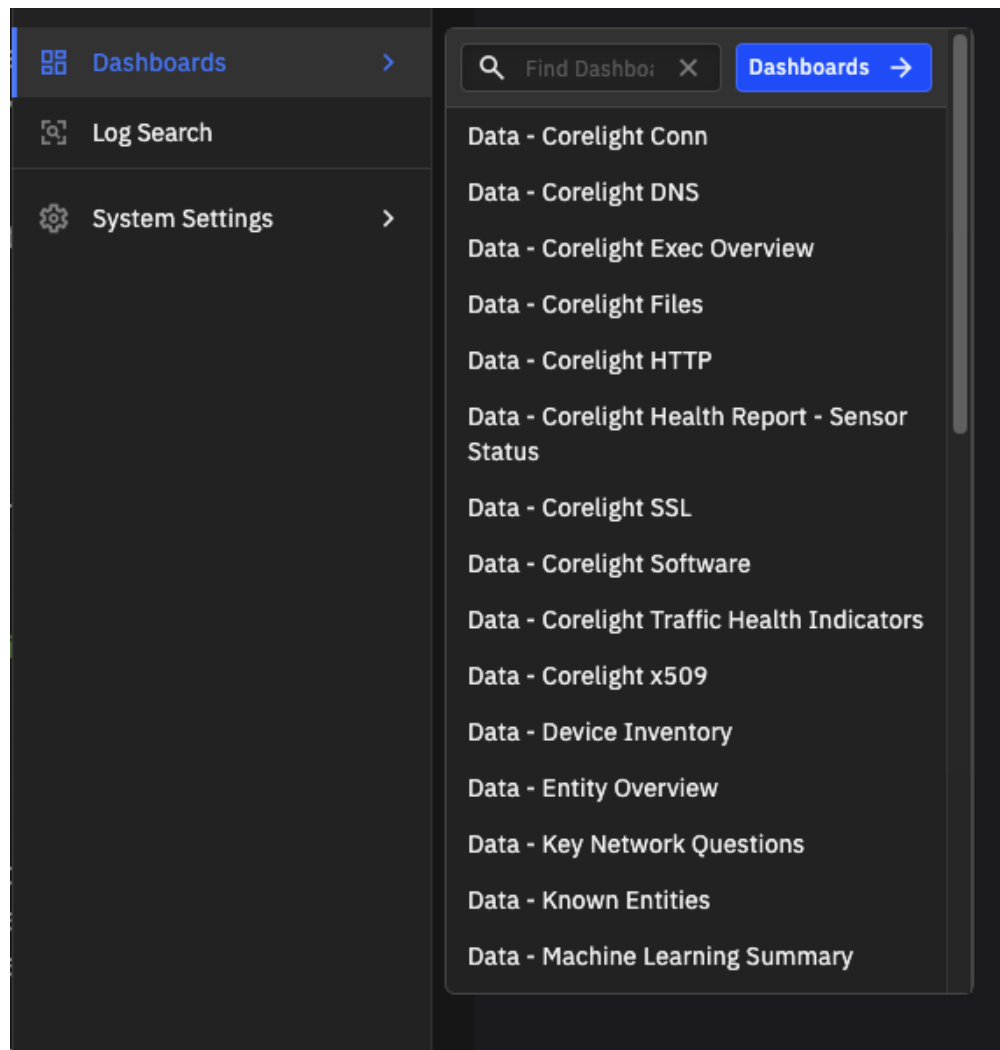
5.5.2 Additional learning resources

Watch a Corelight video on YouTube: [CrowdStrike Host Isolation Integration](#)

EXPLORE DATA THROUGH DASHBOARDS

Dashboards are a convenient way to monitor events and explore data. Investigator provides a set of pre-built [Falcon LogScale \(Humio\)](#) dashboards for common views corresponding to protocols, logs, or specific content collections and detections.

Click **Dashboards** in the left navigation to get started. You can choose a specific dashboard to view, search for a dashboard, or click the **Dashboards** button to go to the primary landing page for all dashboards.



The Corelight packaged dashboards have both data-related dashboards and security-related dashboards. Some dashboards are new to Investigator.

The available dashboards include:

- **Data - Corelight Conn** – Summarizes activity in `conn` logs, showing analytics of top hosts, services, ports, data transfers, and long-lived connections. The dashboard also includes health monitoring of potentially unavailable services.
- **Data - Corelight DNS** – Summarizes host and query information in DNS logs, including details of top query types, reverse queries, queries with no response, queries to non-existent domains, and top hosts.
- **Data - Corelight Exec Overview** – Provides a high level overview across multiple log sources summarizing top applications, usernames, websites, services, data transfer, and geographic information.
- **Data - Corelight Files** – Provides a summary of `files` log activity, including information on MIME types, file protocols, file flow, bytes sent and received, and executables.
- **Data - Corelight HTTP** – Provides a summary of HTTP log activity, including information on referrers, users, hosts, body length, user agents, host headers, originators, and status codes.
- **Data - Corelight Health Report - Sensor Status** – Displays all the sensors that are connected to Investigator and displays sensor status, including log types, connections, and event rates.
- **Data - Corelight SSL** – Summarizes information from SSL and x509 logs, including details on ciphers, TLS versions and validation status, certificate summary, and certificate subjects.
- **Data - Corelight Software** – Summarizes information in software logs, including a breakdown of top software by connection, and software versions and types.
- **Data - Corelight x509** – Summarizes information in x509 logs, including top and rare subjects and certificate expiration information.
- **Data - Entity Overview** – Summarizes the network entities.
- **Data - Key Network Questions** – Shows essential network information to answer important questions, including what network technologies are in use, what systems are providing core services/access services/file transfer services, plus bandwidth measurements and conversation tracking.
- **Data - Known Entities** – The Known entities logs (including `known_users`, `known_hosts`, and `known_devices`) extract and aggregate behavior for individual network entities. The Known Entities dashboard summarizes the entity behavior across all of these logs.
- **Data - Machine Learning Summary** – Provides an overview of all machine learning models producing alerts, including model result output and detection counts.
- **Security - Corelight IP Interrogation** – Provides a summary of protocol usage, and information on top connections, services, user agents and ports from `conn` and `http` logs.
- **Security - Corelight Intel** – Overview of intel logs, including volume over time, indicators, and summary of details.
- **Security - Corelight Log Hunting** – Overview of log volume and data across all log sources.
- **Security - Corelight Notice** – Overview of Notices, including details on volume of alerts over time and alert categories.
- **Security - Corelight RDP Inference Overview** – Provides details inferred from authentication requests to the server in RDP connections. Details include inference type, inferences over time, successful/failed connections, security protocol, and details for connecting users.
- **Security - Corelight SSH Inference Overview** – Provides details inferred from SSH login attempts. Details include inference type, inferences over time, HASSH fingerprint details, SSH host key, SSH authentication, SSH auth success, and inference log data.

- **Security - Corelight Suricata** – Overview of Suricata alerts, including details on volume of alerts over time and alert categories
- **Security - Corelight VPN Insights** – Provides details of VPN protocol connections from the vpn log, including top VPN users and VPN types.

Investigator also includes a set of dashboards with matches to the CrowdStrike Indicator of Compromise (IOC) database.

- **Security - IOC - Overview** – Provides a summary of IOC activity and threat details, including threat by confidence, threat attributes, and severity over time.
- **Security - IOC - IP Overview** – Displays an IOC dashboard based on IP addresses. The dashboard shows IOC geolocation, threat relationships, threat types, malware, actors, and kill chains.

Consult these Falcon LogScale (Humio) documentation topics to learn more about how to work with and manage dashboards.

- [Managing dashboards](#)
- [Using dashboards](#)
- [Editing dashboards](#)

Tip: If you frequently use a dashboard, click the star to the left of the dashboard name to mark it as a favorite. Favorites appear at the top of the list.


FIND DETAILS IN THE LOGS

Investigator provides detailed, searchable access to the data ingested from logs. With this data, you can build graphs, identify anomalies, and create alerts to monitor your systems.

Imported logs are available in the log search page.

7.1 Search the logs

Click **Log Search** in the left navigation to get started.



Key components of the search let you set the time interval, enter a query, and choose how to display the data.

The screenshot shows the Investigator Log Search interface. At the top, there's a navigation bar with 'Event List', 'Queries', 'Language syntax', and 'Event list widget'. A search filter 'Last 24h (Static)' is applied. The main area displays a table of log events with columns like @timestamp, #path, id.orig_h, id.resp_h, id.resp_p, uid, and system_name. A left sidebar shows filter fields and a 'Save as...' button is highlighted in the top right.

Investigator uses Falcon LogScale (based on Humio technology) to provide these extensive search capabilities. Consult the Humio documentation for a complete reference to the functionality.

These topics are a great place to start.

- [Create and customize search queries](#)
(This link to the Language Syntax help is also available at the top of the Log Search page.)
- [Using query functions to transform and aggregate data](#)

As an example, follow these steps to find the top entries for logs and display the results as a pie chart.

1. Use the query `top(path)` to display the top entries for the different log types.

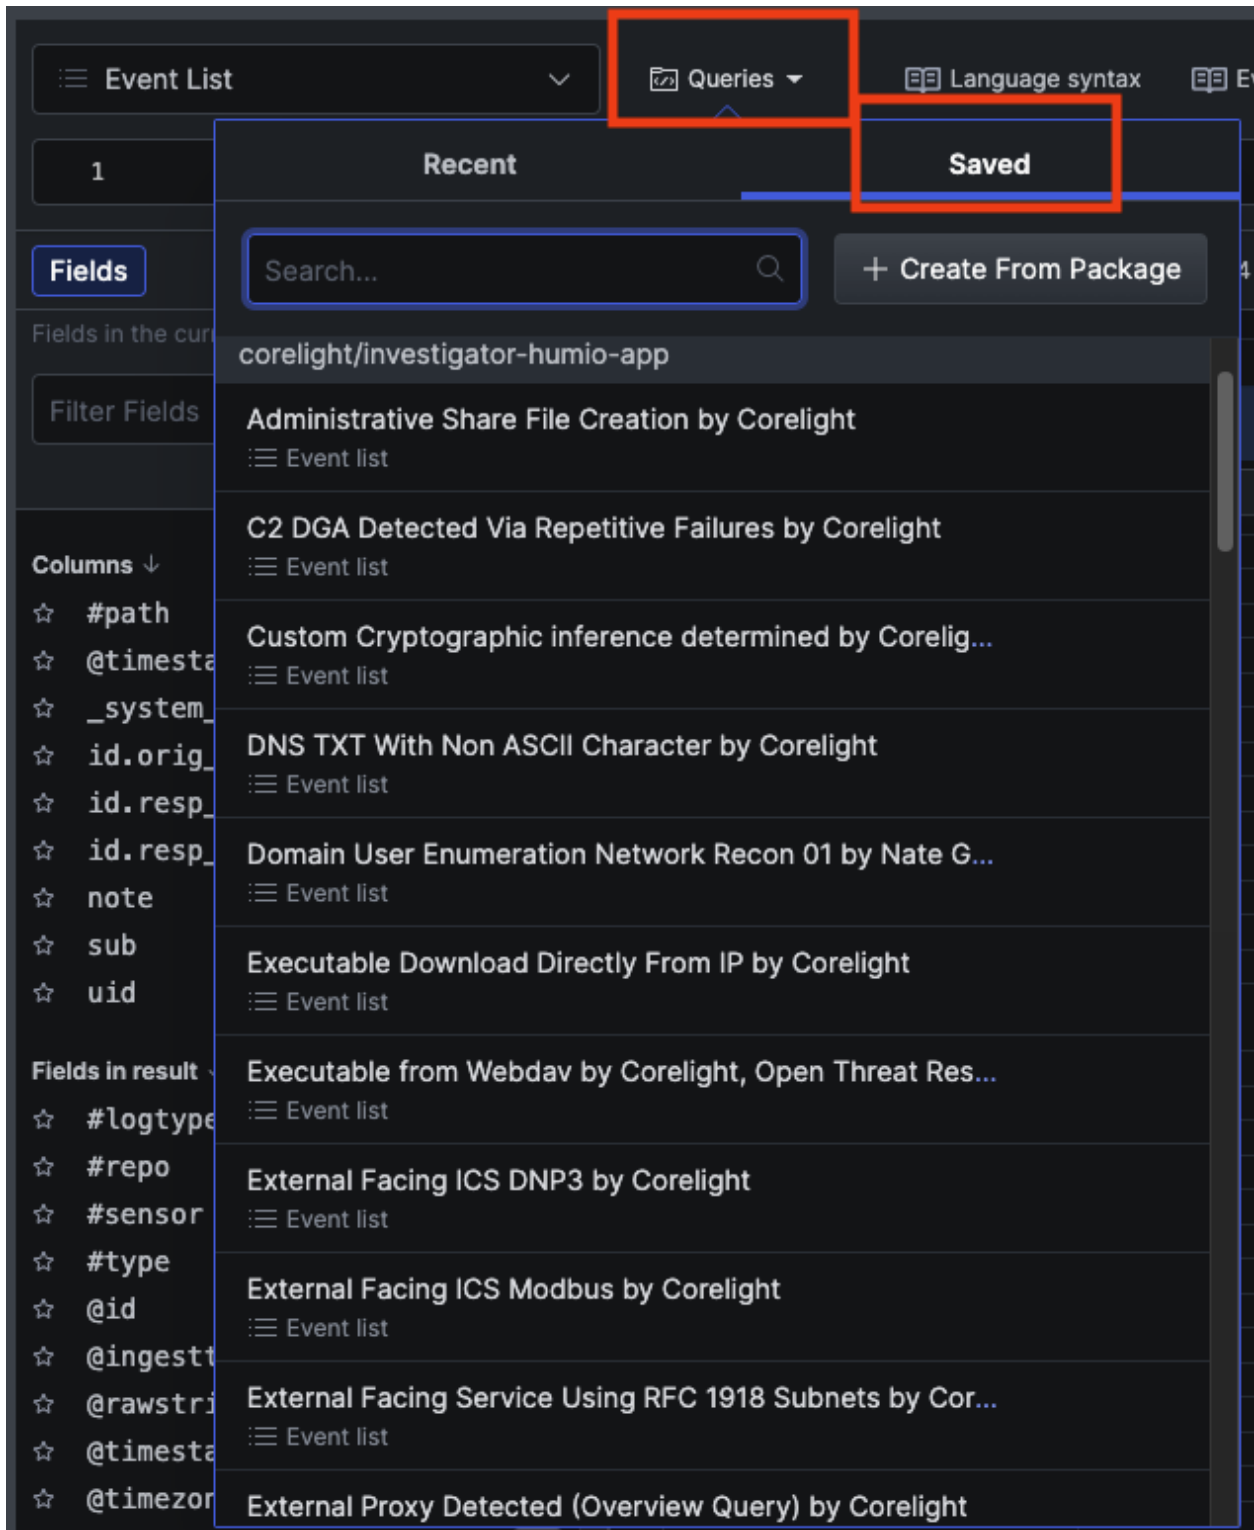
The page shows you the search progress.

2. Add filters to the query to focus the results.
3. From the top left menu, you can change how the results appear.

For example, you can change to a pie chart or a table.

4. Click **Save** to add the query to your saved list, or save in several other formats such as a Dashboard widget.

The Log Search page also includes a set of saved searches to assist in threat hunting. To access the saved searches, click **Queries** and choose the **Saved** tab.

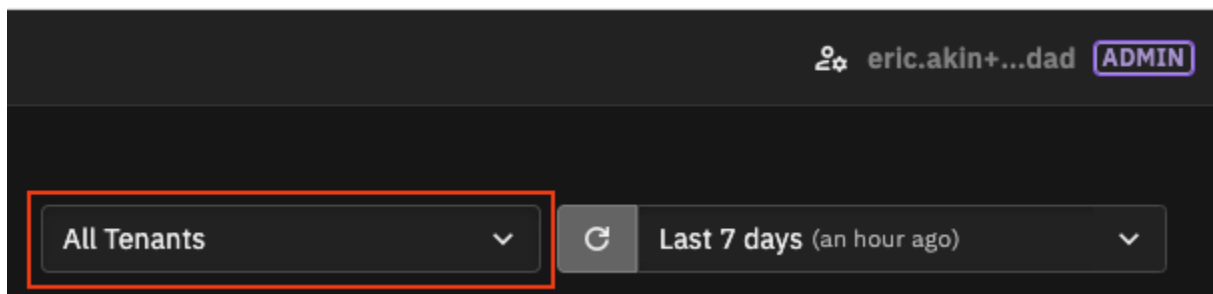


FEDERATED TENANTS

The federated tenants feature provides a collective view of data from configured sub-tenants (child tenants). The feature streamlines the management of tenants across different business units and lets administrators efficiently manage detections from all tenants. The aggregated data provides comprehensive insights into network security within a unified interface.

A federated tenant is an umbrella account that aggregates data on configured sub-tenants; the federated account does not have a sensor directly associated with it and does not ingest its own data.

Investigator pages that support federated views include a Tenant menu that lets you choose the aggregate view (All Tenants) or select an individual tenant to view only their data.



Child tenants can only access data collected from the sensors configured specifically for their respective tenant.

Contact your Corelight account manager to set up federated tenants. You can federate existing tenants or add new tenants.

Federated tenants can view child tenant data on these pages:

- Security Overview Page – provides an aggregate view of all tenants and the ability to switch to individual tenant views, allowing for comprehensive monitoring and management of security across the entire system.
- Detections Page – provides an aggregate view of all tenants and the ability to switch to individual tenant views.

Federated tenant users can take actions such as closing a detection, sending a detection to ServiceNow, or suppressing an entity from the federated tenant.

All actions appear in the security audit and are logged against respective tenants. Detections cannot be assigned to users from a different tenant.

Each detection identifies the associated tenant, both in the list/table view and in the details.

- Dashboards – let you view data on LogScale dashboards across all tenants for a comprehensive view.
- Logs – displays the logs for all tenants. Each log entry includes a sensor tag (system_name) for identification, enabling LogScale queries to retrieve logs from specific child tenants.
- Alert Catalog – provides a per-tenant view only, ensuring admins focus on managing alerts specific to the selected tenant without interference from the data of other tenants. Admins in a federated account can switch the view between child tenants.
- General Settings – admins can configure the tenant name, ensuring accurate and consistent tenant identification across the system. Admins must log in to a specific tenant to change the tenant name.

To change a tenant name

1. Log in to the individual tenant.
2. Go to **General Settings | Tenant Settings** and click the **Edit** icon for the tenant.
3. Edit the display name and click **Save**.

Tenant names can include text, numbers, and special characters.

- Security Audit – provides an aggregate view of all tenants and a per-tenant view of the audit logs.

Note: Account Settings are per tenant and do not have a federated view.

Other changes to note:

- Integrations Page – is hidden for federated tenants. Admins must log in to a specific tenant to manage integrations.

- User Management – displays only users from the federated tenant, ensuring that access to user information is restricted to the relevant tenant for enhanced security and data privacy. Admins must log in to a specific tenant to manage users for that tenant.
- Investigator does not support duplicate email addresses. Federated tenant admins need to modify their email addresses to get access, typically by adding a plus (+) sign and an extra identifier.
- Detections across child tenants do not have a unique detection ID. Within a child, a tenant detection ID is unique.

ACCOUNT SETTINGS

Within account settings, you can review and manage these settings:

9.1 Account alias

Each user account has an alias, initially defined by the account admin. The alias is not a username, but essentially a nickname for the account. You can change your alias in account settings.

To change your alias

1. Click your username in the upper-right corner to display the menu and choose **Account Settings**.
2. In your User Profile, enter a new alias.
Your alias cannot be more than 30 characters.
3. Click **Save**.

9.2 Cookies

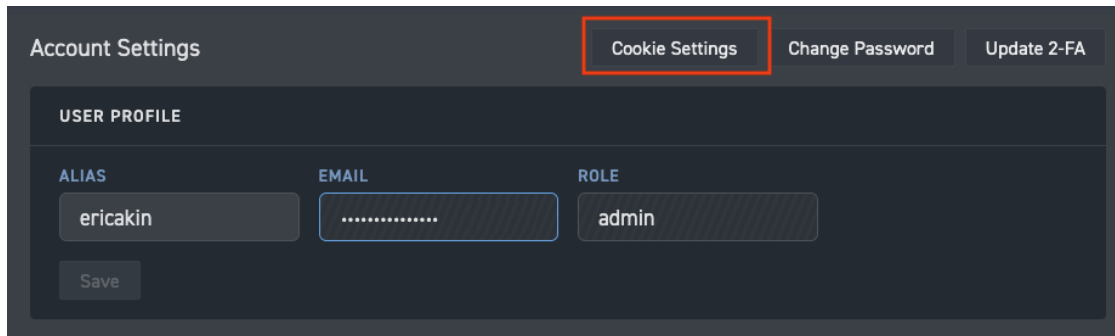
The Corelight Investigator website uses cookies to provide you with personalized content and to facilitate, improve, and analyze site operations. The site stores cookies in two categories:

- Necessary cookies - These cookies are required for the operation of our site and cannot be turned off. The information collected relates to our product operations and lets Corelight provide services.
- Analytical cookies - These cookies are optional. Corelight uses these cookies to collect information on how visitors use the product features. They help us to improve our website and customize it to better fit your needs. All information collected by these cookies is aggregated and anonymous.

When you log in to Investigator for the first time, you are prompted to specify your cookie preferences. You can choose one of these options:

- Accept all Cookies - consent to the use of all cookies, including non-essential cookies, and the related processing of your personal data.
- Cookie Settings - open account settings where you can see details about each type of cookie, and you can specify which cookies you want to accept.

You can manage your cookies at any time from the Cookie settings tab available on the [Account Settings](#) page.



The screenshot shows the 'Account Settings' interface. At the top, there are three tabs: 'Account Settings', 'Cookie Settings' (highlighted with a red box), 'Change Password', and 'Update 2-FA'. Below the tabs is a 'USER PROFILE' section. It contains three input fields: 'ALIAS' with the value 'ericakin', 'EMAIL' with a masked password '.....', and 'ROLE' with the value 'admin'. A 'Save' button is located below the input fields.

Cookies are stored based on the user and not the browser. Corelight uses your browser's local storage for some cookie information.

For more information about Corelight's use of cookies, go to our [Privacy & Cookies Policy](#).

9.3 Password

Any user can update their password. User passwords must meet these requirements:

- At least 8 characters long
- At least one digit
- At least one special character
- At least one upper case character
- At least one lower case character

To change your password

1. Click your username in the upper-right corner to display the menu and choose **Account Settings**.
2. Click **Change Password**.
3. Enter your current password, enter a new password, and confirm the new password.
4. Click **Confirm with 2FA**.
5. Enter your OTP and click **Verify**.
6. Click **Confirm**.

If you forget your password, you can change your password from the login screen.

9.4 Two-factor authentication

Two-factor authentication (2FA) is mandatory for all users, using a TOTP (time-based one-time password).

9.4.1 Update 2FA access

You set up 2FA during registration, but any user can update their preferred method for receiving 2FA verification codes.

Note: If you update your 2FA, you will be logged out and will be required to reconfigure 2FA for the next sign in.

To update your 2FA method

1. Click your username in the upper-right corner to display the menu and choose **Account Settings**.
2. Click **Update 2-FA**.
3. Enter your current password and click **Confirm**.
4. Enter your TOTP to confirm your existing 2FA and click **Verify**.
5. Click **Confirm**.
Investigator logs you out.
6. Enter your login email and password and click **Sign In**.
7. Open an authenticator app (such as Google Authenticator) on your phone and add a new account.
8. Scan the QR code that Investigator displays to connect your app to Investigator.
9. Enter the TOTP from your app in Investigator.
10. Click **Verify OTP** to log into Investigator.
Your 2FA is updated.
11. Click **Login** to sign in to Investigator with your email, password, and new 2FA method.

9.4.2 Recover 2FA access

If you cannot log in to Investigator due to an 2FA issue (for example, if you lose access to the registered authentication app, lose your device, or if there was an error configuring 2FA during registration), you can work with an admin user to recover 2FA access. An admin can reset 2FA access by deactivating and then reactivating your user account.

To recover user access

1. As an admin user, go to **System Settings | Users & Access**.
System Settings are available through the left navigation.
2. Select the user with the access issue and click the **Edit ()** icon.
3. Change their status from Active to Inactive and click **Save**.
4. Select the user again and click the **Edit ()** icon.
5. Change the status from Inactive to Active and click **Save**.

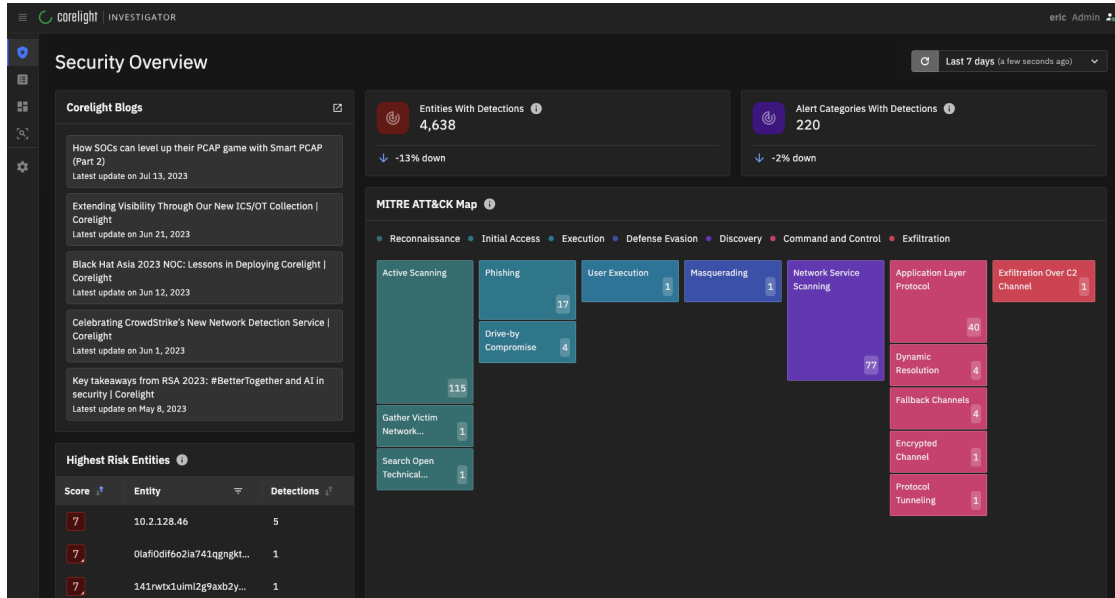
When reactivated, the system sends an email to the user with password and 2FA reset instructions.

9.5 Theme settings

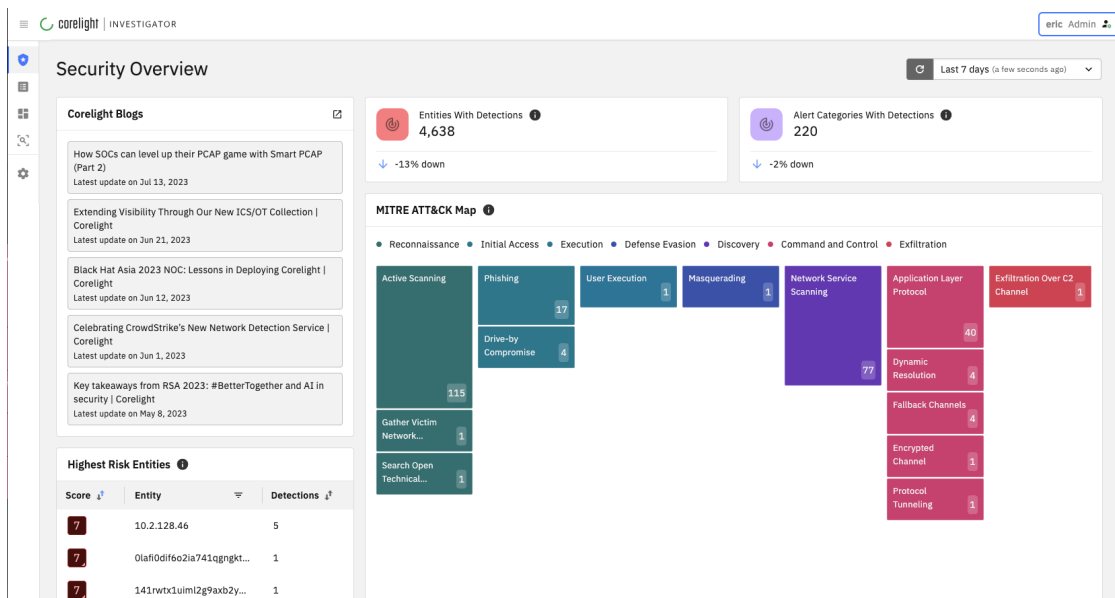
Investigator provides two views: a light setting and a dark setting.

To switch your theme setting, click your username in the upper-right corner to display the menu and choose **Switch to Light Theme** or **Switch to Dark Theme**.

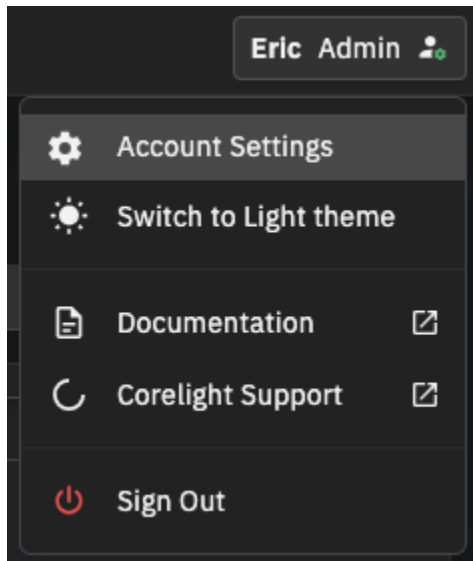
Dark theme



Light theme



You can access account settings from the menu that appears when you click your username in the upper-right corner of the page.



SYSTEM SETTINGS

To access the system settings, click the gear icon in the left navigation.

Depending on your role, you can view and configure these system settings:

General Settings

10.1 Licensing

Corelight Investigator offers an Advanced license and also an evaluation version. The type of license determines the features and functionality.

This table summarizes the features supported by each license type.

Feature	Advanced Eval	Advanced
Incident Response		
Detection triage and workflow	✓	✓
Alert aggregation, prioritization, and tuning	✓	✓
Analytics		
Corelight sensor collections	✓	✓
Suricata IDS + Proofpoint ET Pro ruleset	✓	✓
Cloud-based ML detections	✓	✓
CrowdStrike Falcon X IOC database	✓	✓
Data Retention		
Investigator alerts & detections	90 days	90 days
Full Zeek + Suricata logs	30 days	30 days
Additional Zeek + Suricata log retention	Optional	Optional
Data Export to SIEM/XDR		
Full Zeek + Suricata log export from sensor	✓	✓
Alert export from Investigator	✓	✓
Administration & Integration		
SAML / SSO	✓	✓
Security auditing	✓	✓
Fleet Manager	✓	✓
Smart PCAP	✓	✓
Support & Services		
Standard support	✓	✓
Enterprise support	—	Optional
QuickStart service	✓	✓
Managed threat hunting services	—	Optional

10.1.1 License status

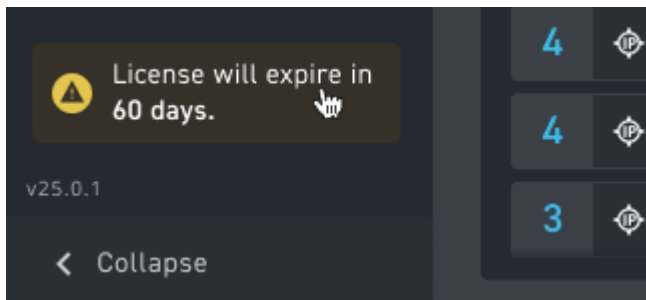
You can view your license status and details at any time. From the **System Settings** in the left navigation, choose **General Settings**.

The License Status section displays your license information, including the start date, the expiration date, and the primary contact for your account. The section also shows the type of license you have and the log retention period.

The license information is read only; contact Corelight Support or your Account Manager to make any changes.

10.1.2 License expiration

Customers will receive warnings starting at 60 days before a license expires. The Investigator interface displays a warning in the left navigation panel and indicates the number of days before license expiration.



The system also sends an email notification to account admins at 60 and 30 days before expiration and when the license expires.

Once a license expires, account users cannot log in to Investigator. Corelight keeps the account infrastructure for a 90-day grace period and after that, deletes all infrastructure.

Contact Corelight Support or your Account Manager to renew your license.

With an Advanced license, Investigator imports all log data. Imported logs are available in the log search page.

10.2 Autoclose detections

By default, the system closes detections without activity after 7 days. An admin user can configure the autoclose time period.

To configure the autoclose time period for detections

1. From **System Settings** in the left navigation, choose **General Settings**.
2. In the Detection Settings section, click the **Edit** () icon for the Detection Auto Close Time.
3. Specify the number of days to wait before autoclosing detections and click **Save**.

You can specify a time period from 7 to 90 days.

Once you customize the autoclose period, you can click the **Reset** icon to restore the default period of 7 days.

The system tracks changes to this setting and records when this setting changes and who made the change.

System Settings

10.3 Alert Catalog

The Alert Catalog is a list of all alert categories in the system. From the Alert Catalog, you can view details about each alert category, customize severity scores, and define entities to exclude from a category. The catalog is available to both analysts and admins.

To access the Alert Catalog

- From **System Settings** in the left navigation, choose **Alert Catalog**.

The Alert Catalog appears as a tab in the System Settings.

The screenshot shows the 'Alert Catalog' interface. At the top, there is a search bar labeled 'Security Alert Name' with a 'Search' button. Below the search bar, there are three filter dropdowns: 'Status: No value set', 'Score: No value set', and 'Type: No value set'. The main content is a table with 10 rows. The table has columns for 'Status', 'Severity', 'Security Alert Name', 'Type', and 'Actions'. The 'Status' column contains toggle switches, 'Severity' contains scores (9 or 8), 'Security Alert Name' contains the alert category names, 'Type' contains the alert types, and 'Actions' contains edit and toggle icons. At the bottom of the table, there is a pagination control showing '10' items per page, '51-60 of 60717 items', and 'Page 6 of 6072'.

Status	Severity	Security Alert Name	Type	Actions
<input checked="" type="checkbox"/>	9	DGA Malware	Machine Learning	
<input type="checkbox"/>	9	SSH::Login_By_Password_Guesser	Notice	
<input checked="" type="checkbox"/>	9	ICMPSpecificTunnelDetectors::icmptx::ICMP_icmptx_Tunnel	Notice	
<input checked="" type="checkbox"/>	9	ICMPSpecificTunnelDetectors::itun::ICMP_itun_Tunnel	Notice	
<input checked="" type="checkbox"/>	9	CVE_2021_44228::LOG4J_ATTEMPT_HEADER	Notice	
<input checked="" type="checkbox"/>	9	CVE_2021_44228::LOG4J_JAVA_CLASS_DOWNLOAD	Notice	
<input checked="" type="checkbox"/>	9	EternalSafety::EternalSynergy	Notice	
<input checked="" type="checkbox"/>	9	ICMPSpecificTunnelDetectors::hans::ICMP_Hans_Tunnel	Notice	
<input checked="" type="checkbox"/>	8	ETPRO MALWARE Win32/Fadok.A Checkin	Suricata	
<input checked="" type="checkbox"/>	8	ET MOBILE_MALWARE Android.YzhcSms CnC Keepalive Message	Suricata	

For each alert category in the catalog, you can view:



- Status** - Active or Inactive. Active alerts are enabled and visible in the dashboard. Inactive alerts are not enabled and are not visible in the dashboard. Click the toggle button to change the status. If an alert is inactive, it is still available in Falcon LogScale (Humio).

You can select multiple alert categories and change their status with a single click.

- Severity** - A score ranging from 1 to 10 with more severe threats having a higher score. Investigator normalizes scores across alert types.
- Security Alert Name** - Click an alert category name for the Alert Category details page, which includes:
 - **Info** - a description of the alert category (if available)

- MITRE ATT&CK mapping to tactics and technique (if available)
- Alert properties, including a custom severity score
- Excluded entities, which are trusted entities you exclude from the alert category based on IP, domain, or CIDR.
- Type - Machine Learning, Notice, Search Based, or Suricata.

Each alert category includes actions. Actions include:

-  **Edit Alert** opens a dialog box that lets you edit the severity score for the alert.
-  **View all detections** lets you pivot from the category to the Detections page to get more information about all detections associated with the alert category and to take the next troubleshooting steps.

You can filter alert categories based on these values:

- Status - Show only Active or Inactive alerts.
- Score - Show only low, medium, or high risk alerts based on their severity score.
- Type - Show only machine learning, Suricata, search based, or notice alerts.

You can select multiple filter values from each filter menu.

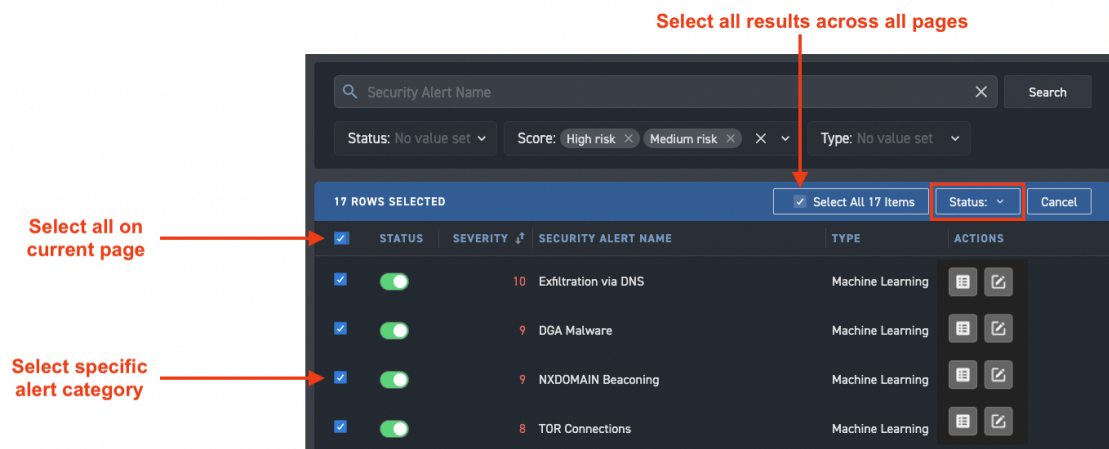
You can also search for a specific alert category.

Click an entry in the Alert Catalog to view additional details and properties.

Note: When Investigator identifies a Notice and Suricata alert not in the Alert Catalog, Investigator adds the alert to the catalog with an inactive status. You can change the status for this alert to appear in your dashboard.

10.3.1 Update multiple alert categories at the same time



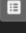
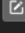

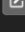
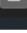
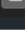
You can select multiple alert categories and update their status with a single click. You can select specific entries, all the entries on the page, or all entries across multiple pages. When you select an alert category, a banner appears that provides selection information and options.



Select all results across all pages

Select all on current page

Select specific alert category

STATUS	SEVERITY	SECURITY ALERT NAME	TYPE	ACTIONS
<input checked="" type="checkbox"/>	10	Exfiltration via DNS	Machine Learning	 
<input checked="" type="checkbox"/>	9	DGA Malware	Machine Learning	 
<input checked="" type="checkbox"/>	9	NXDOMAIN Beaconing	Machine Learning	 
<input checked="" type="checkbox"/>	8	TOR Connections	Machine Learning	 

With multiple alert categories selected, use the **Status** menu to change the status of the selected entries to Active or Inactive.

Tip: Use search and filters to limit the results to make more focused selections.

10.3.2 Exclude entities from an alert category

You can exclude entities from an alert category in the Alert Catalog. Excluded entities do not generate new alerts for the alert category. Typically, you exclude trusted entities so you can focus on other entities.

To exclude an entity from an alert category

1. From the Alert Catalog, click the name of the alert to display the details page.
2. In the Excluded Entities section, click **+ Add New**.

The right side panel opens.

3. Choose an **Entity Type** and enter the corresponding value to identify the entity.

You can exclude Notice and Suricata alerts based on IP address or CIDR. Exclusions for Machine Learning alerts depend on the nature of alert: some Machine Learning alerts let you exclude based on the IP or CIDR and some (such as DNS alerts) let you exclude based on the domain. Similarly, excluding entities from Search Based alerts depends on the nature of alert.

4. Click **Save**.

Once saved, specified entities do not generate new alerts for the alert category. Investigator tracks and displays the user who excluded the entity and the time/date the entity was added to the exclusion list.

Click the **Delete** icon in the Actions column to remove excluded entities.

You can also add an entity to the Excluded Entities list through the **Suppress Entity** button on the Detections, Entity detail, and Alert detail pages. Similarly, you can use the **Unsuppress Entity** button to remove an entity from the Excluded Entities list from the Entity detail and Alert detail pages.

10.3.3 Customize severity scores

Investigator lets you change the severity score of an alert category. A custom score overrides the default score assigned by Corelight for alert categories and sets the priority for alert categories. You can revert the custom severity score to the Corelight assigned score. System changes do not override a custom score to an alert category.

By changing the severity score of an alert, you can prioritize the incidents that need immediate attention over those that are less severe. This helps analysts allocate their time and resources more efficiently and address the most critical issues first.

You can also change a security score to reduce false positives and false negatives. An alert might be triggered due to a false positive, which means that the alert is not indicative of a security incident. In such cases, changing the severity score of the alert can downplay the significance of the alert and prevent it from taking up analysts time.

Conversely, there may be situations where an alert is triggered, but the severity score is too low to warrant immediate attention. In such cases, changing the severity score of the alert to a higher value can help to ensure that the incident receives appropriate attention.

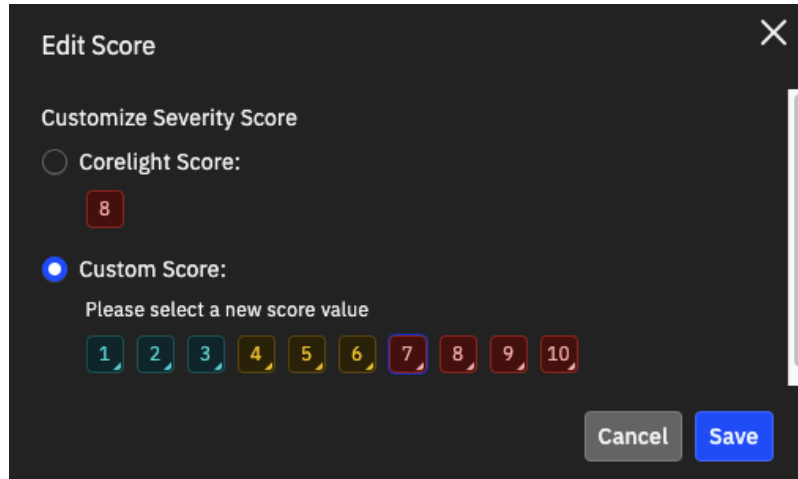
Additionally, changing the severity score of an alert can be part of a broader effort to adjust detection thresholds. By changing the threshold values for alerts, analysts can fine-tune the sensitivity of the alerting system to better match the organization's risk profile and incident response capabilities.

You can update the severity score from the Alert Catalog.

To customize the severity score for an alert category

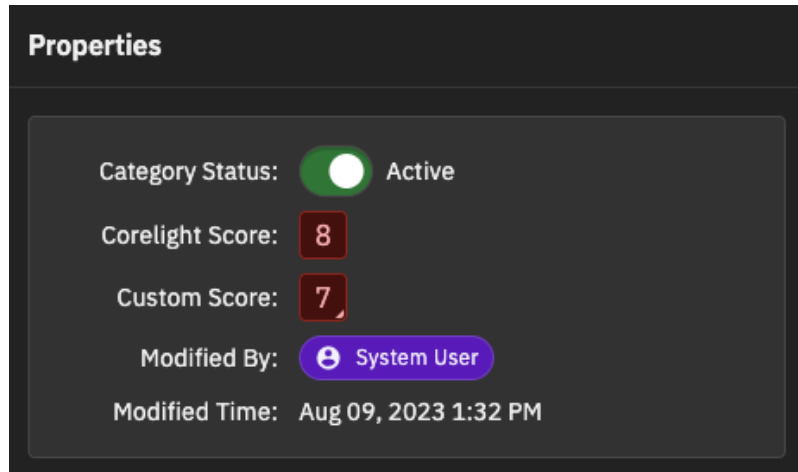
1. From the Alert Catalog, click the name of the alert category to display the details page.
2. Click the **Edit Score** button under the alert category name.

An Edit Score dialog box appears.



3. Click **Custom Score** and click a new score value with 10 being the most critical and 1 being the least important.
4. Click **Save**.

The Properties panel on the alert category details page indicates the custom score for the alert category, and the score icon for an alert category with a custom score changes to include a triangle in the lower-right corner.



When the custom score is set, all alerts generated in the Alert Category will have the custom score.

Changes to the severity score are captured in the Security Audit. The user who modified the alert properties is also identified with the custom severity score.

You can use the **Edit Score** button to revert to the default Corelight score or adjust the custom score at any time.

10.3.4 Search the Alert Catalog

You can search the Alert Catalog based on security alert names. (The search applies to alert titles only and does not include fields and descriptive content.)

The search uses [Kibana query language syntax](#).

Follow these guidelines when creating search queries:

- A query can consist of one or more words or a phrase. A phrase is a group of words surrounded by double quotation marks, such as “test search”. Use a phrase to match an exact string.
- Combine multiple sub-queries with AND and OR operators. Operators such as AND and OR must be capitalized.
- Use parentheses to group sub-queries.
- Use wildcards (*) to match multiple values.
- Use [regular expression \(regex\)](#) to form complex string queries.

Here are some example search terms:

- malware – finds a single phrase
- “et dns” – finds the exact phrase in the double quotes
- et AND dns – finds results with both terms
- et OR dns – finds results with either one of the terms
- (ETPRO MALWARE) AND Payloads – finds results with a compound search term and a single search term
- nim AND *2 – finds results with the literal phrase (nim) and any other term that ends with a 2
- /v1.[0-9]/ – finds results with a literal string (v1.) followed by a number matching a range of numbers from 0 to 9
- /mal[a-z]*re/ – finds results with a literal string (mal) followed by a letter matching a range of lowercase letters and then a wildcard in the middle adding flexibility and ending with another literal string (re)
- /mal[a-z]*/ – finds a term with a range of letters and a wildcard at the end to allow more options
- nim AND /. *2/ – finds the first literal term and the results of the regex construct that finds a term that ends with a 2

10.4 Audit activities through logs

Corelight Investigator provides a record of key user and application activities. Activities are recorded in the Security Audit log and admin users can review them. Security Audit logs help diagnose problems, address security concerns, and comply with regulatory requirements.

Security Audit

Security Audit

User: No value set | Category: No value set | Type: No value set | Last 7 days

Timestamp (UTC)	Author	Category	Type	Activity
Oct 5, 2023 7:12 PM	omesh.kumar+staginganalyst@c...	User	Audit	omesh.kumar+staginganalyst@c...
Oct 5, 2023 6:52 PM	kris.leonard+stagingone@corell...	User	Audit	kris.leonard+stagingone@corell...
Oct 5, 2023 6:51 PM	eric.akin+stagingone@corelight...	Admin	Audit	License viewed
Oct 5, 2023 6:50 PM	omesh.kumar+staginganalyst@c...	User	Audit	omesh.kumar+staginganalyst@c...
Oct 5, 2023 6:50 PM	omesh.kumar+staginganalyst@c...	User	Audit	omesh.kumar+staginganalyst@c...
Oct 5, 2023 6:08 PM	omesh.kumar+staginganalyst@c...	User	Audit	omesh.kumar+staginganalyst@c...
Oct 5, 2023 5:55 PM	eric.akin+stagingone@corelight...	Escalation	Audit	
Oct 5, 2023 5:55 PM	eric.akin+stagingone@corelight...	Admin	Audit	License viewed
Oct 5, 2023 5:55 PM	eric.akin+stagingone@corelight...	Escalation	Audit	
Oct 5, 2023 5:54 PM	eric.akin+stagingone@corelight...	User	Audit	eric.akin+stagingone@corelight...

10 | 1-10 of 370 items | Page 1 of 37

Log entries are grouped in these categories:

- **User** – Actions taken by users such as logins, logouts, password changes, alias updates, 2FA setup and modifications, and cookie settings.
- **Admin** – Actions taken by administrators, such as user creation, permission changes, and configuration changes.
- **System** – Events related to the system itself, such as system errors or warnings, scheduled updates, and Corelight support activity.
- **Alert** – Events or user actions on Alerts such as enabling or disabling an Alert Category.
- **Detection** – Events or user actions on Detections such as assigning or closing a detection.
- **Entity** – Events or user actions on entities, such as adding an entity to an exclusion list.
- **Export** – Events related to Exporters such as creating, updating, or deleting.

Entries in the audit log are read-only and cannot be modified or deleted.

10.4.1 View the audit log

To access the audit log:

- As an admin user, from the **System Settings** in the left navigation, choose **Security Audit**.

The Security Audit page shows logs of recent activity, including a timestamp, user who performed the action, activity category, type of activity, and activity description.

You can filter the logs based on these values:

- **User** – Choose a specific user to limit the results to their activity or choose to view all users. (All users associated with the account appear in the list.)
- **Category** – Choose a category for the results.
- **Type** – Choose Audit to show system activity or Error to show failed requests.

- Date – Specify a time range. The window can range from 1 hour to 3 months. You can also specify a custom date range. The custom range cannot exceed 3 months.

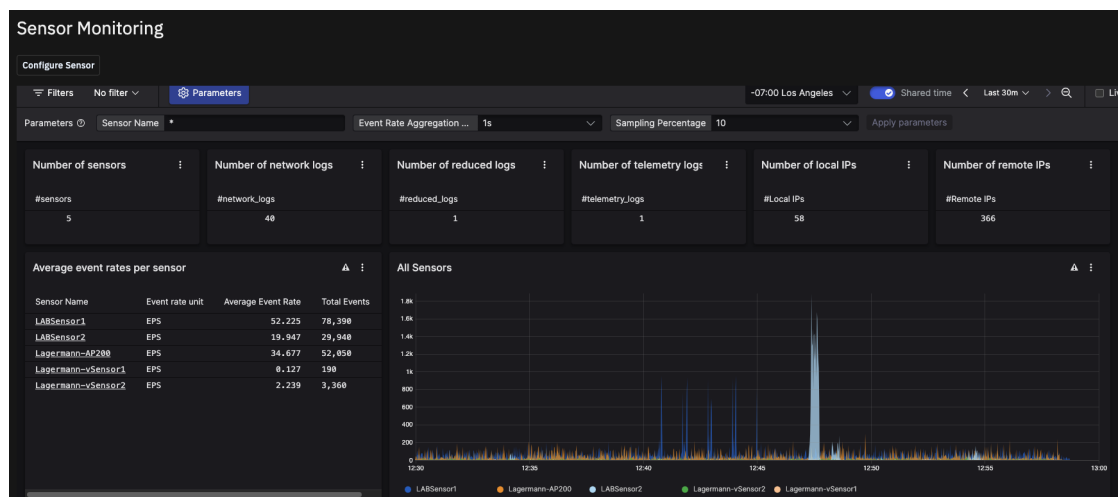
Click the **X** to the right of a filter name to remove it.

You can also sort logs based on timestamp, user, category, or type.

10.5 Sensor monitoring and management

The Sensor Monitoring page guides you through sensor configuration and provides a health dashboard for connected sensors. You can see how many sensors are exporting data to Investigator and view detailed information about the status and the nature of the exports. The dashboard also provides search and filtering options so you can pinpoint problems with your sensor connections.

To access sensor monitoring, click **System Settings** in the left navigation and choose **Sensor Monitoring** from the menu.



If you are an admin, you can configure new sensors from this page. Click the **Configure Sensor** button and follow the instructions in the Quickstart topic *Configure your sensors*.

To remove a sensor from Investigator, turn off the export on the sensor.

10.6 User management

Admin users can manage users and authentication options from **System Settings | Users & Access**.

Investigator provides these authentication options:

10.6.1 Local user management

You manage users and choose authentication methods from the **Users & Access** page, available from **System Settings** in the left navigation.

By default, Investigator provides local authentication for users. You can use the Access tab to switch to *SSO SAML authentication*.

Admin users can view, create, and manage local users. The User tab on the Users & Access page displays information about currently provisioned users including their alias, role, and status.

A user's status can be:

- **Invited.** An invitation email was sent, but the user hasn't completed their account registration. Temporary credentials expire in 7 days.
- **Active.** The user accepted the invitation and completed their account registration. An admin can activate a user account at any time.
- **Inactive.** The user account is suspended. While their account is inactive, the user can't access Corelight Investigator but their account settings and data are preserved. An admin can inactivate a user account at any time.

To add a new user

1. From **System Settings** in the left navigation, choose **Users & Access**.

The User tab appears.

2. Click **+ Add User**.

3. Provide these user details.

- **Alias.** The alias is not a username, but essentially a nickname for the account. This value has a limit of 30 characters. The user can change this alias later in their account settings, but the admin can't edit it once created.
- **Email.** The login name. The initial confirmation email with account access details is sent to this address. The email address can't be changed once the user is created. To change the email, the admin must delete the user and recreate their account with the new address.
- **Role.** Specify if this user is an analyst, if they can perform admin tasks like user management and system configuration, or if they can simply be a viewer of detections without taking action.

4. If you don't want this user to have immediate access, such as a new hire starting at a later date, select **Create User As Inactive**.

You can edit this user information and change the state to active at any time. The user receives their welcome and account confirmation emails as soon as they're marked as active.

5. Click **Create**.

To edit user details

1. For the user entry, click the edit icon () in the Actions column.

An admin can change a user's role and status but can't edit the alias or the email. The user can change their own alias in their account settings. To change the account email, the admin must delete the user and recreate their account with the new address.

Note: Email addresses are obscured in the full user list, but if you edit the user details you can view the configured value.

2. Click **Save**.

When a user role or status changes, the system sends an email to notify the user.

When a user role changes, the user is forced to log out and log in again so the system properly recognizes the new role.

Use the checkboxes to select and change the status of multiple users with a single click.

To delete a user

1. Select the user entry and click **Delete**.
2. Confirm the action.

You can select multiple users to delete more than one at a time.

10.6.2 SAML SSO user management

Investigator provides cross-domain single-sign on (SSO) using Security Assertion Markup Language (SAML) 2.0.

SAML 2.0 is an XML standard that acts as an authentication interface between Investigator and an identity provider (IdP) that manages user credentials. When Investigator receives a login request, it determines if SAML is enabled. If SAML is enabled, Investigator redirects the authentication request to the IdP.

Investigator SSO is compliant with SAML 2.0 systems and has been verified with Okta and Auth0.

Admins can set up SAML in the **Access** tab on the **Users & Access** page.

With SSO enabled, local authentication remains available for all admin users. For analyst users, Investigator does not allow local authentication from the same domain as the SSO, however, analyst users not in the SSO domain can log in.

10.6.2.1 Configure SAML SSO

The Investigator interface provides a wizard to step account admins through the SAML SSO configuration.

To configure SAML SSO

1. From **System Settings** in the left navigation, choose **Users & Access** and click the **Access** tab.
2. Click **Enable** for the **SAML Single Sign On** access option.

The configuration wizard displays two configuration values for your IdP.

3. Copy the **Assertion Consumer Service URL** and the **Remote Manager Audience URL** and add these values to the SAML configuration on your IdP.

These values are specific for your account.

4. From your IdP, ensure the SAML 2.0 assertion contains the `email_address` and `group` attributes. Make sure that the attribute names are exactly as specified (and ensure the IdP does not prepend the namespace to the attribute name).

Define the group attribute to be either `corelight_admin`, `corelight_analyst`, or `corelight_viewer`, which align with the Admin, Analyst, and Viewer roles in Investigator.

For example:

```
<saml2:Attribute Name="email_address">
  <saml2:AttributeValue xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:type="xs:anyType">investigator.user@corelight.com</
saml2:AttributeValue>
</saml2:Attribute>
<saml2:Attribute Name="group">
  <saml2:AttributeValue xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:type="xs:anyType">corelight_admin</saml2:AttributeValue>
</saml2:Attribute>
```

5. From your IdP, find and copy the Entity ID (identity provider), single sign on service URL, and X.509 certificate.

As another approach, you can download a configuration file from your IdP and upload it to Investigator to provide the same information.

6. In Investigator, click **Next** and enter the Entity ID, single sign on service URL, and certificate or upload the configuration file.
7. Click **Check Connection** to ensure Investigator can successfully communicate with your IdP.
8. Click **Next**.
9. Map the Analyst and Admin roles in Investigator to roles in your IdP.

Analyst, Admin, and Viewer roles must be defined in your IdP as `corelight_analyst`, `corelight_admin`, and `corelight_viewer` and associated with SSO users.

10. Click **Enable**.

A message indicates SAML Single Sign On is enabled.

With SAML enabled, you add and manage new users through your IdP. The User tab in Investigator is read-only. (Local authentication users remain in the system but only users not in the SSO domain remain as active.)

If a local authentication email address matches an SSO email, Investigator maintains any defined user preferences.

SAML users log in to Investigator using the **Sign In SSO** button. Users need to provide their email domain to redirect to their IdP for authentication.

When configured, you can use the SAML configuration wizard from the **Access** tab to modify the configuration.

10.6.2.2 Delete SAML SSO configuration

You can delete the SAML configuration by enabling Local Authentication.

When you switch to Local Authentication, all SSO user information is disabled and local user accounts are enabled. You will need to add local user accounts again when you disable SAML. If the user email matches a previously configured user, their profile settings are restored. See *Local user management* for more information.

Integrations

10.7 GPT settings

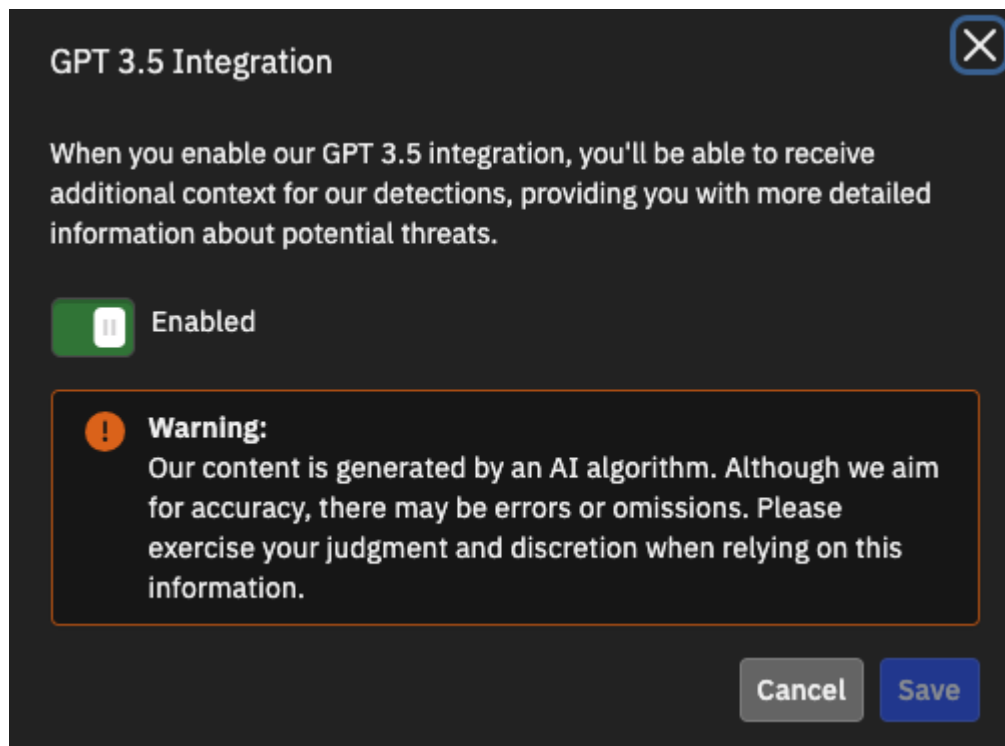
Investigator uses GPT 3.5 to provide additional context and details for detections. An AI algorithm generates this content and there might be errors or omissions. We recommend using your best judgement with AI generated content. (GPT generated content is identified in the interface an icon.)

By default, GPT integration is enabled. Admin users can turn off GPT content through the system settings.

Note: To configure this integration, you need to have admin access. (Analyst users can view the integration but cannot make changes.)

To disable or enable GPT integration

1. From **System Settings** in the left navigation, choose **Integrations**.
2. In the Integrations tab, click the GPT card.
3. Toggle the GPT 3.5 Integration value to **Enabled** or **Disabled** and click **Save**.



10.8 ServiceNow integration

Investigator lets analysts send detections to ServiceNow for further analysis and case management. To enable this functionality, you need to configure the integration in the settings. The integration sets up a REST API connection between Investigator and ServiceNow.

To configure the integration, you need to have a ServiceNow license and Investigator admin access. (Analyst users can view the integration but cannot make changes.)

Important: Once configured, users send detections manually on a case-by-case basis.

Before you configure the integration in Investigator, you need to create a table in your ServiceNow instance to manage detections. When configuring the integration, you specify the table name to ensure this integration is associated with the correct table.

To create a table your ServiceNow instance

1. From your app dashboard page in ServiceNow, create a table with these fields for the Investigator detection details.

Column Label	Column Name	Type	Max Length
alert_category	alert_category	String	500
score	score	Integer	
description	description	String	1000
detection_url	detection_url	URL	
detection_status	detection_status	String	10
entity	entity	String	40
entity_type	entity_type	String	40
detection_created_at	detection_created_at	UTC Time	
detection_updated_at	detection_updated_at	UTC Time	
no_of_alerts	no_of_alerts	Integer	
assignee	assignee	String	500

For more information, see [Create a table](#) in the ServiceNow documentation.

2. Note the table name and prefix. You will need to provide this value in the Investigator settings.

To integrate your ServiceNow instance with Investigator

1. From **System Settings** in the left navigation, choose **Integrations**.
2. In the Integrations tab, click the **ServiceNow** card.

An integration dialog box appears.

3. Toggle the integration value to **Enabled**.
4. Enter your ServiceNow credentials, including username, password, and instance name.

These values are available in ServiceNow in My Instance > Manage Instance Password.

The ServiceNow instance name (or instance ID) is a unique identifier for your ServiceNow instance. It's important to provide the correct value to ensure that actions and data are associated with the correct ServiceNow environment.

For details, see the [ServiceNow documentation](#).

5. Provide the table name (including the prefix) for the table you created with the detection-specific fields.
6. Click the **Verify Connection** button to ensure Investigator can communicate with your ServiceNow instance.

You cannot save your connection until you verify it.

7. Click **Save**.

With ServiceNow configured and enabled, analysts can manually send individual detections to ServiceNow.

If you want to pause the integration, toggle the integration setting to **Disabled**. This preserves your connection details.

If you want to disable the integration and delete your connection details, click the **Delete** icon.

10.8.1 Additional learning resources

Watch a Corelight video on YouTube: [How Corelight's ServiceNow integration speeds response](#)

10.9 CrowdStrike EDR integration

With a CrowdStrike EDR (Endpoint Detection and Response) integration, you can receive entity enrichment for Investigator detections.

The CrowdStrike integration seamlessly blends CrowdStrike EDR with Investigator network detection capabilities and maps Corelight IP addresses from detections to CrowdStrike's host information. When integrated, detections show valuable entity context within the detection details. The expanded data provides additional context to analyze threats and helps analysts make informed decisions during the triage process.

To configure the integration, you need to have CrowdStrike OAuth2 API access with read access to the managed device information. Within Investigator, you need admin access. (Analyst users can view the integration but cannot make changes.)

If you want to enable the CrowdStrike Network Containment integration, you need read and write API access.

To integrate CrowdStrike with Investigator

1. From **System Settings** in the left navigation, choose **Integrations**.
2. In the Integrations tab, click the **CrowdStrike** card.
An integration dialog box appears.
3. Toggle the integration value to **Enabled**.
4. If you want to allow admin users to isolate entities from the network, toggle the **Isolate Entity** slider to **Enabled**.

For hosts to enable Network Containment integration, your CrowdStrike API client needs write permissions so Investigator can enable the CrowdStrike Falcon Network Containment feature.

5. For the **URL**, enter the API URL for your region.

US-1: <https://api.crowdstrike.com>

US-2: <https://api.us-2.crowdstrike.com>

EU-1: <https://api.eu-1.crowdstrike.com>

US-GOV-1: <https://api.laggar.gcw.crowdstrike.com>

US-GOV-2: <https://api.us-gov-2.crowdstrike.mil>

6. Enter your CrowdStrike Client ID and Client Secret Key.

These values are available from the CrowdStrike Falcon Console. For more information on creating and obtaining these values from the CrowdStrike Falcon Console, see the documentation within the Falcon console or [Getting Access to the CrowdStrike API](#).

7. Specify the **Polling Time** in minutes.

Use this field to customize the interval for updates to the EDR data. The minimum interval is 5 minutes.

8. Click **Verify Connection** to ensure Investigator can access your CrowdStrike data.

You cannot save your connection until you verify it.

9. Click **Save**.

With CrowdStrike configured and enabled, detection details show *detailed entity information* and indicate that CrowdStrike is the source of the content.

If you want to pause the integration, toggle the integration setting to **Disabled**. This preserves your connection details.

If you want to disable the integration and delete your connection details, click the **Delete** icon.

10.10 Alert Exports

As an admin user, the Alert Exports integrations page lets you set up alert exports to CrowdStrike Falcon LogScale, Elastic, and Splunk HTTP Event Collector (HEC) as well as a generic HTTP exporter. Admins can configure one or more alert exports, and can export to multiple instances of the same type of exporter.

Analyst users can view the alert exports but cannot make changes.

Exported alerts include a URL to the detection details page within Investigator.

Note: If needed, update your firewall rules to allow alerts from Investigator. The source IP address for the alerts is fixed for a region: the North America (us-west-2) IP address is 35.81.184.144 and the EU IP address is 35.157.240.249. For endpoint details, see the AWS help topic [Amazon Kinesis Data Streams endpoints and quotas - AWS General Reference](#).

Once configured, you can click the name of an exporter to open a side panel where you can edit, disable, or delete it.

Important: If you already configured log export on one or more sensors, exporting alerts from Investigator duplicates some notice and Suricata alerts.

10.10.1 CrowdStrike Falcon LogScale

To configure alert export to Falcon LogScale

1. From **System Settings** in the left navigation, choose **Integrations** and click the **Alert Exports** tab.
2. Click **CrowdStrike Falcon LogScale**.
3. Provide a name for the exporter.
4. Toggle on **Enabled**.
5. Provide information about your instance, including the URL, the token, and the LogScale index.
6. Optionally, enable **Verify SSL** to verify the certificate and hostname provided by the exporter.

If you enable this option, the certificates must be current and not expired and must be issued by a trusted issuer. Additionally, the hostname used to connect to the remote host must be present in the TLS certificate presented by the remote host, either as the Common Name or as an entry in the Subject Alternative Name extension.

7. Click **Save**.

For more information, see the [Falcon LogScale documentation](#).

10.10.2 Elastic

To configure alert export to Elastic

1. From **System Settings** in the left navigation, choose **Integrations** and click the **Alert Exports** tab.
2. Click **Elastic**.
3. Provide a name for the exporter.
4. Toggle on **Enabled**.
5. Provide information about your Elastic instance, including the URL, the username and password, and the index.
6. Optionally, enable **Verify SSL** to verify the certificate and hostname provided by the exporter.

If you enable this option, the certificates must be current and not expired and must be issued by a trusted issuer. Additionally, the hostname used to connect to the remote host must be present in the TLS certificate presented by the remote host, either as the Common Name or as an entry in the Subject Alternative Name extension.

7. Click **Save**.

For more information about Elastic HTTP exporter, see the [Elastic documentation](#).

10.10.3 Splunk HEC

To configure alert export to Splunk HEC

1. From **System Settings** in the left navigation, choose **Integrations** and click the **Alert Exports** tab.
2. Click **Splunk HEC**.
3. Provide a name for the exporter.
4. Toggle on **Enabled**.
5. Provide information about your HEC instance, including the URL, the token, and the Splunk index.

6. Optionally, enable **Verify SSL** to verify the certificate and hostname provided by the exporter.

If you enable this option, the certificates must be current and not expired and must be issued by a trusted issuer. Additionally, the hostname used to connect to the remote host must be present in the TLS certificate presented by the remote host, either as the Common Name or as an entry in the Subject Alternative Name extension.

7. Click **Save**.

For more information about Splunk HEC, see the [Splunk documentation](#).

10.10.4 Generic HTTP exporter

To configure alert export to a generic HTTP exporter

1. From **System Settings** in the left navigation, choose **Integrations** and click the **Alert Exports** tab.
2. Click **HTTP Exporter**.
3. Provide a name for the exporter.
4. Toggle on **Enabled**.
5. Provide the URL for your exporter.
6. Specify the authentication type.
 - None.
 - Basic. Provide a username and password for authentication.
 - Bearer. Provide a token value for authentication.

7. Optionally, add custom headers in the key/value format.

You can add one or more custom headers to export additional information. Each header must have a key (without spaces) and a corresponding value.

8. Optionally, enable **Verify SSL** to verify the certificate and hostname provided by the exporter.

If you enable this option, the certificates must be current and not expired and must be issued by a trusted issuer. Additionally, the hostname used to connect to the remote host must be present in the TLS certificate presented by the remote host, either as the Common Name or as an entry in the Subject Alternative Name extension.

9. Click **Save**.

Note: Users with an Analyst role can see the Integrations page in view-only mode and can see available integrations, their status, and alert exports.
