

# Exabeam Advanced Analytics Administration Guide

---

Exabeam SOC Platform - Cloud-Delivered Release Only

August 6, 2024

**Exabeam**

1051 E. Hillsdale Blvd, 4th Floor  
Foster City, CA 94404

650.209.8599

Have feedback on this guide? We'd love to hear from you!  
Email us at [docs@exabeam.com](mailto:docs@exabeam.com)

**Disclaimer:** Ensure that you are viewing the most up-to-date version  
of this guide by visiting the Exabeam Documentation Portal.

## **Copyright**

All content in this document, including text, graphics, logos, icons, images, and video clips, is the exclusive property of Exabeam or its content suppliers and is protected by U.S. and international copyright laws. The compilation (meaning the collection, arrangement, and assembly) of all content in this document is the exclusive property of Exabeam and is also protected by U.S. and international copyright laws. The content in this document may be used as a resource. Any other use, including the reproduction, modification, distribution, transmission, republication, display, or performance, of the content in this document is strictly prohibited.

Copyright ©2024 Exabeam, Inc. All Rights Reserved.

## **Trademarks**

Exabeam and the Exabeam logo are trademarks or registered trademarks of Exabeam. All third-party trademarks mentioned in this document are the property of their respective owners. The trademarks and logos displayed in this document may not be used without the prior written consent of Exabeam or their respective owners.

## **Products**

Exabeam owns and retains all right, title and interest in and to the Exabeam's products and services and portions, features and/or functionality of Exabeam's products may be protected under Exabeam's patents, as well as patents pending.

## **Other Policies**

For information regarding Exabeam's treatment of personally identifiable information, please review Exabeam's current privacy policy at [www.exabeam.com/privacy](http://www.exabeam.com/privacy).

## Table of Contents

Advanced Analytics .....	6
Understand the Basics of Advanced Analytics .....	7
Advanced Analytics Overview .....	7
How Advanced Analytics Works .....	7
Fetch and Ingest Logs .....	8
Add Context .....	9
Detect Anomalies .....	9
Assess Risk .....	9
Data Flow Diagrams .....	9
Data Retention in Advanced Analytics .....	10
Configure Log Management .....	12
Log Ingestion Settings .....	12
i63 and Later .....	12
i60 to i62 .....	12
View Insights About Syslog-Ingested Logs .....	13
Ingest Logs from Google Cloud Pub/Sub into Advanced Analytics .....	14
Log Feeds .....	15
Set up a Log Feed .....	15
Training and Scoring .....	15
Notifications .....	15
Set Up Notifications to a Log Repository, Ticketing System, or SIEM .....	16
Set Up Notifications to Email .....	17
Syslog Notifications Key-Value Pair Definitions .....	18
Audit Logs .....	22
Set Up Admin Operations .....	24
Exabeam Licenses .....	24
License Lifecycle .....	24
Types of Exabeam Product Licenses .....	25
View Version Information .....	26
Cluster Authorization Token .....	26
Exabeam Engines .....	28
i63 and Later .....	28
i60 to i62 .....	29
Generate a Support File .....	29
Set Up Authentication and Access Control .....	30
What Are Accounts & Groups? .....	30
What Are Assets & Networks? .....	30
Universal Role-Based Access .....	31
Migrate to Universal Role-Based Access .....	31
Legacy Role-Based Access Control .....	36
Role-Based Access Control .....	36
Manage Users .....	36
Third-Party Identity Provider Configuration .....	45
Set Up LDAP Server .....	71
Set Up LDAP Authentication .....	73

Azure AD Context Enrichment .....	73
Set Up Azure AD Context Enrichment .....	74
Set Up Context Management .....	80
Out-of-the-Box Context Tables .....	80
Azure Active Directory Context Tables .....	82
Threat Intelligence Service Context Tables .....	82
Custom Context Tables .....	82
Prepare Context Data .....	83
Create Custom Lookups .....	84
Create a Context Table .....	84
Import Data into a Context Table Using CSV .....	85
Import Data into a Context Table Using an LDAP Connection .....	87
Mask Data Within the Advanced Analytics UI .....	91
Mask Data for Notifications .....	92
Additional Configurations .....	94
Display a Custom Login Message .....	94
Reprocess Jobs .....	96
Reparse Raw Logs to Create New Events .....	96
Run the Analytics Engine to Reprocess Log Feeds .....	97
Configure Job Status Notifications .....	98
User Engagement Analytics Policy .....	98
Configure Rules .....	100
What Is an Exabeam Rule? .....	100
How Exabeam Models Work .....	101
Model Aging .....	101
Rule Naming Convention .....	101
View Rules in Advanced Analytics .....	102
Filter Rules .....	102
Disable or Enable a Rule .....	103
Create a Fact-Based Rule .....	103
Edit a Rule .....	105
Edit a Rule Using the Advanced Editor .....	106
Edit a Fact-Based Rule Using the Simple Editor .....	106
Clone a Rule .....	107
Reprocess a Rule .....	107
Revert Out-of-the-Box Rules .....	107
Exabeam Threat Intelligence Service .....	109
Threat Intelligence Service Prerequisites .....	111
View Threat Intelligence Feeds .....	111
Threat Intelligence Context Tables .....	112
View Threat Intelligence Context Tables .....	113
Assign a Threat Intelligence Feed to a New Context Table .....	114
Individual Feed Assignment .....	114
Bulk Feed Assignment .....	115
Create a New Context Table from a Threat Intelligence Feed .....	116
Create a Table from a Single Feed .....	116
Create a Table from Multiple Feeds .....	117

Check ExaCloud Connector Service Health Status .....	118
Exabeam Cloud Telemetry Service .....	120
Manage Security Content in Advanced Analytics .....	121
Manually Install a Content Package .....	121
Automatically Install Content Packages .....	122
Manually Check for New Content Packages .....	122
Automatically Check for New Content Packages .....	123
Upload a Content Package .....	123
Uninstall a Custom Content Package .....	124
Health Status Page .....	125
Proactive and On-Demand System Health Checks .....	125
Advanced Analytics Specific Health Checks .....	126
Configure Alerts for Worker Node Lag .....	127
System Health Checks .....	128
Alerts for Storage Use .....	128
Default Data Retention Settings .....	129
System Health Alerts for Low Disk Space .....	129
System Health Alerts for Paused Parsers .....	129
View Storage Usage and Retention Settings .....	130
Set Up System Optimization .....	130
Disabled Models .....	131
Paused Parsers .....	131
Disabled Event Types .....	133
Automatically Redistribute System Load .....	133
Automatic Shutdown .....	134

## Advanced Analytics

Exabeam collects and processes data from a log management system and other external context data sources in order to identify advanced security attacks.

Advanced Analytics can identify compromised, malicious insiders, and advanced threats by leveraging logs and contextual information. High risk behaviors in your organization are tracked across networks and assets then articulated into comprehensive timelines so you can focus your investigation from point to point of action rather than manually amassing and sifting data.

## Understand the Basics of Advanced Analytics

This section provides information about the components of Advanced Analytics, and how they work together. Some aspects of the components vary depending on which version of Advanced Analytics you are using. These differences are noted throughout the guide as follows:

- **i63 and later** – This designation applies to the latest SaaS version of Advanced Analytics. In this version:
  - The Log Ingestion and Messaging Engine (LIME) has been replaced by a unified ingestion pipeline (UIP) that centralizes log ingestion activities for all Exabeam products. Visibility into the unified ingestion pipeline is provided in the cloud-native [Log Stream](#) functionality.
  - Log parsing, field compliance, and event building are all based on a hierarchical [common information model](#) that informs the data structure for all Exabeam products.
  - Search and presentation functionality that used to be performed in Data Lake are now available in cloud-native [Search](#) and [Dashboard](#) applications.
- **i60 to i62** – This designation applies to the legacy SaaS version of Advanced Analytics. In this version:
  - LIME continues to handle log ingestion functionality.
  - The information model in use is specific to Advanced Analytics only.
  - Data Lake still provides log management, search, and presentation functionality.

### Advanced Analytics Overview

Advanced Analytics provides user and entity behavior intelligence on top of existing SIEM and log management data repositories. Advanced Analytics can detect compromised and rogue insiders and can present a complete picture of both the user session and lateral movement use within the attack chain.

Exabeam pulls logs from a variety of data sources and enriches this data with identity information collected from Active Directory (LDAP). This information provides an identity context for credential use. Through behavior modeling and analytics, Advanced Analytics learns normal user credential activities and access characteristics. By automatically comparing incoming data to these normal behaviors, anomalous activity can be exposed.

Advanced Analytics places all user credential activities and characteristics on a timeline with scores assigned to anomalous access behavior. Traditional security alerts are also scored, attributed to identities, and placed on the activity timeline. All systems touched by compromised credentials of insiders are identified to reveal the attacker's path through the IT environment.

### How Advanced Analytics Works

Exabeam uses a two-layer approach to identifying incidents. The first layer involves collecting data from log sources and from external context data sources. The second layer involves processing data through the Analysis Engine and the Exabeam Stateful User Tracking™ technology.

In the first layer, events are normalized and enriched with contextual information about users and assets in order to understand entity activities within the environment across a variety of dimensions. At this level, statistical modeling profiles the behaviors of network entities while machine learning is applied in the areas of context estimation, for example, to distinguish between users and service accounts.

In the second layer, as events are processed, Exabeam uses Stateful User Tracking to connect user activities across multiple accounts, devices, and IP addresses. It places all user credential activities and characteristics on a timeline with scores assigned to anomalous access behavior. Traditional security alerts are also scored, attributed to identities, and placed on the activity timeline. The Analysis Engine updates risk scores based on input from the anomaly detection process as well as from other external data sources. The Analysis Engine brings in the security knowledge and expertise in order to bubble up significant anomalies. Incidents are generated when scores exceed defined thresholds.

### Fetch and Ingest Logs

The way that logs are fetched and ingested depends on the version of Advanced Analytics you are using:

- **i63 and Later**

In this version, log data enters Exabeam through a set of collector services. These services collect data from servers, applications, databases, and other devices across an infrastructure, whether the source is local, remote, or cloud-based. Logs can be collected from on-premises sites using [Site Collectors](#) or from third-party cloud vendors using [Cloud Collectors](#).

The ingested logs are processed into events by the unified ingestion pipeline (UIP). These events conform to the hierarchical [common information model](#) that informs the data structure for all Exabeam products. These events are then processed through a UIP Advanced Analytics Integration service that transforms them into events that are readable by Advanced Analytics.

- **i60 to i62**

In this version, Exabeam can fetch logs from SIEM log repositories and also ingest logs via Syslog. Currently log ingestion is supported from Splunk, Microfocus ArcSight, IBM QRadar, McAfee ESM and RSA Security Analytics, as well as other data sources such as Data Lake. For Splunk and QRadar, log ingestion is via external APIs and Syslog is used for all others. For SIEM solutions, such as LogRhythm, McAfee ESM, and LogLogic, ingestion is via Syslog forwarding. The ingested logs are processed into events by the LIME (Log Ingestion and Message Extraction) engine.

### **Data Volume Limitations**

With a Fusion license, you must use Advanced Analytics in a way that does not interfere with or disrupt the SaaS environment, servers, or networks. Additionally, you must not engage in any activity that interferes with the integrity or proper working of the service.

To ensure the proper operation of Advanced Analytics, the following limitations are in place for each 1GB of daily average consumption purchased:

- You can not exceed 17 EPS.



- If you exceed 10 EPS in any 30 minute time period, you can not exceed 10 EPS for at least four hours afterward.

For example, if you purchased 100 GB daily average consumption but exceed 1,000 EPS for over 30 minutes in a 4-hour timeframe, this can cause the performance and functionality of cloud-delivered Advanced Analytics to decrease.

**⚠ CAUTION**

Failure to adhere to operational limitations, can result in downtime or errors. In the event that you exceed the limitations, Exabeam will not be responsible for any resulting issues.

## Add Context

Logs tell us what the users and entities are doing while context tells us who the users and entities are. Context data typically comes from identity services such as Active Directory. This data enriches the logs to help with the anomaly detection process or the data can be used directly by the Analysis Engine for fact-based rules. Regardless of where this context data is used, it goes through the anomaly detection process as part of an event. Examples of context information include the location for a given IP address, ISP name for an IP address, and department for a user. Contextual information can also include data from HR Management Systems, Configuration Management Databases, Identity Systems, etc. Threat intelligence feeds are another example of contextual data which can be used by the anomaly detection process to identify activity from a known malicious domain or IP address.

## Detect Anomalies

In this part of the process, machine learning algorithms are used to identify anomalous behaviors. The anomalies may be relative to a single user, session, or device, or relative to group behavior. For example, some anomalies may refer to a behavior that is anomalous for a user relative to their past history. Other anomalies may take into account anomalous behaviors relative to people with roles similar to the individual (peer group), location, or other grouping mechanisms.

The algorithms are constantly improved upon to increase the speed and accuracy of numerical data calculations. This in turn improves the performance of Advanced Analytics.

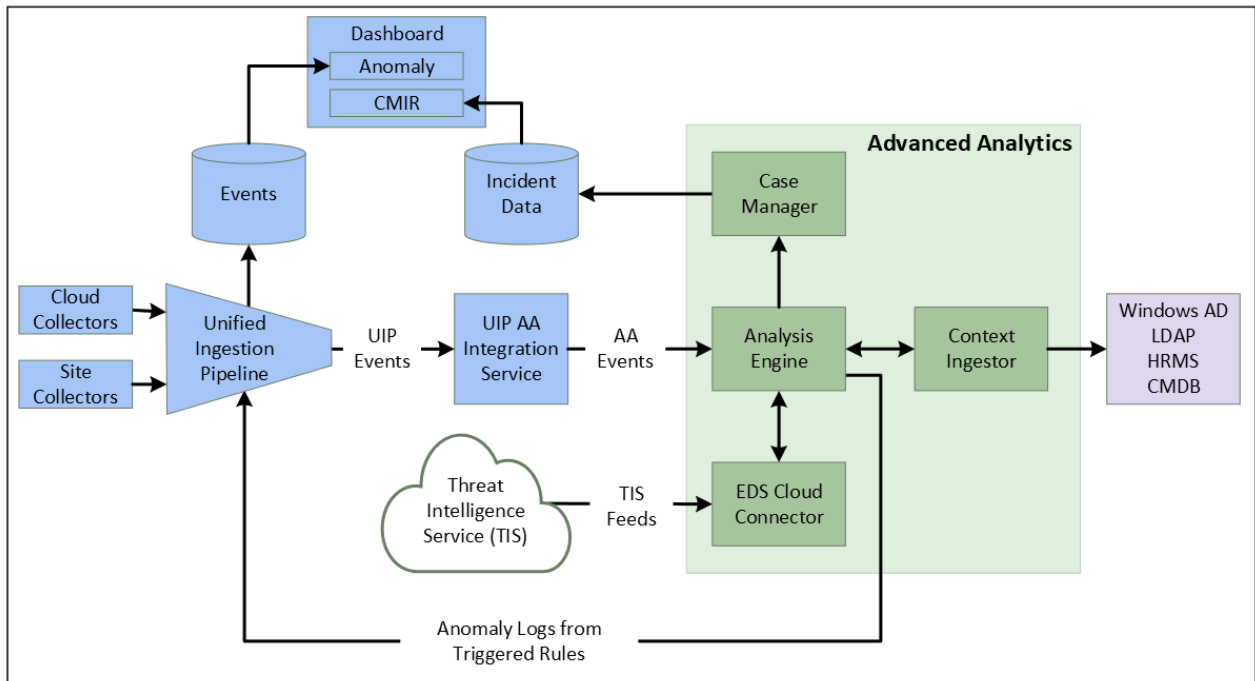
## Assess Risk

The Analysis Engine treats each session as a container and assigns risk scores to the events that are flagged as anomalous. As the sum of event risk scores reach a threshold (a default value of 90), incidents are automatically generated within the Case Management module or escalated as incidents to an existing SIEM or ticketing systems. Event scores are also available to be queried through the user interface. In some cases, these scores reflect more than just information considered anomalous on the basis of behavior log feeds alone. In some cases, the scores are provided in connection with other security alerts that may be generated by third-party sources (for example, FireEye or CrowdStrike alerts). These security alerts are integrated into user sessions and scored as factual risks.

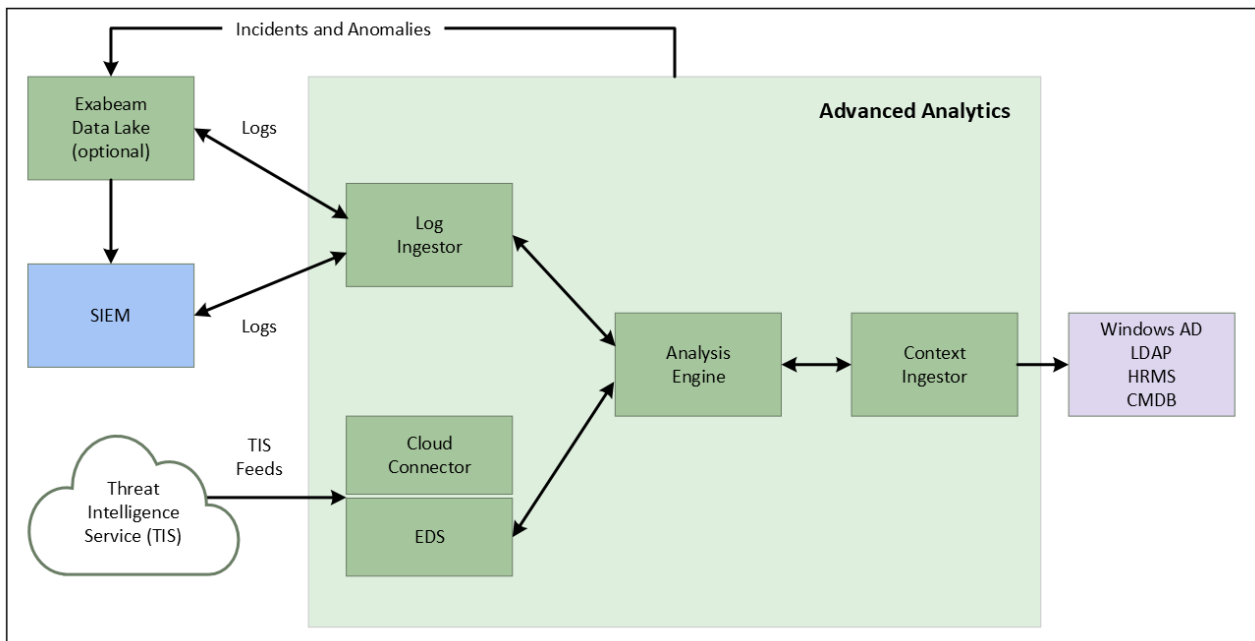
## Data Flow Diagrams

The flow of data through the components of Advanced Analytics depends on which version you are using. Select the relevant version below to view the appropriate diagram.

**Diagram for i63 and later**



**Diagram for i60 to i62**



## Data Retention in Advanced Analytics

Advanced Analytics retains both event log and session data for limited periods of time. Retention times depend on the retention categories and the time periods defined in your purchased license.

Data in Advanced Analytics is divided into the following retention categories:

- *Raw logs.* The original event logs sent to Advanced Analytics.



**NOTE**

Your [Event Selection](#) policy determines which event logs are sent to Advanced Analytics.

- *Enriched events.* The event logs created by Advanced Analytics when the raw logs are received and enriched with contextual data.



**NOTE**

Until a raw event log is purged from the system, you can view the event in both its original and enriched forms.

- *Events that triggered rules.* Enriched events that have triggered or helped to trigger one or more rules.
- *User and Asset Sessions.* The containers that Advanced Analytics creates for both users and assets to represent the different timeframes of the enriched events attributed to them. Sessions are retained for the same amount of time as the enriched events that comprise them. If a session includes one or more events that were involved in triggering rules, the session is retained for as long as the event(s) that triggered the rules are retained; however, any events in the session that did not trigger rules are removed from the session when their retention period expires.

When the date of an event log exceeds the retention period of its category, the event is purged from the system. Likewise, when all the event logs associated with a session have been purged, the session is purged.

For details on the retention periods included with your license, see the [Product Entitlement](#) page on the Exabeam Community.


## Configure Log Management

Large enterprise environments generally include many server, network, and security technologies that can provide useful activity logs to trace who is doing what and where. In Advanced Analytics, several aspects of log management can be configured. For more information, see the links below:

- [Log Ingestion Settings](#)
- [Log Feeds](#)
- [Training and Scoring](#)
- [Notifications](#)
- [Audit Logs](#)

### Log Ingestion Settings

To configure log ingestion settings in Advanced Analytics:

1. In the sidebar, click **SETTINGS** , then select **Analytics**.
2. Under **Log Management**, select **Log Ingestion Settings**.
3. Log ingestion settings are different depending on which version of Advanced Analytics you are using. Click the appropriate links below for more information.

#### i63 and Later

Log ingestion is handled via a unified ingestion pipeline (UIP). Visibility into the UIP is provided through the cloud-native [Log Stream](#) and [Live Tail](#) functionality.

Log data enters the UIP through a set of collector services. These services collect data from servers, applications, databases, and other devices across an infrastructure, whether the source is local, remote, or cloud-based.

In the UIP, the ingested logs are processed into events that conform to the hierarchical [common information model](#). These events are then processed through a UIP Advanced Analytics Integration service that transforms them into events that are readable by Advanced Analytics.

For information about setting up specific collector services, see the following:

- For third-party cloud sources, see [Cloud Collectors](#).
- For on-premises sources, see [Site Collectors](#).

#### i60 to i62

Log ingestion is handled by the LIME engine (Log Ingestion and Message Extraction). LIME can fetch data from SIEM log repositories or ingest it via Syslog. It normalizes the raw logs into events that the rest of the Advanced Analytics pipeline can process.

Currently, log ingestion is supported from the sources listed below.

- 
- |              |                |                          |
|--------------|----------------|--------------------------|
| • Data Lake  | • HP ArcSight  | • RSA Security Analytics |
| • Splunk     | • IBM QRadar   | • Sumo Logic             |
| • ServiceNow | • McAfee Nitro | • Google Cloud Pub/Sub   |
- 

On the **Log Ingestion Settings** page in Advanced Analytics, you can do the following:

- **Enable Syslog Ingestion**
- Configure Syslog **Options**
- View **Syslog Stats**
- **Add** a New Log Source

For Splunk and QRadar, log ingestion occurs via external APIs. Syslog is used for the other sources. For SIEM solutions, such as LogRhythm, McAfee ESM, and LogLogic, ingestion occurs via Syslog forwarding.



### NOTE

The Syslog destination is your site collector IP/FQDN. Only TLS connections are accepted in port TCP/515.

For more information about about configuring log ingestion in this version of Advanced Analytics, see the following subsections:

- [View Insights about Syslog-Ingested Logs](#) – Test the data pipeline of incoming logs.
- [Ingest Logs from Google Cloud Pub/Sub into Advanced Analytics](#) – Configure Google Pub/Sub as a log source.
- [Set up a Log Feed](#) – Configure log feeds from a SIEM source.

## View Insights About Syslog-Ingested Logs



### NOTE

The information in this section applies to Advanced Analytics versions i60–i62.

Advanced Analytics has the ability to test the data pipeline of logs coming in via Syslog.



### NOTE

This option is only available if the **Enable Syslog Ingestion** button is toggled on.

Click the **Syslog Stats** button to view the number of logs fetched, the number of events parsed, and the number of events created. A warning is displayed that lists any event types that were not created within the Syslog feed that was analyzed.

In this view, you can also select **Options** to configure the time range and number of log events tested.

## Ingest Logs from Google Cloud Pub/Sub into Advanced Analytics



### NOTE

The information in this section applies to Advanced Analytics versions i60–i62.


To create events from Google Cloud Pub/Sub topics, you must configure Google Pub/Sub as an Advanced Analytics log source.

### *Prerequisites to Configure Google Pub/Sub*

- [Create a Google Cloud service account](#) with **Pub/Sub Publisher** and **Pub/Sub Subscriber** permissions.
- [Create and download a JSON-type service account key](#). You use this JSON file later.
- [Create a Google Cloud Pub/Sub topic](#) with **Google-managed key** encryption.
- For the Google Cloud Pub/Sub topic you created, [create a subscription](#) with specific settings:
  - **Delivery type** – Select **Pull**.
  - **Subscription expiration** – Select **Never expire**.
  - **Retry policy** – Select **Retry immediately**.

Save the subscription ID to use later.

### *Procedure to Configure Google Pub/Sub*

1. In the sidebar, click **SETTINGS** , then select **Analytics**.
2. Under **Log Management**, select **Log Ingestion Settings**.
3. Click **ADD**, then from the **Source Type** list, select **Google Cloud Pub/Sub**.
4. Enter information about your Google Cloud Pub/Sub topic:
  - **Description** – Describe the topic, what kinds of logs you're ingesting, or any other information helpful for you to identify this as a log source.
  - **Service key** – Upload the Google Cloud service account key JSON file you [downloaded](#).
  - **Subscriptions**
    - **Subscription name** – Enter the Google Cloud Pub/Sub subscription ID you [created](#).
    - **Description** – Describe the subscription, to which Google Cloud Pub/Sub topic it was created, or what messages it receives.
5. To verify the connection to your Google Cloud Pub/Sub topic, click **TEST CONNECTION**. If you see an error, verify the information you entered then retest the connection.
6. Click **SAVE**.
7. [Restart](#) Log Ingestion and Messaging Engine (LIME).  
To ingest specific logs from your Google Cloud Pub/Sub topic, [configure](#) a log feed.

## Log Feeds

**NOTE**

The log feed setup information in this section applies to Advanced Analytics versions i60–i62.

Advanced Analytics can be configured to fetch log data from a SIEM. Administrators can configure log feeds that can be queried during ingestion. Exabeam provides out-of-the-box queries for various log sources; or you can edit them and apply your own.

Once a log feed is set up, you can perform a test query that fetches a small sample of logs from the log management system. You can also parse the sample logs to make sure that Advanced Analytics is able to normalize the logs. If the system is unable to parse the logs, reach out to Customer Success and the Exabeam team will create a parser for those logs.

### Set up a Log Feed


To set up a log feed, navigate to **Settings > Log Feeds** and complete the workflow to create a new log feed. You will be prompted to publish the feed to let the Advanced Analytics processing engine know that the feed is ready for consumption. You have two options:

- **Publish** – If you choose to publish the feed is placed into publishing mode and will be picked up by the processing engine at the top of the hour.
- **Draft** – If you choose not to publish, the feed remains in draft mode and will not be picked up by the processing engine. Draft mode allows you to add multiple feeds and text queries without worrying that the feed will be picked up for processing or that errors will be caused if the feed is deleted. You can always choose to publish a draft feed at another time.

## Training and Scoring

To establish a baseline, Advanced Analytics extensively profiles the people, asset usage, and sessions in your environment. For example, in a typical deployment, Advanced Analytics begins by examining 60–90 days of an organization's logs. After the initial baseline analysis is done, Advanced Analytics begins assigning scores to each session based on the amount and type of anomalies in the session.

To configure the Advanced Analytics initial training period:

1. In the sidebar, click **SETTINGS** , then select **Analytics**.
2. Under **Log Management**, select **Training & Scoring**.
3. Configure the processing and scoring parameters.



## Notifications

You can configure Advanced Analytics to send notification about system health, notable sessions, anomalies, and other important system information. You can configure notification to be sent in the following formats:

- **Log repository** – Notifications can be sent to a log repository in a structured data format using the Syslog protocol. These notifications are formatted so machines, like your log repository, can easily understand them.
- **Email** – Notifications can be sent to an email account in a format that's more human-readable.

### Set Up Notifications to a Log Repository, Ticketing System, or SIEM

To configure notifications to a log repository, log ticketing system, or SIEM using the Syslog protocol:

1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **NOTIFICATIONS**, select **Setup Notifications**.
3. Click add , then select **Syslog Notification**.
4. Configure the following notification settings:
  - **IP / Hostname** – Enter the IP or hostname of your Syslog server.
  - **Port** – Enter the port your Syslog server uses.
  - **Protocol** – Select the network protocol your Syslog server uses to send messages: TCP, SSL\_TCP, or UDP.
  - **Syslog Security Level** – Assign a severity level to the notification:
    - **Informational** – Normal operational events, no action needed.
    - **Debug** – Useful information for debugging, sent after an error occurs.
    - **Error** – An error has occurred and must be resolved.
    - **Warning** – Events that will lead to an error if you don't take action.
    - **Emergency** – Your system is unavailable and unusable.
    - **Alert** – Events that should be corrected immediately.
    - **Notice** – An unusual event has occurred.
    - **Critical** – Some event, like a hard device error, has occurred and your system is in critical condition.
  - **Notifications by Product** – Select the events for which you want to be notified:
    - **Advanced Analytics**:
      - **System Health** – All system health alerts for Advanced Analytics.
      - **Notable Sessions** – A user or asset has reached a risk threshold and become notable. This notification describes which rule was triggered and contains any relevant information.
      - **Anomalies** – A rule has been triggered.
      - **AA/CM/OAR Audit** – An Exabeam user does something in Advanced Analytics, Case Manager, or Incident Responder that's important to know when auditing their activity history; for example, when someone modifies rule behavior, changes log sources, or changes user roles and permissions.



- **Job Start** – Data processing engines have started processing a log.
- **Job End** – Data processing engines have stopped processing a log.
- **Job Failure** – Data processing engines have failed to process a log.

5. Click **ADD NOTIFICATION**.

6. **Restart** the Analytics Engine.



### Set Up Notifications to Email

You can configure email notifications for both Advanced Analytics and for some Incident Responder actions. The Incident Responder notifications can include:

- Notify User By Email Phishing
- Phishing Summary Report
- Send Email
- Send Template Email
- Send Indicators via Email

If you configure these settings correctly, Incident Responder uses *IRNotificationSMTPService* as the service for these actions. If you configure these settings incorrectly, these actions won't work correctly.

To configure human-readable email notifications:

1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **NOTIFICATIONS**, select **Setup Notifications**.
3. Click add , then select **Email Notification**.
4. Configure the following notification settings:
  - **IP / Hostname** – You must enter `cloudrelay1.connect.exabeam.com`.
  - **Port** – Enter the port number for your outgoing mail server.
  - **SSL** – You must select this box.
  - **Username Required** – If your mail server requires a username, select the box, then enter the username.
  - **Password Required** – If your mail server requires a password, select the box, then enter the password.
  - **Sender Email Address** – Enter `<yourinstance>@notify.exabeam.com`.
  - **Recipients** – List the email addresses to receive these email notifications, separated by a comma.
  - **E-mail Signature** – Enter text that's automatically added to the end of all email notifications.

- **Notifications by Product** – Select the events for which you want to be notified.
  - Incident Responder:
    - **System Health** – All system health alerts for Case Manager and Incident Responder.
  - Advanced Analytics:
    - **System Health** – All system health alerts for Advanced Analytics.
    - **Notable Sessions** – A user or asset has reached a risk threshold and become notable. This notification describes which rule was triggered and contains any relevant information.
    - **Anomalies** – A rule has been triggered.
    - **AA/CM/OAR Audit** – An Exabeam user does something in Advanced Analytics, Case Manager, or Incident Responder that's important to know when auditing their activity history; for example, when someone modifies rule behavior, changes log sources, or changes user roles and permissions.
    - **Job Start** – Data processing engines have started processing a log.
    - **Job End** – Data processing engines have stopped processing a log.
    - **Job Failure** – Data processing engines have failed to process a log.

5. Click **ADD NOTIFICATION**.

### Syslog Notifications Key-Value Pair Definitions

The Incident Notifications sent via Syslog to your SIEM use the following parameters in their key-value pairs. The pairs are separated by a space (for example, reasons\_count="3" score="135").

The following tables define each extension field key in Syslog messages by Event Type for Advanced Analytics v.i48 and later.

#### Event Type: System Health

Key Name	Key Value	Description	Syslog Key Value Example
service	Service of the system health	The service of the system health.	service="Analytics Log Ingestion"
status	Status of the system health	The status of the system health, either "running" or "stopped".	status="stopped"

#### Event Type: Notable Sessions / Anomalies

Key Name	Key Value	Description	Syslog Key Value Example
id	Session, sequence, or asset sequence ID	The ID of the session, sequence, or asset sequence ID.  Also the triggered rule containerId of a session or sequence.	id="npage-201905151524"

## Configure Log Management

Key Name	Key Value	Description	Syslog Key Value Example
url	URL to the session, sequence, or asset sequence	The URL to the session, sequence, or asset sequence in the timeline.	url="https://<ExabeamAA>:443/uba/#user/admin/timeline/admin-201906241700"
entity_value	Asset sequence entity value	The asset sequence entity value.	entity_value="dev_kr"
score	Session, sequence, or asset sequence risk score	The session, sequence, or asset sequence risk score.  Also the triggered rule risk score.	score="316"
sequence_type	Type of sequence	The type of sequence or asset sequence entity name.	sequence_type="lockout"
start_time	loginTime in UTC format	The loginTime for the session or sequence.  Also the day for the asset sequence.	start_time="2019-05-25T08:05:24-07:00"
end_time	logoutTime in UTC format	The logoutTime for the session or sequence.  Also the start_time + 24 hours for the asset sequence.	end_time="2019-05-25T18:26:24-07:00"
status	Status of the session or sequence	The status of the session or sequence, either "open" or "closed". It is closed if logoutTime is != 0, otherwise it is open.	status="closed"
user	Session or sequence username	The session or sequence user.  Also the triggered rule username.	user="admin"
src_host	Session or sequence loginHost	The session or sequence loginHost.  Also the triggered rule src host. (optional)	src_host="dc_464"
src_ip	Session src_ip	The session src_ip.  Also the triggered rule src_ip. (optional)	src_ip="10.55.0.123"
accounts	Session or sequence accounts	The session or sequence accounts.	accounts="achan, tmiles, dgreen, mbridges"
labels	Session or sequence labels	The session or sequence labels.	labels="TIME"
assets_count	Session or sequence assets count	The total number of assets recorded for the session or sequence.	assets_count="23"
assets	Session or sequence assets	The names of all unique assets for the session or sequence.	assets="srv_123_dev, 10.23.123.56, tks_en_0b_jt"
zones	Session zones	The session zones.  Also the asset sequence zones. (optional)	zones="atlanta office"

## Configure Log Management

Key Name	Key Value	Description	Syslog Key Value Example
reasons	All unique rule definition descriptions for any triggered rules for the session	By default, only takes the top three by triggered rule risk score. Configurable through <code>exabeam_custom_config</code> ScoreManager. IncidentReasonsLimit.	reasons="It is abnormal for this user to perform account management activities (a user created and added to a group) on this day of the week. Account management events are notable because they can provide a path for an attacker to move laterally through a system."
reasons_count	Session or sequence reasons count	The total number of reasons recorded for the session or sequence.	reasons_count="8"
events_count	Session or sequence events count	The total number of events recorded for the session or sequence.	events_count="779"
alerts_count	Session or sequence alerts count	The total number of security events recorded for the session or sequence.	alerts_count="7"
asset_labels (optional)	Labels for the asset	The labels for the asset.	asset_labels="SOURCE HOST"
asset_locations	Locations for the asset	The locations for the asset.	asset_locations="los angeles office"
top_users	Top users for the asset	The top users for the asset.	top_users="adonald"
host_name (optional)	Hostnames for the asset	The hostnames for the asset.	host_name="adonald-win7"
ip_address (optional)	IP addresses for the asset	The IP addresses for the asset.	ip_address="10.0.0.3"
dest_host (optional)	dest_host of the triggered rule	The dest_host of the triggered rule.	dest_host="It-adonald-123"
dest_ip (optional)	dest_ip of the triggered rule	The dest_ip of the triggered rule.	dest_ip="10.23.121.87"
event_time	Time of the event	When the event took place.	event_time="8:26:00"
event_type	Type of the event	The type of the event.	event_type="remote-access"
host (optional)	Host for the triggered rule	The host for the triggered rule.	host="atl-file-01"
domain (optional)	Domain for the triggered rule	The domain for the triggered rule.	domain="kt_cloud"
raw (optional)	Raw "payload"	The raw "payload" for the triggered rule or event.	raw="http://exampleurl"
rule_id	Rule definition ID	The rule definition ID for the triggered rule.	rule_id="WL-HG-F"
rule_name	Rule definition name	The rule definition name for the triggered rule.	rule_name="Account switch by new user"
rule_description	Rule definition description	The rule definition description for the triggered rule.	rule_description="User for which we have insufficient data logged onto a Domain Controller"

### Event Type: Advanced Analytics/Case Manager/OAR Audit

## Configure Log Management

Key Name	Key Value	Description	Syslog Key Value Example
time	Time in Unix timestamp	When the event took place.	time="1559076154132"
user	Admin or user's username	The user performing the activity in Exabeam.	user="admin"
host	Hostname, and if not available, IP address	The machine logging the event.	host="10.0.0.4"
src_ip	Client IP address	The IP from which the user connected to Exabeam.	src_ip="10.0.0.6"
event_type	Event type	The event type, such as "app-activity", "remote-logon", and "logout".	event_type="failed-app-login"
app	Name of the application	The name of the application that logs the event, such as "Exabeam Advanced Analytics". If it is not possible to get the app type then "Exabeam" is used.	app="Exabeam"
event_subtype	Exabeam Audit Event	The Exabeam Audit Event identifier.	event_subtype="Exabeam Audit Event"
activity	Activity type	The activity type, such as "User added", "Role updated", and "LDAP server removed".	activity="Failed log in"
additional_info	Activity details	The activity details, such as those pertaining to user, role, and LDAP updates.	additional_info="User 'admin' failed to login"

## Event Type: Job Status

Key Name	Key Value	Description	Syslog Key Value Example
job_status	Job status	The job status, either "Started", "Failed", or "Completed".	job_status="Started"
job_details	Job details	The job details, which includes modified rules and reprocess times.	job_details="Modified rules: rule AM-OG-A has new score 40.0 ,rule AM-GOU-A has new score 40.0 ,rule AM-GA-AC-A has new score 40.0. Reprocess starts from May 5 2014, 7:00AM (UTC), ends on May 7 2018, 6:59AM (UTC)."
job_id	Job ID	The job ID string.	job_id="5c1ace5c123 b3801207481f"
created_by	Admin or user's username	The user who created the job in Exabeam.	created_by="admin"
timestamp	Timestamp in UTC format	When the event took place.	timestamp="December 19 2018, 11:05PM (UTC)"
start_time (optional)	Timestamp in UTC format	When the event started.	start_time="March 12 2019, 10:05PM (UTC)"
end_time (optional)	Timestamp in UTC format	When the event ended.	end_time="March 13 2018, 12:05AM (UTC)"

## Audit Logs

Audit logs represent user, object, or setting events in your organization. Specific events related to all Exabeam users are logged, including activities within the user interface as well as configuration activities.

Advanced Analytics audit logs are stored. The entire auditing history is stored and you cannot purge audit logs or set retention limits.

To access the activity data, you can forward audit logs via Syslog to an existing SIEM, to Data Lake, or to Search. Exabeam sends the Advanced Analytics activity data every five minutes. To access audit logs via Syslog, follow the notification setup procedure in [Set Up Notifications to a Log Repository, Ticketing System, or SIEM](#).



### NOTE

Advanced Analytics activity log data is not masked or obfuscated when sent via Syslog. It is your responsibility to upload the data to a dedicated index which is available only to users with appropriate privileges.

The following events are logged:

- 
- |                                       |  |
|---------------------------------------|--|
| • Log in and log out                  | • Adding or editing of secured resources           |
| • Failed log in                       | • API cluster authorization events                 |
| • User addition, update, and removal  | • Log source addition, update, and deletion        |
| • Role addition, update, and deletion | • Log feed addition, update, and deletion          |
| • Permission addition and deletion    | • Syslog enable and disable                        |
| • Audit being turned on or off        | • Full and partial acceptance of a session         |
| • Token create, read, and update      | • Full and partial acceptance of a lockout         |
| • Reindex job create and initiate     | • Full and partial acceptance of an asset sequence |
| • Threat Hunter Search                | • Starting of a session                            |
| • Component restart                   | • Starring of an asset sequence                    |
| • SAML events                         | • Watchlist addition, update, and delete           |
- 

An additional type of audit logging is available for applications in the Exabeam Security Operations Platform. Access to these stored audit logs is available in Search. For ease of use, an **Audit Logs** tab is accessible in the Search query builder. For information about using the **Audit Logs** tab, see [Basic Search](#) in the Search Feature Guide.

Events from the following Exabeam Security Operations Platform applications are logged:

- Authentication
- Threat Center
- Correlation Rules
- Search
- Settings, including

## Configure Log Management

- Users
- Roles
- Single sign-on
- API keys

## Set Up Admin Operations

This section contains information about certain Advanced Analytics administrative operations. For more information, see the links below:

- [License Management](#)
- [Cluster Authorization Token](#)
- [Exabeam Engines](#)

### Exabeam Licenses

Exabeam products require a license in order to function. These licenses determine which Exabeam products and features you can use. You are not limited by the amount of external data you can ingest and process.

There are multiple types of Exabeam product licenses available. Exabeam bundles these licenses together and issues you one key to activate all purchased products. For more information on the different licenses, see [Types of Exabeam Product Licenses](#).

#### License Lifecycle

When you first install Exabeam, the installed instance uses a 30 day grace period license. This license allows you to try out all of the features in Exabeam for 30 days.

#### Grace Period

Exabeam provides a 30-day grace period for expired licenses before products stop processing data. During the grace period, you will not experience any change in product functionality. There is no limit to the amount of data you can ingest and process.

When the license or grace period is 14 days away from expiring, you will receive a warning alert on the home page and an email.

You can request a new license by contacting your Exabeam account representative or by opening a support ticket.

#### Expiration Period

When your grace period has ended, you will start to experience limited product functionality. Contact your Exabeam representative for a valid license to restore all product features.

For Advanced Analytics license expirations, the Log Ingestion Engine will continue to ingest data, but the Analytics Engine will stop processing. Threat Hunter and telemetry will also stop working.


You will receive a critical alert on the home page and an email.

#### License Alerts



License alerts are sent via an alert on the home page and in email when the license or grace period is 14 days away from expiring and when the grace period expires.

The home page alert is permanent until resolved. You must purchase a product license or renew your existing license to continue using Exabeam.

To check the status and details of your license, navigate to **Settings**  > **Admin Operations** > **Licenses**.

### Types of Exabeam Product Licenses

Exabeam licenses specify which products you have access to and for how long. For information about specific license types, and detailed tables of supported features for each license type, click the appropriate links below.

- [Exabeam Security Operations Portfolio License Types](#) – Includes the latest cloud-delivered subscription product licenses.
- [Fusion License Types](#) – Includes the cloud-delivered Fusion XDR and Fusion SIEM product licenses.
- Legacy Licenses – Includes the following legacy product licenses:
  - **User Analytics** – This is the core product of Advanced Analytics . Exabeam’s user behavioral analytics security solution provides modern threat detection using behavioral modeling and machine learning.
  - **Threat Hunter** – Threat Hunter is a point and click advanced search function which allows for searches across a variety of dimensions, such as Activity Types, User Names, and Reasons. It comes fully integrated with User Analytics.
  - **Exabeam Threat Intelligence Services (TIS)** – TIS provides real-time actionable intelligence into potential threats to your environment by uncovering indicators of compromise (IOC). It comes fully integrated with the purchase of an Advanced Analytics V3 license. TIS also allows access to telemetry.
  - **Entity Analytics (EA)** – Entity Analytics offers analytics capabilities for internet-connected devices and entities beyond users such as hosts and IP addresses within an environment. Entity Analytics is available as an add-on option. If you are adding Entity Analytics to your existing Advanced Analytics platform, you will be sent a new license key. Note that you may require additional nodes to process asset oriented log sources.
  - **Incident Responder** – Also known as Orchestration Automation Response. Incident Responder adds automation to your SOC to make your cyber security incident response team more productive.  
Incident Responder is available as an add-on option. If you are adding Incident Responder to your existing Advanced Analytics platform, you will be sent a new license key. Note that you may require additional nodes to support automated incident responses.
  - **Case Manager** – Case Manager can fully integrate into Advanced Analytics enabling you to optimize analyst workflow by managing the life cycle of your incidents.

Case Manager is available as an add-on option. If you are adding Case Manager to your existing Advanced Analytics platform, you will be sent a new license key. Note that you may require additional nodes to support this module extension.

## View Version Information

View the Advanced Analytics release version of your deployment, as well as build numbers that constitute the release version (Advanced Analytics product, Incident Responder product, Data Lake product, Exabeam platform common services, and Security content).

To view version information:

1. Click the options menu, **☰**, in the top right corner of any page. A menu expands.
2. Scroll to the **About this version** section at the bottom of the expanded menu.

## Cluster Authorization Token

The cluster authorization token is used to verify identities between clusters that have been deployed in phases as well as HTTP-based log collectors. Each peer cluster in a query pool must have its own token. You can set expiration dates during token creation or manually revoke tokens at any time.

To generate a token:


1. Navigate to **Settings > Core > Admin Operations > Cluster Authorization Token**. The Cluster Authorization Token page is displayed.
2. In the **Cluster Authorization Token** page, click **+** to open the **Setup Token** dialog box.

### Setup Token

**Token Name \***

**Expiry Date \***

Manual Date Entry



Permanent (no expiry date)

**Permission Level \***

**Default Roles**

<input type="checkbox"/>	Administrator	▼
<input type="checkbox"/>	Tier 3 Analyst	▼
<input type="checkbox"/>	Tier 1 Analyst	▼
<input type="checkbox"/>	Auditor	▼
<input type="checkbox"/>	Data Privacy Officer	▼

**Custom Roles**

[CANCEL](#) [ADD TOKEN](#)

3. Enter a **Token Name**, and then select an **Expiry Date**.

**! IMPORTANT**


Token names can contain only letters, numbers, and spaces.

4. Select the **Default Roles** for the token.
5. Click **Add Token**.  
Use this generated file to allow your APIs to authenticate by token. Ensure that your API uses `ExaAuthToken` in its requests. For curl clients, the request structure resembles the following:

```
curl -H "ExaAuthToken:<generated_token>" https://<external_host>:<api_port>/<api_request_path>
```

## Exabeam Engines

To start any Exabeam engines, navigate to the **Exabeam Engine** page in Advanced Analytics:

1. From the left sidebar, click **SETTINGS** , then select **Analytics**.
2. Under **Admin Operations**, select **Exabeam Engine**.
3. Procedures for starting Exabeam engines vary depending on the Advanced Analytics version you are using. Click the appropriate links below for more information.

### i63 and Later

In this Advanced Analytics version, you can use the UIP reprocessing option or restart the analytics engine:

- **Exabeam Unified Log Ingestion Engine** – Click **UIP Log Reprocessing** to open the cloud-based [Log Stream](#) functionality. Log Stream provides visibility into the unified ingestion pipeline with the following tabs:
  - **Re-parsing Jobs** – You can opt to re-parse logs by scheduling a re-parsing job.
  - **Live Tail** – View samples of incoming data in real time to ensure proper processing.
- **Exabeam Analytics Engine** – Restart the analytics engine so that it can process log feeds.
  1. In the **Exabeam Analytics Engine** panel, click **Restart Processing** and select a restart option from the following settings:
    - **Restart the engine** – The engine continues processing from where it left off.
    - **Restart from the initial training period** – The engine continues processing from the initial training period.
    - **Restart from a date** – The engine chooses the nearest snapshot available for the specified date and reprocesses from this date.
  2. Click **Process** to start the engine. The system validates any changes and checks for errors. If errors are identified, they are listed and the engine does not start processing. If no errors are identified, the engine starts.

## i60 to i62

In this Advanced Analytics version, you can restart the LIME engine or the analytics engine:

- **Exabeam Log Ingestion Engine** – Restart the LIME engine so that it ingests logs from log feeds that are defined in the Advanced Analytics **Log Feeds** settings.
  1. Click **Ingest Log Feeds**, and select specific log feeds for restart.
  2. Select a restart option from the following settings:
    - **Restart the engine** – The engine continues processing from where it left off.
    - **Restart from the initial training period** – The engine continues processing from the initial training period.
    - **Restart from a date** – The engine continues processing from a specified date.
  3. Click **Ingest feeds** to start the engine.
  
- **Exabeam Analytics Engine** – Restart the analytics engine so that it can process log feeds.
  1. Click **Restart Processing** and select a restart option from the following settings:
    - **Restart the engine** – The engine continues processing from where it left off.
    - **Restart from the initial training period** – The engine reprocesses all configured log feeds from the initial training period.
    - **Restart from a date** – The engine chooses the nearest snapshot available for the specified date and reprocesses from this date.
  2. Click **Process** to start the engine. The system validates any changes and checks for errors. If errors are identified, they are listed and the engine does not start processing. If no errors are identified, the engine starts.

## Generate a Support File

Generate the technical support information Exabeam Customer Success needs to understand a problem in your system. Once you have these files, open a support ticket in the Exabeam Community.

1. Click the options menu, **☰**, in the top right corner of any page, and then select **Generate Support File**.
2. Select the Exabeam products for which you want to generate technical support file. If Exabeam Customer Success requests a support file that includes the last five rotated logs from the Analytics Engine or Log Ingestion Messaging Engine (LIME), select **ADDITIONAL ROTATED LOGS**.
3. Click **GENERATE SUPPORT**. It may take up to five minutes to generate a support file. The file downloads automatically. Attach the file to your case ticket on the [Exabeam Community](#).

## Set Up Authentication and Access Control

### What Are Accounts & Groups?

#### Peer Groups

Peer groups can be a team, department, division, geographic location, etc. and are defined by the organization. Exabeam uses this information to compare a user's behavior to that of their peers. For example, when a user logs into an application for the first time Exabeam can evaluate if it is normal for a member of their peer group to access that application. When Dynamic Peer Grouping is enabled, Exabeam will use machine learning to choose the best possible peer groups for a user for different activities based on the behaviors they exhibit.

#### Executives

Exabeam watches executive movements very closely because they are privileged and have access to sensitive and confidential information, making their credentials highly desirable for account takeover. Identifying executives allows the system to model executive assets, thereby prioritizing anomalous behaviors associated with them. For example, we will place a higher score for an anomaly triggered by a non-executive user accessing an executive workstation.

#### Service Accounts

A service account is a user account that belongs to an application rather than an end user and runs a particular piece of software. During the setup process, we work with an organization to identify patterns in service account labels and uses this information to classify accounts as service accounts based on their behavior. Exabeam also adds or removes points from sessions based on service account activity. For example, if a service account logs into an application interactively, we will add points to the session because service accounts should not typically log in to applications.

### What Are Assets & Networks?

#### Workstations & Servers

Assets are computer devices such as servers, workstations, and printers. During the setup process, we will ask you to review and confirm asset labels. It is important for Exabeam to understand the asset types within the organization - are they Domain Controllers, Exchange Servers, Database Servers or workstations? This adds further context to what Exabeam sees within the logs. For example, if a user performs interactive logons to an Exchange Server on a daily basis, the user is likely an Exchange Administrator. Exabeam automatically pulls in assets from the LDAP server and categorizes them as servers or workstations based on the OS property or the Organizational Units they belong to. In this step, we ask you to review whether the assets tagged by Exabeam are accurate. In addition to configuration of assets during setup, Exabeam also runs an ongoing classifier that classifies assets as workstations or servers based on their behavior.

#### Network Zones

Network zones are internal network locations defined by the organization rather than a physical place. Zones can be cities, business units, buildings, or even specific rooms. For example, "Atlanta" can refer to a network zone within an organization rather than the city itself (all according to an organization's preference). Administrators can upload information regarding network zones for their internal assets via CSV or add manually one at a time.

## Asset Groups

Asset Groups are a collection of assets that perform the same function in the organization and need to be treated as a single entity from an anomaly detection perspective. An example of an asset group would be a collection of Exchange Servers. Grouping them this way is useful to our modeling processing because it allows us to treat an asset group as a single entity, reducing the amount of false positives that are generated when users connect to multiple servers within that group. As a concrete example, if a user regularly connects to email exchange server #1 then Exabeam builds a baseline that says this is their normal behavior. But exchange servers are often load-balanced, and if the user then connects to email exchange server #2 we can say that this is still normal behavior for them because the exchange servers are one Asset Group. Other examples of asset groups are SharePoint farms, or Virtual Desktop Infrastructure (VDI).

## Universal Role-Based Access

Universal role-based access centralizes user identity and access management (IAM) for applications across the entire Exabeam Security Operations Platform. If you are a new Advanced Analytics customer, you should set up IAM exclusively with universal role-based access. For information on configuring universal role-based access, refer to [Universal Role-Based Access](#) in the *Exabeam Security Operations Platform Administration Guide*.

For information on migrating from legacy authentication to universal role-based access, see [Migrate to Universal Role-Based Access](#).



### NOTE

Existing customers can continue to use legacy role-based access control until they are prepared to migrate to universal role-based access. For information, see [Legacy Role-Based Access Control](#).

## Migrate to Universal Role-Based Access

If you are an existing Exabeam Security Operations Platform customer, you are encouraged to migrate from the individualized identity and access management (IAM) of your Exabeam products to universal role-based access. You can continue to use legacy authentication until your organization is prepared to migrate.

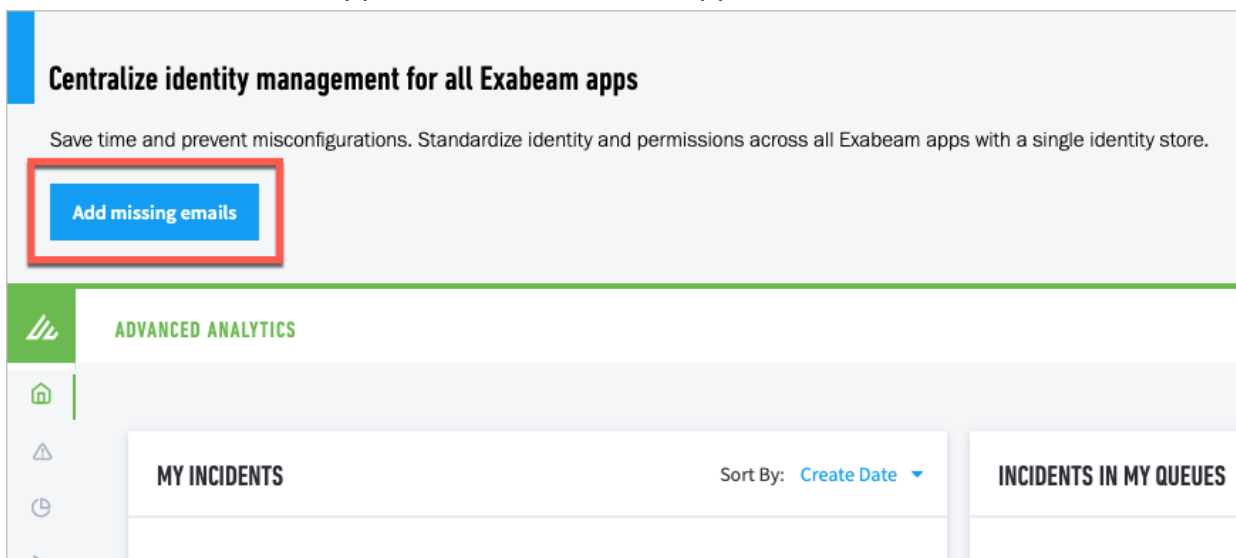


### WARNING

Migration to universal role-based access cannot be reversed.

## Prerequisites

- Unlike the legacy IAM, universal role-based access requires a unique email address for each user account. If your legacy account configurations do not include email addresses, you need to be prepared to add email addresses to the accounts to complete their migration.
- For customers using a third-party identity provider (IdP), you need to have an IdP administrator available to modify the IdP configuration for it to authenticate with universal role-based access.
- Universal role-based access does not support migration from LDAP directories. If LDAP is enabled, it needs to be disabled.
- You should see a notification about centralizing identity management up the upper-left side of the home page, along with a button to add missing emails. If you do not see this notification, try clearing your cookies or logging in with your browser in incognito/private mode. If the notification still does not appear, contact customer support.



To migrate to universal role-based access:

1. In the centralized identity management notification banner, click **Add missing emails**. The Add Missing Emails page opens. The page lists any accounts that do not have email addresses attached to them.
2. Click the name of each of the listed users to either add their unique email addresses or delete their accounts.

**!** **IMPORTANT**








You should delete the accounts of former employees and/or inactive users.



### Add missing emails

7 out of 10 of local users are missing emails

**i** We will carry over all your users' permissions. You may still need to re-configure your identity provider.


 <b>Cori Director</b> Roles: Administrator	Inactive	
<input type="text" value="Email"/>	<a href="#">Add email</a>	<a href="#">Delete user</a>
 <b>Esteban Security Engineer</b> Roles: Administrator	Inactive	
 <b>Hunter Tier 3 Analyst</b> Roles: Tier 3 Analyst	Inactive	
 <b>Tim Tier 1 Analyst</b> Roles: Tier 1 Analyst	Inactive	
 <b>Alice Auditor</b> Roles: Auditor	Inactive	
 <b>Debbie Data Privacy Officer</b> Roles: Data Privacy Officer	Inactive	
 <b>Carl Custom Role</b> Roles: Administrator, Auditor, Data Privacy Officer, Tier 1 Analyst, Tier 3 Analyst	Inactive	

[Continue later](#) [Next](#)

3. Click **Next** and repeat step 2 if needed.  
When all of the accounts are ready for migration, the Enable Unified Login page appears.

## Enable Unified Login

---



### Before you continue!

**1**

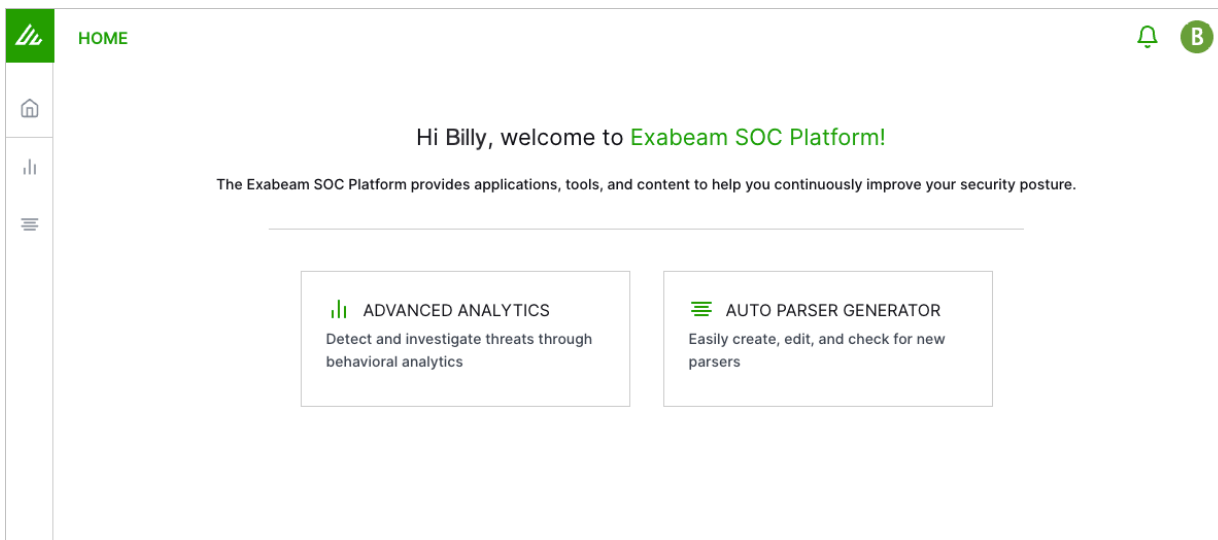
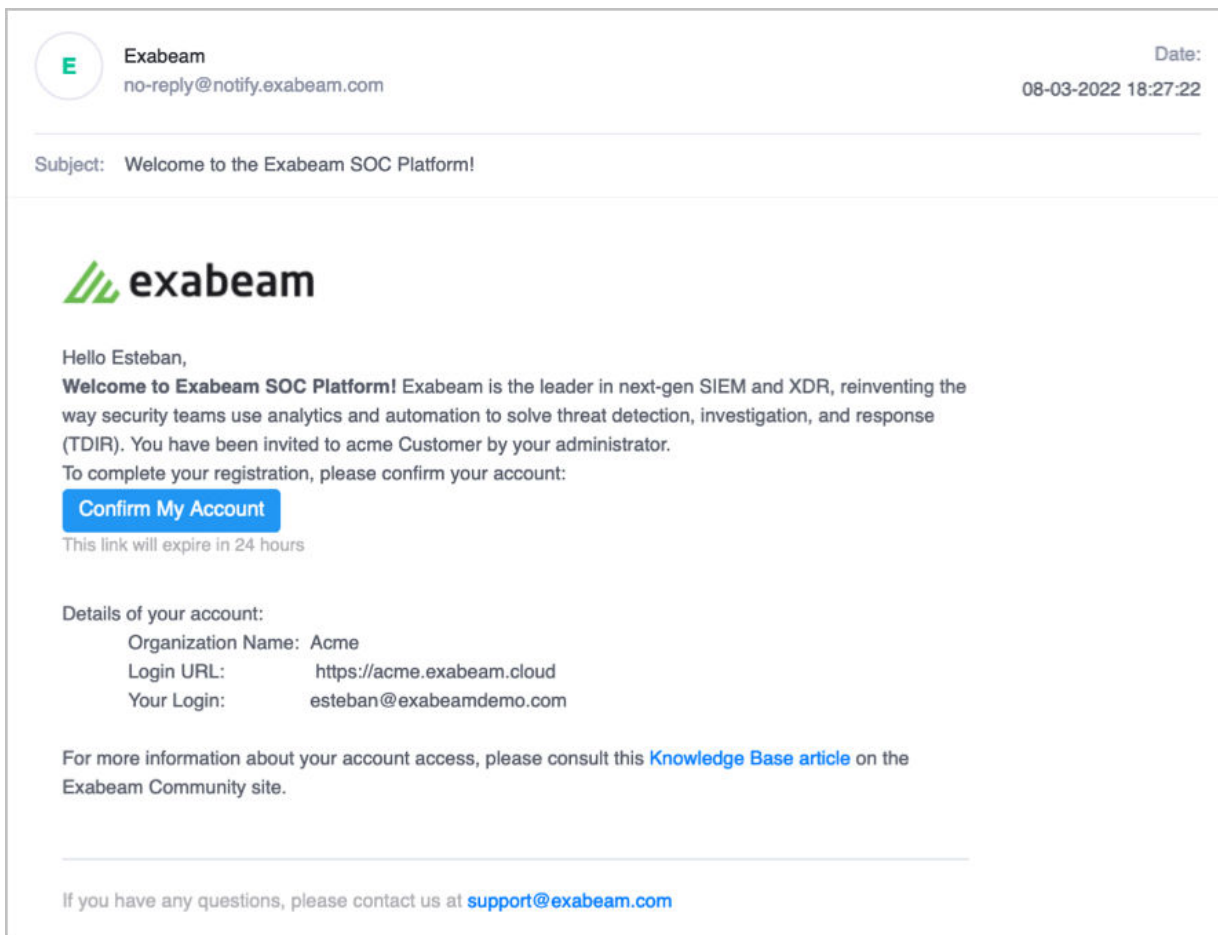
Unified logon affects all users on your team, and we will notify them by email with the change as soon as you enable it.


**2**

You will need to re-configure SAML after enabling. We suggest having your identity provider's settings page open before enabling.

[Continue later](#)   [Enable](#)

4. To proceed with the migration, click **Enable**.  
An email is sent to the migrated users to welcome them to the Exabeam Security Operations Platform. The email includes a link for them to confirm their accounts and set their passwords.



5. Verify the permissions for your users in the Exabeam Security Operations Platform. The Exabeam Security Operations Platform provides tailored roles for Advanced Analytics users such as Administrator, Tier 3 Analyst, and Data Privacy Officer. To see all roles, view the **Settings**  **> Roles** page.

For each user that should have access to Advanced Analytics, review and adjust any roles as desired. For more information, see [Administrative Users](#) in the Exabeam Security Operations Platform Guide.

## Legacy Role-Based Access Control

With its introduction in Advanced Analytics i61, universal role-based access is the recommended method for identity and access management (IAM). If you are a new Advanced Analytics customer, you should set up IAM exclusively with universal role-based access. Existing customers can continue to use role-based access control until they are prepared to migrate to universal role-based access. For information on migrating, see [Migrate to Universal Role-Based Access](#).

### Role-Based Access Control

You can control the responsibilities and activities of your SOC team members with Role-Based Access Control (RBAC). To tailor access, you can assign local users, LDAP users, or SAML authenticated users one or more roles within Exabeam.

The responsibilities of those roles are determined by the permissions the role allows. If users are assigned more than one role, that user receives the permissions of each role.



#### NOTE

If a user is assigned multiple roles with conflicting permissions, Exabeam enforces the role having more permission. For example, if a role with read-only permissions and a role with full permission are both assigned to a user, then the user will have full permission.

To access the Roles page, navigate to **Settings > User Management > Roles**.

### Manage Users

Understand the difference between Roles and Users. Configure the analysts that have access to the Exabeam User Interface, add the analyst's information, assign them roles, and set up user permissions and access based on your organization's needs.

#### Users

Users are the analysts that have access to the Exabeam UI to review and investigate activity. These analysts have specific roles, permissions, and can be assigned Exabeam objects within the platform. They also have the ability to accept sessions. Exabeam supports local authentication or authentication against an LDAP server.

#### Default Roles

Advanced Analytics includes preconfigured default roles with standard permission sets for common user types, such as analysts, administrators, and auditors.



#### NOTE

Default roles cannot be edited or deleted.

### Administrator

This role is intended for administrative access to Exabeam. Users assigned to this role can perform administrative operations on Exabeam, such as configuring the appliance to fetch logs from the

SIEM, connecting to Active Directory to pull in contextual information, and restarting the analytics engine. The default admin credential belongs to this role. This is a predefined role provided by Exabeam and cannot be deleted.

Default permissions include:

Permission	Description
Manage Users and Context Sources	Manage users and roles in the Exabeam Security Intelligence Platform, as well as the context sources used to enhanced the logs ingested (e.g. assets, peer groups, service accounts, executives).
Manage context tables	Manage users, assets or other objects within Context Tables.
Manage Content Packages	Users add/remove/configure content packages for automatic installation.
View Metrics	View the IR Metrics page.
Manage Data Ingest	Configure log sources and feeds and email-based ingest.
Add IR comments	Add IR comments.
Upload Custom Services	Upload custom actions or services.
Create incidents	Create incidents.
Delete incidents	Delete incidents.
Manage Custom Services and Packages	User can manage custom services and related packages
Manage Data Ingest	Configure log sources and feeds and email-based ingest.
Manage ingest rules	Add, edit, or delete rules for how incidents are assigned, restricted, and prioritized on ingest.
Manage Queues	Create, edit, delete, and assign membership to queues
Manage Templates	Create, edit, or delete playbook templates.
Manage Triggers	Create, update, or delete playbook triggers.
Run Actions	Launch individual actions from the user interface.
Manage Bi-directional Communication	Configure inbound and outbound settings for Bi-Directional Communications.
Manage Incident Configuration	Users can manage the Incident Configurations including Incident Types, Fields, Layouts, and Checklists.
Manage Playbooks	Create, update, or delete playbooks.
Manage Services	Configure, edit, or delete services (3rd party integrations).
Run Playbooks	Run a playbook manually from the workbench.
Reset Incident Workbench	User can reset incident workbench
All Admin Ops	Perform all Exabeam administrative operations such as configuring the appliance, connecting to the log repository and Active Directory, setting up log feeds, managing users and roles that access the Exabeam UI, and performing system health checks.
View comments	View comments.
View health	View health.
View Raw Logs	View the raw logs that are used to built the events on AA timeline
View Rules	View configured rules that determine how security events are handled
View API	View API.
View incidents	View incidents.
View Metrics	View the IR Metrics page.
Edit incidents	Edit an incident's fields, edit tasks, entities & artifacts.

Permission	Description
Manage Rules	Create/Edit/Reload rules that determine how security events are handled
Bulk edit	Users can edit multiple incidents at the same time.
Search Incidents	Can search keywords in IR via the search bar.
Basic Search	Perform basic search on the Exabeam homepage. Basic search allows you to search for a specific user, asset, session, or a security alert.
View Restricted Incidents	View incidents restricted to other users or groups.

### *Auditor*

Users assigned to this role have only view privileges within the Exabeam UI. They can view all activities within the Exabeam UI, but cannot make any changes such as add comments or approve sessions. This is a predefined role provided by Exabeam.

Default permissions include:

Permission	Description
View Comments	View comments
View Activities	View all notable users, assets, sessions, and related risk reasons in the organization.
View Global Insights	View the organizational models built by Exabeam. The histograms that show the normal behavior for all entities in the organization can be viewed.
View Executive Info	View the risk reasons and the timeline of the executive users in the organization. You will be able to see the activities performed by executive users along with the associated anomalies.
View Incidents	View incidents.
View Infographics	View all the infographics built by Exabeam. You will be able to see the overall trends for the organization.
View Insights	View the normal behaviors for specific entities within the organization. The histograms for specific users and assets can be viewed.
Search Incidents	Can search keywords in Incident Responder via the search bar.
Basic Search	Perform basic search on the Exabeam homepage. Basic search allows you to search for a specific user, asset, session, or a security alert.
View Search Library	View the Search Library provided by Exabeam and the corresponding search results associated with the filters.
Threat Hunting	Perform threat hunting on Exabeam. Threat hunting allows you to query the platform across a variety of dimension such as find all users whose sessions contain data exfiltration activities or a malware on their asset.
View Restricted Incidents	View incidents restricted to other users or groups.

### *Tier 1 Analyst*

Users assigned to this role are junior security analysts or incident desk responders who supports the day-to-day enterprise security operation and monitoring. This type of role will not be authorized to make any changes to Exabeam system except for making user, session and lockout comments. Users in this role cannot approve sessions or lockout activities. This is a predefined role provided by Exabeam.

Default permissions include:

## Set Up Authentication and Access Control

Permission	Description
Add Advanced Analytics Comments	Add comments for the various entities (users, assets and sessions) within Exabeam.
Add Incident Responder Comments	Add Incident Responder comments.
Create Incidents	Create incidents.
Run Playbooks	Run a playbook manually from the workbench.
Run Actions	Launch individual actions from the user interface.
View comments	View comments.
View Global Insights	View the organizational models built by Exabeam. The histograms that show the normal behavior for all entities in the organization can be viewed.
View incidents	View incidents.
View Infographics	View all the infographics built by Exabeam. You will be able to see the overall trends for the organization.
View Activities	View all notable users, assets, sessions and related risk reasons in the organization.
View Executive Info	View the risk reasons and the timeline of the executive users in the organization. You will be able to see the activities performed by executive users along with the associated anomalies.
View Insights	View the normal behaviors for specific entities within the organization. The histograms for specific users and assets can be viewed.

### *Tier 3 Analyst*

Users assigned to this role will be performing more complex investigations and remediation plans. They can review user sessions, account lockouts, add comments, approve activities and perform threat hunting. This is a predefined role provided by Exabeam and cannot be deleted.

Default permissions include:

Permission	Description
Add Advanced Analytics Comments	Add comments for the various entities (users, assets and sessions) within Exabeam.
Add Incident Responder Comments	Add Incident Responder comments.
Upload Custom Services	Upload custom actions or services.
Create incidents	Create incidents.
Delete incidents	Delete incidents.
Manage Playbooks	Create, update, or delete playbooks.
Manage Queues	Create, edit, delete, and assign membership to queues
Manage Services	Configure, edit, or delete services (3rd party integrations).
Manage Triggers	Create, update, or delete playbook triggers.
Run Actions	Launch individual actions from the user interface.
Manage Bi-directional Communication	Configure inbound and outbound settings for Bi-Directional Communications.
Manage Data Ingest	Configure log sources and feeds and email-based ingest.
Manage ingest rules	Add, edit, or delete rules for how incidents are assigned, restricted, and prioritized on ingest.
Manage Templates	Create, edit, or delete playbook templates.
Run Playbooks	Run a playbook manually from the workbench.

Permission	Description
View Activities	View all notable users, assets, sessions and related risk reasons in the organization.
View Comments	View comments.
View Executive Info	View the risk reasons and the timeline of the executive users in the organization. You will be able to see the activities performed by executive users along with the associated anomalies.
View Global Insights	View the organizational models built by Exabeam. The histograms that show the normal behavior for all entities in the organization can be viewed.
View Infographics	View all the infographics built by Exabeam. You will be able to see the overall trends for the organization.
View Rules	View configured rules that determine how security events are handled.
View Insights	View the normal behaviors for specific entities within the organization. The histograms for specific users and assets can be viewed.
Bulk Edit	Users can edit multiple incidents at the same time.
Delete entities and artifacts	Users can delete entities and artifacts.
Manage Watchlist	Add or remove users from the Watchlist. Users that have been added to the Watchlist are always listed on the Exabeam homepage, allowing them to be scrutinized closely.
Approve Lockouts	Accept account lockout activities for users. Accepting lockouts indicates to Exabeam that the specific set of behaviors for that lockout activity sequence are whitelisted and are deemed normal for that user.
Accept Sessions	Accept sessions for users. Accepting sessions indicates to Exabeam that the specific set of behaviors for that session are whitelisted and are deemed normal for that user.
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p><b>⚠ WARNING</b></p> <p>This permission should be given only sparingly, if at all. Accepting sessions is not recommended. The best practice for eliminating unwanted alerts is through tuning the rules and/or models.</p> </div>	
Edit incidents	Edit an incident's fields, edit entities & artifacts.
Sending incidents to Incident Responder	Send incidents to Incident Responder.
Basic Search	Perform basic search on the Exabeam homepage. Basic search allows you to search for a specific user, asset, session, or a security alert.
Search Incidents	Can search keywords in IR via the search bar.
Threat Hunting	Perform threat hunting on Exabeam. Threat hunting allows you to query the platform across a variety of dimensions such as find all users whose sessions contain data exfiltration activities or a malware on their asset.
Manage Search Library	Create saved searches as well as edit them.
View Search Library	View the Search Library provided by Exabeam and the corresponding search results associated with the filters.

### *Data Privacy Officer*

This role is needed only when the data masking feature is turned on within Exabeam. Users assigned to this role are the only users that can view personally identifiable information (PII) in an unmasked form. They can review user sessions, account lockouts, add comments, approve activities and perform threat hunting. This is a predefined role provided by Exabeam.

See [Mask Data Within the Advanced Analytics UI](#) for more information on this feature.

Default permissions include:




Permission	Description
Add Advanced Analytics Comments	Add comments for the various entities (users, assets and sessions) within Exabeam.
View comments	View comments.
View Global Insights	View the organizational models built by Exabeam. The histograms that show the normal behavior for all entities in the organization can be viewed.
View incidents	View incidents.
View Infographics	View all the infographics built by Exabeam. You will be able to see the overall trends for the organization.
View Activities	View all notable users, assets, sessions and related risk reasons in the organization.
View Executive Info	View the risk reasons and the timeline of the executive users in the organization. You will be able to see the activities performed by executive users along with the associated anomalies.
View Insights	View the normal behaviors for specific entities within the organization. The histograms for specific users and assets can be viewed.
Manage Watchlist	Add or remove users from the Watchlist. Users that have been added to the Watchlist are always listed on the Exabeam homepage, allowing them to be scrutinized closely.
Sending incidents to Incident Responder	Sending incidents to Incident Responder.
Accept Sessions	Accept sessions for users. Accepting sessions indicates to Exabeam that the specific set of behaviors for that session are whitelisted and are deemed normal for that user.
Basic Search	Perform basic search on the Exabeam homepage. Basic search allows you to search for a specific user, asset, session, or a security alert.
Threat Hunting	Perform threat hunting on Exabeam. Threat hunting allows you to query the platform across a variety of dimensions such as find all users whose sessions contain data exfiltration activities or a malware on their asset.
Manage Search Library	Create saved searches as well as edit them
View Search Library	View the Search Library provided by Exabeam and the corresponding search results associated with the filters.
View Unmasked Data (PII)	Show all personally identifiable information (PII) in a clear text form. When data masking is enabled within Exabeam, this permission should be enabled only for select users that need to see PII in a clear text form.
View Restricted Incidents	View incidents restricted to other users or groups.

### Create a Custom User Role

If the preconfigured default roles do not meet all your needs, you can create your own roles.

To add a new role:

1. Click **Settings** , and then click **Analytics > Exabeam User Management > Roles**. The User Management page opens with the Roles tab selected.
2. Click **Create Role**.

## Set Up Authentication and Access Control

**SETTINGS**

**USER MANAGEMENT**

ROLES | **USERS** | CONFIGURE SAML | LDAP AUTHENTICATION

Configure roles for users to perform various actions on the Exabeam User Interface. [+ Create Role](#)

Default Roles	Available Permissions for Administrator Role
<b>Administrator</b> (5 Users)	Users assigned to this role can perform administrative operations on Exabeam such as configuring the appliance to fetch logs from the SIEM, connecting to Active Directory to pull in contextual information and restarting the analytics engine. The default admin credential belongs to this role. This is a predefined role provided by Exabeam and cannot be deleted.
<b>Auditor</b> (2 Users)	Filter: All Products ▾
<b>Data Privacy Officer</b> (1 User)	▼ Core
<b>Tier 1 Analyst</b> (1 User)	ADMINISTRATION
<b>Tier 3 Analyst</b> (3 Users)	<b>Manage Users and Context Sources</b> Manage users and roles in the Exabeam Security Intelligence Platform, as well as the context sources used to enhance the logs ingested (e.g. assets, <a href="#">More</a> )
<b>Custom Roles</b>	<b>Manage Content Packages</b> Users add/remove/configure content packages for automatic installation.
	<b>Manage context tables</b> Manage users, assets or other objects within Context Tables.
	▼ ADVANCED ANALYTICS
	COMMENTS

3. In the **Create a new role** dialog box, enter a **Role name** and ensure that the **Advanced Analytics** tab is selected.

**Create a new role**

Demo Role

Demo description.

**ADVANCED ANALYTICS** CORE

Search

**COMMENTS** [Select All](#) | [Deselect All](#)

**Add AA Comments**  
Add comments for the various entities (users, assets and sessions) within [More](#)

**Add IR comments**  
Add IR comments.

**CANCEL** **SAVE**

4. Select the permissions that you want to assign to the role.  
To quickly find specific permissions, use the search bar. Otherwise, scroll down to view the available permissions.

### Create a new role

☰ DELETE
Select All | Deselect All

**Delete incidents**  
Delete incidents.

☰ MANAGE
Select All | Deselect All

<input checked="" type="checkbox"/> <b>Manage Bi-directional Communication</b> Configure inbound and outbound settings for Bi-Directional <a href="#">More</a>	<input type="checkbox"/> <b>Manage Custom Services and Packages</b> User can manage custom services and related packages
<input checked="" type="checkbox"/> <b>Manage Incident Configuration</b> Users can manage the Incident Configurations including Incident Types, <a href="#">More</a>	<input checked="" type="checkbox"/> <b>Manage Data Ingest</b> Configure log sources and feeds and email-based ingest.
<input type="checkbox"/> <b>Manage ingest rules</b>	<input type="checkbox"/> <b>Manage Playbooks</b>

CANCEL
SAVE

5. When you are finished, click **Save**.

The new role appears in the list of Custom Roles on the Roles tab of the User Management page. From the Users tab, you can assign the role to both new and existing users.

#### **Add an Exabeam User**

1. Navigate to **Settings > Exabeam User Management > Users**.
2. Click **Add User**.
3. Fill the new user fields and select role(s), and then click **SAVE**.

Your newly created user should appear in the **Users** UI.

#### **User Password Policies**

Exabeam users must adhere to the following default password security requirements:

- Passwords must:
  - Be between 8 to 32 characters
  - Contain at least one uppercase, lowercase, numeric, and special character

- Contain no blank space
- User must change password every 90 days
- New passwords cannot match last 5 passwords
- SHA256 hashing is applied to store passwords
- Only administrators can reset passwords and unblock users who have been locked out due to too many consecutive failed logins

### Third-Party Identity Provider Configuration

Exabeam supports integration with SAML 2.0 compliant third-party identity providers (IdPs) for single sign-on (SSO), multi-factor authentication, and access control. Once an IdP is added to your product, you can make IdP authentication mandatory for users to log in to the product, or you can allow users to log in through either the IdP or local authentication.



#### NOTE

You can add multiple IdPs to your Exabeam product, but only one IdP can be enabled at a time.

### **Add Exabeam to Your SAML Identity Provider**

This section provides instructions for adding Exabeam to your SAML 2.0 compliant identity provider (IdP). For detailed instructions, refer to your IdP's user guide.

### *General Instructions*

The exact procedures for configuring IdPs to integrate with Exabeam vary between vendors, but the general tasks that need to be completed include the following (not necessarily in the same order):

1. Begin the procedure to add a new application in your IdP for Exabeam (if needed, refer to your IdP's user guide for instructions).
2. In the appropriate configuration fields, enter the Exabeam Entity ID and the Assertion Consumer Service (ACS) URL as shown in the following:

**Entity ID:**

**ACS URL:**

**!** **IMPORTANT**

Make sure that you replace `<exabeam_primary_host>` with the IP address or domain name of your primary host. The only acceptable values for `<identity_provider>` are the following:

- **adfs**
- **google**
- **ping**
- **okta**
- **others**

If you are using Microsoft AD FS, Google IdP, Ping Identity, or Okta, enter the corresponding value from the preceding list. For all other IdPs, enter **others**. All of the values are case sensitive.

3. In the attribute mapping section, enter descriptive values for the following IdP user attributes:

- Email address
- First name
- Last name
- Group
- Username (this attribute is optional)

**📌** **NOTE**

The actual names of these user attributes may vary between the different IdPs, but each IdP should have the corresponding attributes.

For example, if **Primary email** is the user email attribute in your IdP, you could enter **EmailAddress** as the descriptive value. The following is an example of a completed attribute map in Google IdP:

**SAML attribute mapping** ^

**Attributes** Add and select user fields in Google Directory, then map them to service provider attributes. Attributes marked with \* are mandatory. [Learn more](#)

Google Directory attributes	→	App attributes	
Basic Information > <b>Primary email</b>	→	<b>EmailAddress</b>	✕
Basic Information > <b>First name</b>	→	<b>FirstName</b>	✕
SAML Group > <b>saml_group</b>	→	<b>Group</b>	✕
Basic Information > <b>Last name</b>	→	<b>LastName</b>	✕

[ADD MAPPING](#)

[CANCEL](#)   [SAVE](#)

**! IMPORTANT**

When you [Configure Exabeam for SAML Authentication](#), you need to use the same descriptive values to map the Exabeam query attributes with the corresponding IdP user attributes.

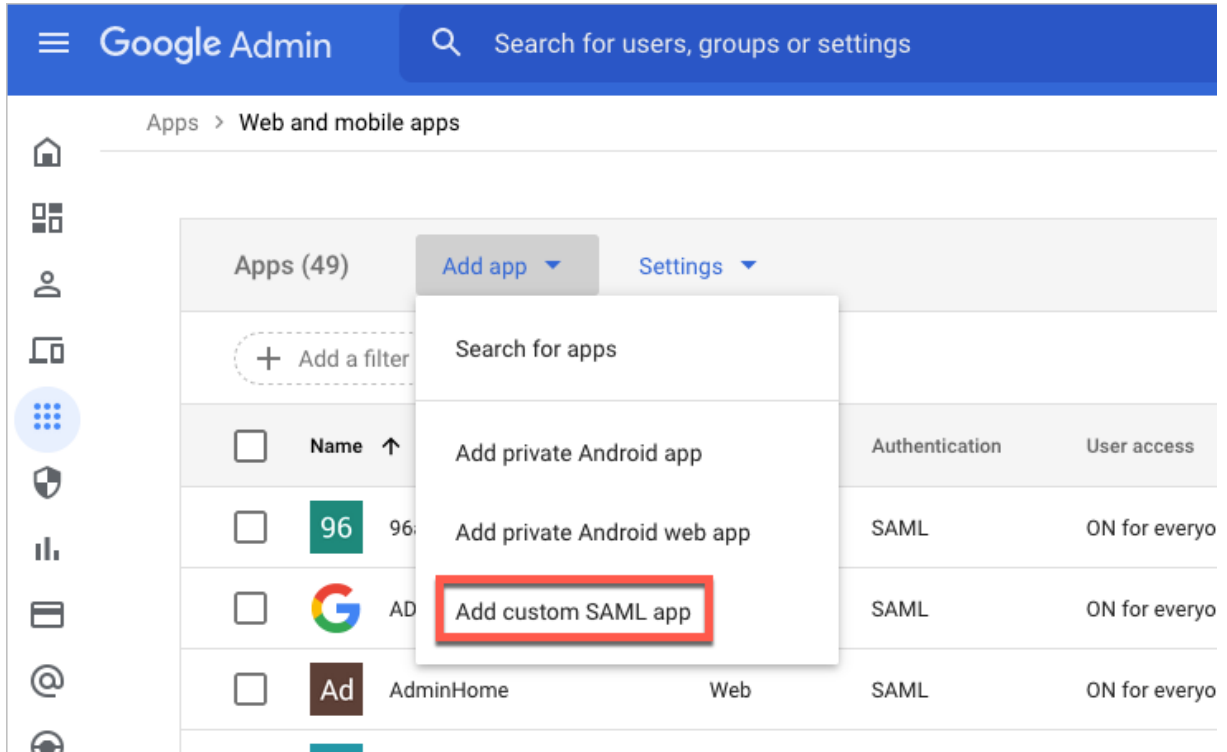
4. Complete any additional steps in your IdP that are necessary to finish the configuration. Refer to your IdP user guide for details.
5. Copy the IdP's connection details and download the IdP certificate or, if available, download the SAML metadata file.

**📌 NOTE**

You need either the connection details and the IdP certificate or the SAML metadata file to complete the integration in Exabeam.

### Google IdP

1. From the main menu on the left, select **Apps** and then click **Web and mobile apps**.
2. From the **Add app** drop-down menu, click **Add custom SAML app**.



The App Details section opens.

3. In the **App name** field, enter a name.
4. Under **App icon**, click the blue circle, navigate to an image file that can be used as an icon and click to upload it.




### App details

Enter details for your custom SAML app. This information is shared with app users. [Learn more](#)

App name  
Exabeam


Description

App icon  
Attach an app icon. Maximum upload file size: 4 MB




CANCEL CONTINUE

5. Click **Continue**.  
The Google Identity Provider Details section opens.
6. Click **Download IdP Metadata**.

 **NOTE**  
The IdP metadata file needs to be uploaded to Exabeam when you [Configure Exabeam for SAML Authentication](#).

7. Click **Continue**.  
The Service Provider Details section opens.
8. Enter the **ACS URL** and **Entity ID** as shown in the following:  
**ACS URL:**  
**Entity ID:**

 **NOTE**  
Make sure that you replace `<exabeam_primary_host>` with the IP address or domain name of your primary host.

9. Click **Continue**.  
The Attribute Mapping section opens.
10. Click **Add Mapping**, and then from **Select field** drop-down menu, select **Primary email**.

11. Repeat the previous step for each of the following attributes:
  - Primary email
  - First name
  - Last name
  - Group
12. In the **App attributes** fields, enter descriptive values for the attributes. For example, for the **Primary email** attribute, you could enter **EmailAddress** for the descriptive value. The following is an example of a completed attribute map:

SAML attribute mapping ^

---

**Attributes** Add and select user fields in Google Directory, then map them to service provider attributes. Attributes marked with \* are mandatory. [Learn more](#)

Google Directory attributes	→	App attributes	
Basic Information > <b>Primary email</b>	→	<b>EmailAddress</b>	✕
Basic Information > <b>First name</b>	→	<b>FirstName</b>	✕
SAML Group > <b>saml_group</b>	→	<b>Group</b>	✕
Basic Information > <b>Last name</b>	→	<b>LastName</b>	✕

[ADD MAPPING](#)

---

CANCEL SAVE

**! IMPORTANT**

When you [Configure Exabeam for SAML Authentication](#), you need to use the same descriptive values to map the Exabeam query attributes with the corresponding IdP user attributes.

13. Click **Continue**.  
The details page opens for your Exabeam app.
14. In the User Access panel, click the **Expand panel** icon to begin assigning the appropriate organizational units and groups to your Exabeam app and manage its service status.

## Set Up Authentication and Access Control

Apps > Web and mobile apps > Exa-Docs-Test

SAML

**Ex** Exa-Docs-Test

- TEST SAML LOGIN
- DOWNLOAD METADATA
- EDIT DETAILS
- DELETE APP

### User access

To make the managed app available to select users, choose a group or organizational unit. [Learn more](#)

[View details](#)

OFF for everyone

### Service provider details

Certificate	ACS URL	Entity ID
Google_2026-5-22-4359_SAML2_0 (Expires May 22, 2026)	https://exadoc-e2e.dl.exabeam.com:443 /api/auth/saml2/google/handle-assertion	https://exadoc-e2e.dl.exabeam.com:443 /api/auth/saml2/google/login

### SAML attribute mapping

Map Google directory user profile fields to SAML service provider attributes.

EmailAddress	FirstName	Group
Basic Information > Primary email	Basic Information > First name	SAML Group > saml_group
LastName		
Basic Information > Last name		

You are now ready to [Configure Exabeam for SAML Authentication](#).

## Azure AD



### NOTE

The following instructions include procedural information for configuring both Azure AD and Exabeam to complete the IdP setup.

1. Log in to Microsoft Azure and navigate to **Enterprise Applications**.
2. Create an Exabeam enterprise application by doing the following:
  - a. Click **New application**, and then click **Create your own application**. The Create your own application dialog box appears.
  - b. In the **What's the name of your app** field, type a name for the app (for example, "Exabeam-SAML").

**Create your own application** ×

What's the name of your app?

Exabeam-SAML

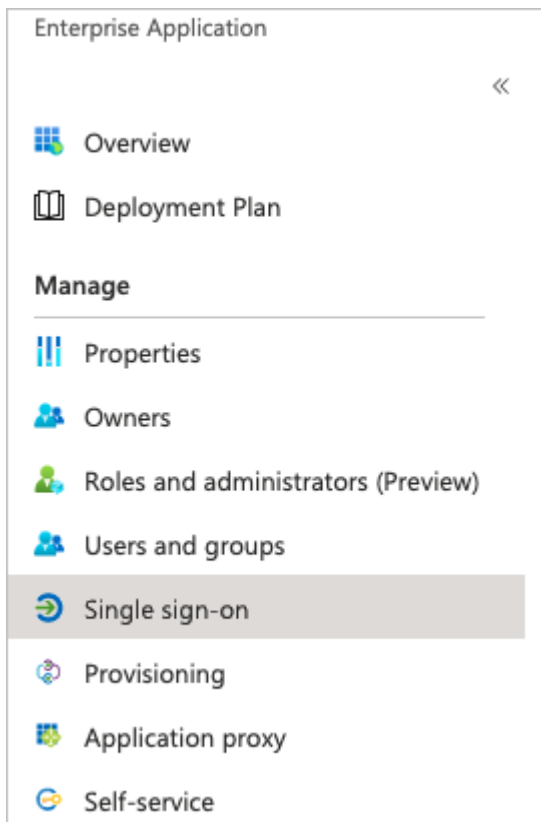
What are you looking to do with your application?

Configure Application Proxy for secure remote access to an on-premises application

Register an application to integrate with Azure AD (App you're developing)

Integrate any other application you don't find in the gallery (Non-gallery)

- c. Select **Integrate any other application you don't find in the gallery (Non-gallery)**.
  - d. Click **Create**.
3. On the Enterprise Application page, locate and click the application that you added in step 2.
4. In the Manage section, click **Single sign-on**.



5. Click the **SAML** tile.


## Set Up Authentication and Access Control

Home > Exadev Directory > Enterprise applications > Browse Azure AD Gallery > Exabeam-saml2

The screenshot shows the Azure AD portal interface for configuring the 'Exabeam-saml2' application. The breadcrumb path is 'Home > Exadev Directory > Enterprise applications > Browse Azure AD Gallery > Exabeam-saml2'. The page title is 'Exabeam-saml2 | Single sign-on'. The left navigation pane includes 'Overview', 'Deployment Plan', 'Manage' (with sub-items: Properties, Owners, Roles and administrators (Preview), Users and groups, Single sign-on, Provisioning, Application proxy, Self-service), and 'Security' (with sub-item: Conditional Access). The main content area is titled 'Select a single sign-on method' and features three options: 'Disabled' (with a red 'X' icon and text: 'Single sign-on is not enabled. The user won't be able to launch the app from My Apps.'), 'SAML' (with a puzzle piece icon and text: 'Rich and secure authentication to applications using the SAML (Security Assertion Markup Language) protocol.'), and 'Password-based' (with a lock icon and text: 'Password storage a web browser exten'). The 'SAML' option is highlighted with a red border. Below these options is a 'Linked' section with a link icon and text: 'Link to an application in My Apps and/or Office 365 application launcher.'

6. In the Basic SAML Configuration box (1), click **Edit**, and then do the following:

a.

 **NOTE**

Make sure that you replace `<exabeam_primary_host>` with the IP address or domain name of your primary host.

b.

 **NOTE**

Make sure that you replace `<exabeam_primary_host>` with the IP address or domain name of your primary host.

c. Click **Save**.

7. In the User Attributes & Claims box (2), click **Edit**, and then map the Azure objects to your Exabeam field attributes.

a. Click the row for the user.mail claim.

The Manage claim dialog box appears.

b. In the **Name** field, type the name of the appropriate Exabeam field attribute.

- c. If needed, clear the value in the **Namespace** field to leave it empty.
- d. Click **Save**.
- e. Repeat steps a through d as needed for the following claims:
  - user.givenname
  - user.userprincipalname
  - user.surname
- f. Click **Add a group claim**.

- g. In the Group Claims dialog box, select **Groups assigned to the application**.
- h. From the **Source attribute** drop-down list, select **Group ID**.
- i. In the Advanced Options section, select the checkbox for **Customize the name of the group claim**.
- j. In the **Name (required)** field, type **Group**.

## Group Claims ✕

Manage the group claims used by Azure AD to populate SAML tokens issued to your app

Which groups associated with the user should be returned in the claim?

None  
 All groups  
 Security groups  
 Directory roles  
 Groups assigned to the application

Source attribute \*

Group ID
▾

### Advanced options

Customize the name of the group claim

Name (required)

Group

Namespace (optional)

Emit groups as role claims ⓘ

- k. Click **Save**.  
 The Group claim is added to the User Attributes & Claims box.

User Attributes & Claims ✎ Edit

EmailAddress	user.mail
FirstName	user.givenname
LastName	user.surname
Username	user.userprincipalname
Group	user.groups
Unique User Identifier	user.userprincipalname

8. In the SAML Signing Certificate box (3), download the **Federation Metadata XML** certificate to upload to Exabeam.

**3** SAML Signing Certificate  Edit

Status	Active
Thumbprint	749BE03BC97A714F8DEE9EB2BF8F5D1BCD8C5B26
Expiration	3/31/2024, 10:35:29 AM
Notification Email	jammy@exadevdirectory.onmicrosoft.com
App Federation Metadata Url	<input type="text" value="https://login.microsoftonline.com/428b791c-67a8..."/>
Certificate (Base64)	<a href="#">Download</a>
Certificate (Raw)	<a href="#">Download</a>
Federation Metadata XML	<a href="#">Download</a>

9. In Exabeam, navigate to **Settings > User Management > Configure SAML**, and then click **Add Identity Provider**.  
The New Identity Provider dialog box appears.
10. From the **SAML Provider** drop-down list, select **Custom/Generic IdP**.
11. Under SSO Configuration, select **Upload the XML metadata filed provided by your IdP**, and then choose the Federation Metadata XML file that was downloaded in step 8.
12. In the **Name of IdP** field, type a name (for example, "Azure").
13. In the **Upload IdP logo** field, click **Choose File**, and then select a PNG file of the logo that you want to use.

**NOTE**

The PNG logo file size cannot exceed 1 MB.



## NEW IDENTITY PROVIDER

Use this Identity Provider for SSO  
Only one IdP can be enabled at a time.

IdP Enabled

**SAML Provider\***

Custom/Generic IdP...

**SSO Configuration\***

Upload the XML metadata file provided by your IdP

ExaSAML-(2).xml ✓ [CHOOSE FILE](#)

Configure SSO manually


**Authentication Method ?**

Select your authentication method or leave blank for default

**Name of IdP\* ?**

Azure

**Current Icon Logo**



[CANCEL](#) [SAVE](#)

- In the Query Attributes section, enter the appropriate IdP attribute values for each field that you defined in step 7.

**! IMPORTANT**

The IdP attribute values must match the values that you defined in step 7.

### Query Attributes\*

Exabeam Attributes	IdP Attributes
Email Address	<input type="text" value="EmailAddress"/>
Username	<input type="text" value="Username"/>
First Name	<input type="text" value="FirstName"/>
Last Name	<input type="text" value="LastName"/>
Group	<input type="text" value="Group"/>

15. Click **Save**.

Azure now appears as an identity provider in the Configure SAML tab of the User Management page, and a Group Mappings section also appears.

**USER MANAGEMENT**      ROLES    USERS    LDAP AUTHENTICATION    **CONFIGURE SAML**

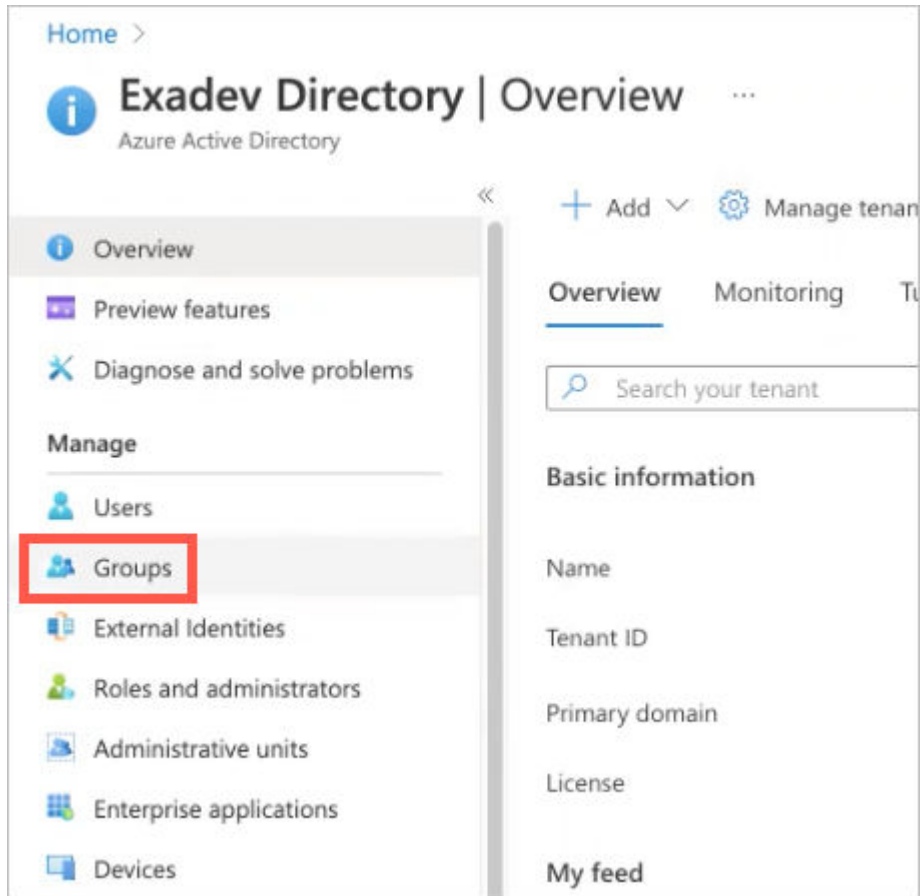
**SAML Status**  
Allowed  
Users can login with their SAML or Exabeam credentials.  
[EDIT](#)

**Identity Providers** [ADD NEW](#)

Name	Status
Azure	Enabled

**Group Mappings**  
Map your SAML groups to Exabeam user roles  
[ADD GROUP](#)

16. To map a SAML group to Exabeam user roles, do the following:
  - a. On the home page of Azure, click **Groups**.



- b. From the **Object Id** column, copy the ID for the Azure group that you want to map.

	Name	Object Id	Group Type
<input type="checkbox"/>	<b>AD</b> AAD DC Administrators	faf52ee7-347a-4cfc-ae5a-b95ff42c8e54	Security
<input type="checkbox"/>	<b>AD</b> ADSyncBrowse	a450df0b-90fd-4520-90c8-25f341e03178	Security
<input type="checkbox"/>	<b>AD</b> ADSyncOperators	718c6232-886f-4f5a-a4f4-47724a4ad301	Security
<input type="checkbox"/>	<b>AD</b> ADSyncPasswordSet	e005315e-4dd0-4cca-aa7b-7e76ef664b46	Security
<input type="checkbox"/>	<b>DE</b> DevOps	80b02c12-f58b-4172-827c-0dd6381cf906	Security
<input type="checkbox"/>	<b>DN</b> DnsAdmins	085c2161-a9ed-4c1b-9215-79944550be4c	Security
<input type="checkbox"/>	<b>DN</b> DnsUpdateProxy	36a92507-67b1-41e9-822c-3ccaaf4ec59f	Security

- c. In Exabeam, on the Configure SAML tab of the User Management page, click **Add Group**.  
The Edit Group Mapping dialog box appears.
- d. From the **Identity Provider** drop-down menu, select **Others**.
- e. In the **Group Name** field, paste the object ID that you copied in step b.

**EDIT GROUP MAPPING**

Identity Provider\*

Others

Group Name\*

e35f57ae-7ff1-4c2c-9e26-34d0ce162c1e

Exabeam User Roles\* ?

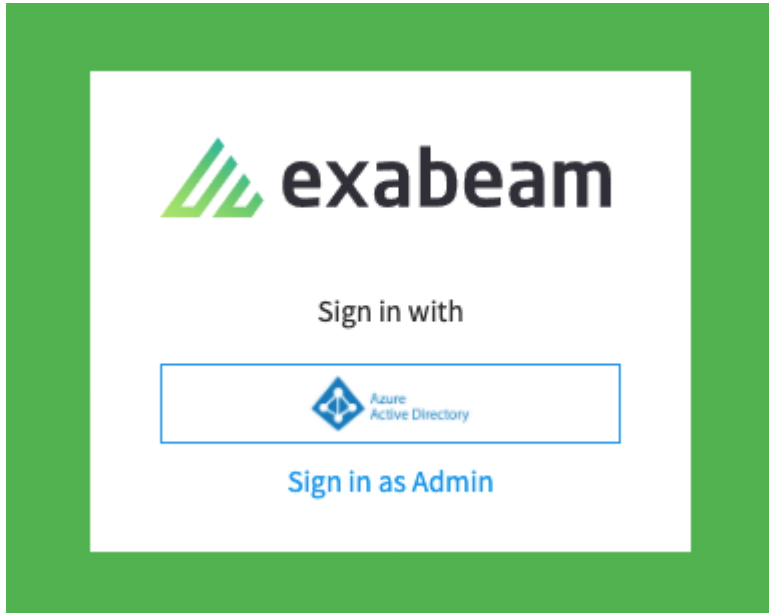
**Default Roles**

<input checked="" type="checkbox"/>	Administrator	▼
<input type="checkbox"/>	Tier 3 Analyst	▼
<input type="checkbox"/>	Tier 1 Analyst	▼
<input type="checkbox"/>	Auditor	▼
<input type="checkbox"/>	Data Privacy Officer	▼

**Custom Roles**

CANCEL SAVE


- f. Select the Exabeam User Roles that you want to assign to the group.
  - g. Click **Save**.
  - h. Repeat steps a through g for each Azure group that you want mapped to user roles.
17. To verify that Azure has been successfully configured, log out of Exabeam and look for the Azure Active Directory option on the sign-on screen.

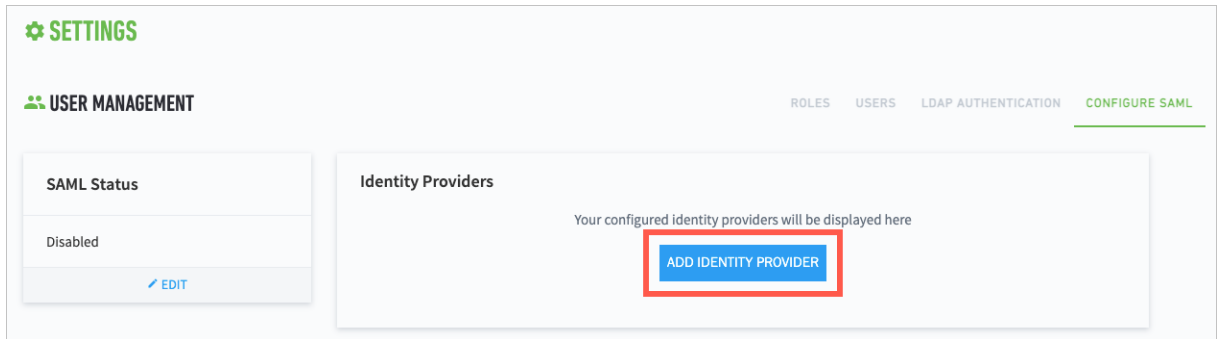


### Configure Exabeam for SAML Authentication

#### IMPORTANT

Before you begin this procedure, you need to [Add Exabeam to Your SAML Identity Provider](#).

1. Log in to your Exabeam product.
2. Navigate to **Settings**  > **Core** > **User Management** > **Configure SAML**.
3. Click **Add Identity Provider**.



4. From the **SAML Provider** drop-down menu, select your IdP.

#### NOTE

If your IdP is not listed, select **Custom/Generic IdP**.

## NEW IDENTITY PROVIDER

Use this Identity Provider for SSO IdP Disabled   
Only one IdP can be enabled at a time.

**SAML Provider\***

Select your IdP ▼

**SSO Configuration\***

Upload the XML metadata file provided by your IdP

No file chosen CHOOSE FILE

Configure SSO manually

5. With the information that you collected in step 5 of [Add Exabeam to Your SAML Identity Provider](#), do one of the following:
  - If you have an XML metadata file from your IdP, select **Upload the XML metadata provided by your IdP**, and then click **Choose File** to locate and upload the file from your computer.
  - If you do not have a metadata file, select **Configure SSO manually** and then do the following:
    1. Click **Choose File** to locate and upload the IdP certificate from your computer.

## NEW IDENTITY PROVIDER

Use this Identity Provider for SSO IdP Disabled   
 Only one IdP can be enabled at a time.

**SAML Provider\***

Okta

**SSO Configuration\***

Upload the XML metadata file provided by your IdP

Configure SSO manually

**IdP Certificate\***

No file chosen ! CHOOSE FILE

Please upload a file.

**Single Sign-on URL\***

Enter URL HTTP POST

**Single Log-Out URL**

Enter URL HTTP POST

**Redirect to URL after Log-Out**

2. In the **Single Sign-on URL** field, enter the appropriate URL, and then select either **HTTP POST** or **HTTP REDIRECT** as needed from the drop-down menu.
  3. *(Optional)* In the **Single Log-Out URL** and **Redirect to URL after Log-Out** fields, enter the appropriate URLs.
6. If you selected Custom/Generic IdP in the previous step, do the following:
- a. In the **Name of IdP** field, enter a name.
  - b. Under **Upload IdP Logo**, click **Choose File** to locate and upload an IdP logo image in PNG format.

The screenshot shows the 'SSO Configuration\*' interface. It has two radio buttons: 'Upload the XML metadata file provided by your IdP' (selected) and 'Configure SSO manually'. Below is a file upload field with 'No file chosen' and a 'CHOOSE FILE' button. An 'Authentication Method' dropdown menu is set to 'Select your authentication method or leave blank for default'. A red box highlights the 'Name of IdP\*' field (with a question mark icon) and the 'Upload IdP logo\*' field (with a question mark icon). The 'Name of IdP\*' field contains the placeholder text 'Enter IdP name'. The 'Upload IdP logo\*' field contains 'No file chosen' and a 'CHOOSE FILE' button. Below these is a 'Query Attributes\*' section.

7. (Optional) From the **Authentication Method** drop-down menu, select an authentication method.





**NOTE**

Leave the field blank to accept the IdP's default method.

8. If you are using **AD FS** and want to enable encryption, click the **Encryption Disabled** toggle to enable it (the toggle turns blue when enabled), and then configure the following encryption options that apply to your environment:



**Authentication Method** 

Select your authentication method or leave blank for default 

Encryption Enabled

Signature Disabled

Internal Keys  
[Download certificate](#) for Internal Encryption keys

Custom Keys

**Query Attributes\***

Exabeam Attributes	IdP Attributes

- In the **Query Attributes** table, map the Exabeam query attributes to the corresponding IdP user attributes by entering the same descriptive values that you did in [Add Exabeam to Your SAML Identity Provider](#), as demonstrated in the following example:

**Query Attributes\***

Exabeam Attributes	IdP Attributes
Email Address	<input type="text" value="EmailAddress"/>
Username	<input type="text" value="Username"/>
First Name	<input type="text" value="FirstName"/>
Last Name	<input type="text" value="LastName"/>
Group	<input type="text" value="Group"/>

- (Optional)* If you are ready to enable the IdP, click the **IdP Disabled** toggle. When the IdP is enabled, the toggle turns blue.

**NOTE**

You can add multiple IdPs to your Exabeam product, but only one IdP can be enabled at a time.

## EDIT IDENTITY PROVIDER

Use this Identity Provider for SSO  
Only one IdP can be enabled at a time.

**IdP Disabled**

SAML Provider\*

Okta

SSO Configuration\*

Upload the XML metadata file provided by your IdP

OktaMetadata.xml  CHOOSE FILE

Configure SSO manually


11. Click **Save**. Your identity provider now appears in the Identity Providers table.

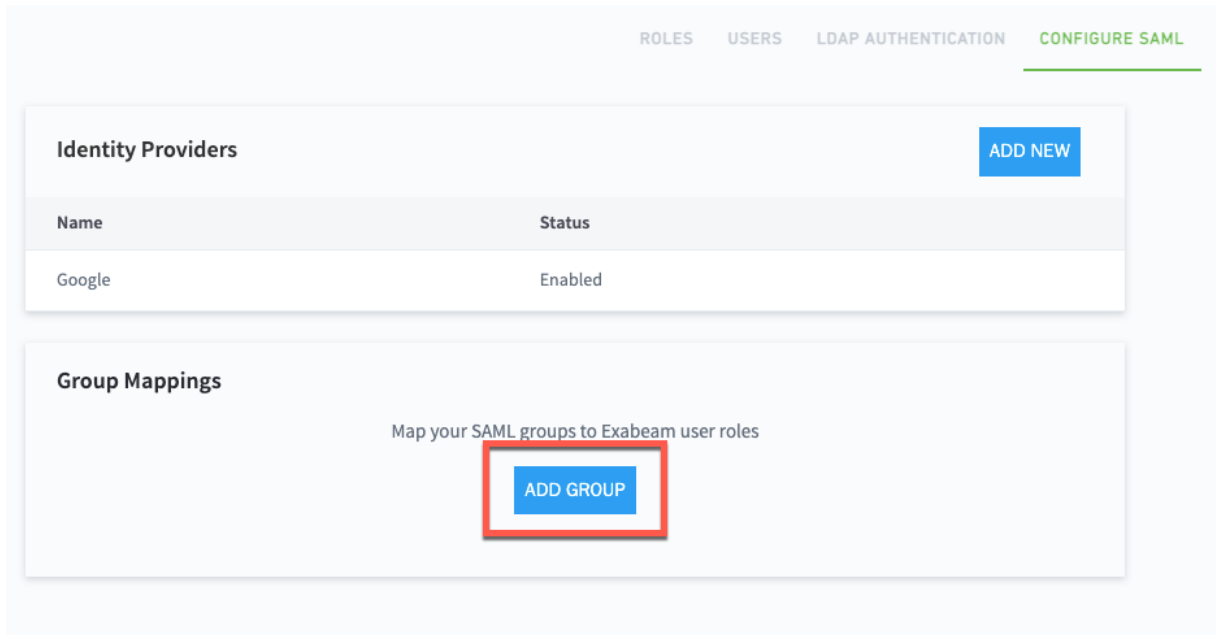
Identity Providers		ADD NEW
Name	Status	
Google	Enabled	
Okta	Disabled	
Ping	Disabled	

12. To complete the configuration, you need to map your SAML groups to Exabeam user roles. For instructions, see [Map SAML Groups to Exabeam User Roles](#).

**Map SAML Groups to Exabeam User Roles**

After adding a third-party identity provider (IdP) to your Exabeam product, you need to map the IdP user groups to the appropriate user roles in Exabeam. For example, if in your IdP you have an "Advanced Analyst" user group that needs the permissions included in the *Tier 3 Analyst (Advanced Analytics)* role, you can map the group to that role. Each group can be mapped to one or more roles as needed.

1. Navigate to **Settings**  > **Core** > **User Management** > **Configure SAML**.
2. In the Group Mappings section (which appears below the Identity Providers table), click **Add Group**.



The New Group Mapping dialog box appears.

3. From the **Identity Provider** drop-down menu, select the IdP that you want to map.

## NEW GROUP MAPPING

**Identity Provider\***

Select your IdP ▼

**Group Name/ID\* ?**

Enter SAML group name

**Exabeam User Roles\* ?**

Default Roles	
<input type="checkbox"/> Administrator	▼
<input type="checkbox"/> Tier 3 Analyst	▼
<input type="checkbox"/> Tier 1 Analyst	▼
<input type="checkbox"/> Auditor	▼
<input type="checkbox"/> Data Privacy Officer	▼

**Custom Roles**

CANCEL SAVE

- In the **Group Name/ID** field, enter the group name or ID as it is listed in the IdP.


**! IMPORTANT**

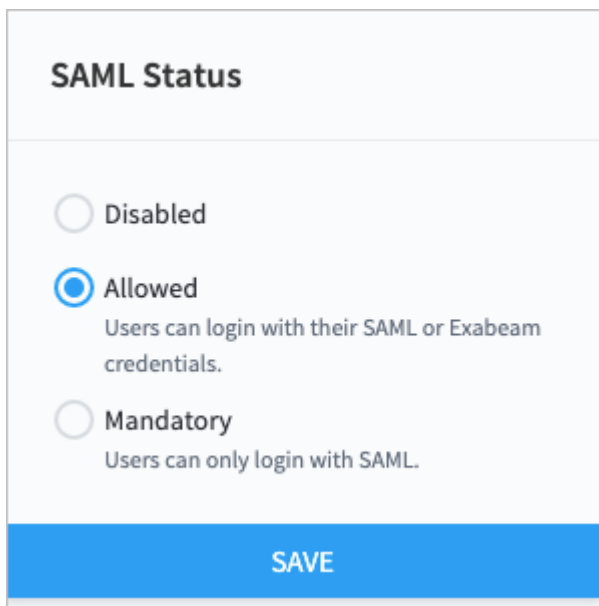
Group names are case sensitive.

- In the **Exabeam User Roles** list, select the checkboxes for the role(s) that you want to assign to the group.
- Click **Save**.

### Manage SAML Login Status

You can make authentication through your selected identity provider (IdP) mandatory for users to log in, or you can allow users to log in through either the IdP or local authentication. You can also disable your selected IdP so that users can only log in through local authentication.

1. Navigate to **Settings**  > > **User Management** > **Configure SAML**.
2. In the **SAML Status** box, select a login status for your IdP.



**SAML Status**

Disabled

**Allowed**  
Users can login with their SAML or Exabeam credentials.

Mandatory  
Users can only login with SAML.

**SAVE**

3. Click **Save**.

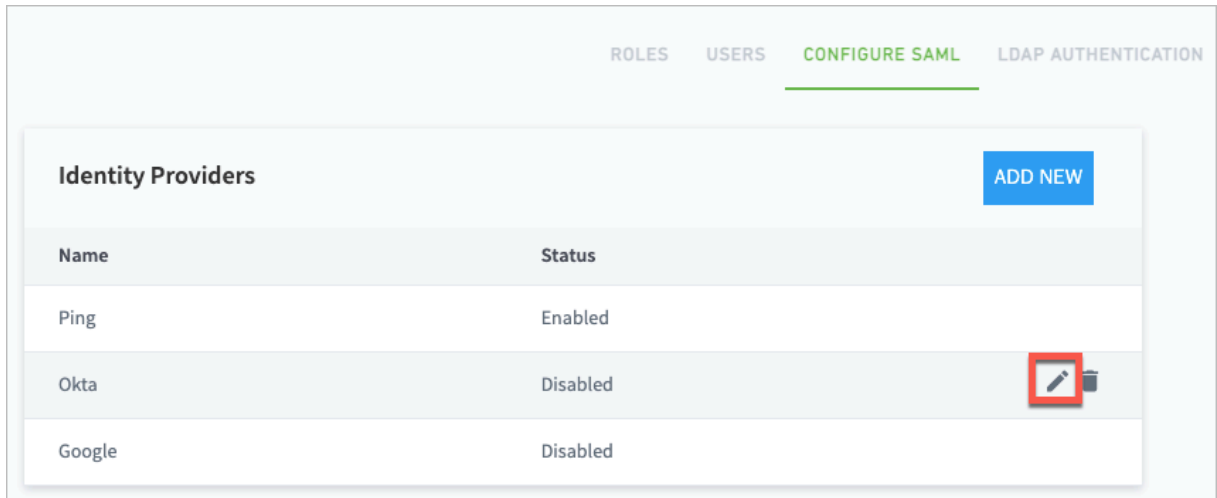
### Enable or Disable Identity Providers



#### NOTE

You can add multiple identity providers (IdPs) to your Exabeam product, but only one IdP can be enabled at a time.

1. Navigate to **Settings** > > **User Management** > **Configure SAML**.
2. Move your pointer over the IdP that you want to enable or disable, and click the edit icon.



The Edit Identity Provider dialog box opens.

3. Click the **IdP Enabled/Disabled** toggle to enable or disable the IdP as needed. The toggle is blue when the IdP is enabled and gray when it is disabled.

## EDIT IDENTITY PROVIDER

Use this Identity Provider for SSO  
Only one IdP can be enabled at a time.

IdP Enabled

SAML Provider\*

Okta

SSO Configuration\*

Upload the XML metadata file provided by your IdP

OktaMetadata.xml ✓

Configure SSO manually

Authentication Method ?

Select your authentication method or leave blank for default

Query Attributes\*

Exabeam Attributes      IdP Attributes

## Set Up LDAP Server

If you are adding an LDAP server for the first time, then the **ADD CONTEXT SOURCE** page displays when you reach the **CONTEXT MANAGEMENT** settings page. Otherwise, a list of LDAP Server appears, click **Add Context Source** to add more.

Select a **Source Type**:

- 
- Microsoft Active Directory
- NetIQ eDirectory
- Microsoft Azure Active Directory

The add/edit **CONTEXT MANAGEMENT** page displays the fields necessary to query and pull context information from your LDAP server(s), depending on the source chosen.

- •
- 

For Microsoft Active Directory:

- **Primary IP Address or Hostname** – Enter the LDAP IP address or hostname for the primary server of the given server type.



**NOTE**

For context retrieval in Microsoft Active Directory environments, we recommend pointing to a Global Catalog server. To list Global Catalog servers, enter the following command in a Windows command prompt window: `nslookup -querytype=srv gc.tcp.acme.local`. Replace `acme.local` with your company's domain name.

- **Secondary IP Address or Hostname** – If the primary LDAP server is unavailable, Exabeam falls back to the secondary LDAP server if configured.
- **TCP Port** – Enter the TCP port of the LDAP server. Optionally, select **Enable SSL (LDAPS)** and/or **Global Catalog** to auto-populate the TCP port information accordingly.
- **Bind DN** – Enter the bind domain name, or leave blank for anonymous bind.
- **Bind Password** – Enter the bind password, if applicable.
- **LDAP attributes for Account Name** – This field auto-populated with the value `sAMAccountName`. Please modify the value if your AD deployment uses a different value.

For NetIQ eDirectory:

- **Primary IP Address or Hostname** – Enter the LDAP IP address or hostname for the primary server of the given server type.
- **Secondary IP Address or Hostname** – If the primary LDAP server is unavailable, Exabeam falls back to the secondary LDAP server if configured.
- **TCP Port** – Enter the TCP port of the LDAP server. Optionally, select **Enable SSL (LDAPS)** and/or **Global Catalog** to auto-populate the TCP port information accordingly.
- **Bind DN** – Enter the bind domain name, or leave blank for anonymous bind.
- **Bind Password** – Enter the bind password, if applicable.
- **Base DN** – .
- **LDAP Attributes** – The list of all attributes to be queried by the Exabeam Directory Service (EDS) component is required. When testing the connection to the eDirectory server, EDS will collect from the server a list of the available attributes and display that list as a drop down menu. Select the name of the attribute from that list or provide a name of your own. Only names for the LDAP attributes you want EDS to poll are required (i.e., not necessarily the full list). Additionally, EDS does not support other types of attributes, therefore you cannot add “new attributes” on the list below.

For Microsoft Azure Active Directory:



- **Application Client ID** — In **App Registration** in Azure Active Directory, select the application and copy the Application ID in the **Overview** tab.
- **Application Client Secret** — In **App Registration** in Azure Active Directory, select the application and click on **Certificates & Secrets** to view or create a new client secret.
- **Tenant ID** — In **App Registration** in Azure Active Directory, select the application and copy the Tenant ID in the **Overview** tab.

Click **Validate Connection** to test the LDAP settings.



**NOTE**

If you selected **Global Catalog** for either Microsoft Active Directory or NetIQ eDirectory, this button displays as **Connect & Get Domains**.

Click **Save** to save your context source,

## Set Up LDAP Authentication

In addition to local authentication Exabeam can authenticate users via an external LDAP server.

When you arrive at this page, by default the 'Enable LDAP Authentication' is selected and the LDAP attribute name is also populated. To change the LDAP attribute, enter the new account name and click Save. To add an LDAP group, select Add LDAP Group and enter the DN of the group you would like to add. Test Settings will tell you how many analysts Exabeam found in the group. From here you can select which role(s) to assign. It is important to note that these roles are assigned to the group and not to the individual analysts; if an analyst changes groups their role will automatically change to the role(s) associated with their new group.

## Azure AD Context Enrichment



**IMPORTANT**

For the Azure AD context enrichment feature to function, your organization must have a hybrid Active Directory deployment that uses Azure AD and either Microsoft AD or Microsoft ADDS.

Organizations using Azure Active Directory (AD) can enrich their event logs by adding user context. This feature automatically pulls user attribute information from Azure AD on a daily basis and enriches logs in real time. Pulled attributes include the following:

- ID
- userType
- userPrincipalName
- mailNickname
- onPremisesSamAccountName
- displayName
- mail

For descriptions of the attributes, see [Azure Active Directory Context Tables](#).



**NOTE**

While context information from Azure AD is pulled daily, you can also perform manual pulls from Azure AD to immediately update information after changes to user accounts.

The following table lists the events that can be enriched with context from Azure AD:

Office 365	Azure	Windows Defender	Windows
Failed Sign in Alert	App Activity	EventHubs Login	Auth Events
Failed App Login	App Login	PIM Activity	App Login
App Login	Core Directory	Security Alert	Activity
Sign in Alert			
Account Unlocked			
Account Password Changed			
Account Disabled			
Security Alert 1			
Security Alert 3			
Member Added			
Member Removed			
PowerBI Activity			
Hub Network Connection			
App Activity			

### Set Up Azure AD Context Enrichment

1. Navigate to **Settings > Core > Context Management > Add Context Source**. The Context Management page opens.
2. Click **+ Add Context Source**.
3. From the **Source Type** drop-down menu, select **Microsoft Azure Active Directory**.

**SETTINGS**

**CONTEXT MANAGEMENT** ADD CONTEXT SOURCE GENERATE CONTEXT CONTEXT TABLES

Exabeam queries the primary and if necessary the secondary context server for the context information.

Server Type: Microsoft Azure Active Directory

Application Client ID\*

Application Client Secret\*

Tenant ID\*

Validate Connection

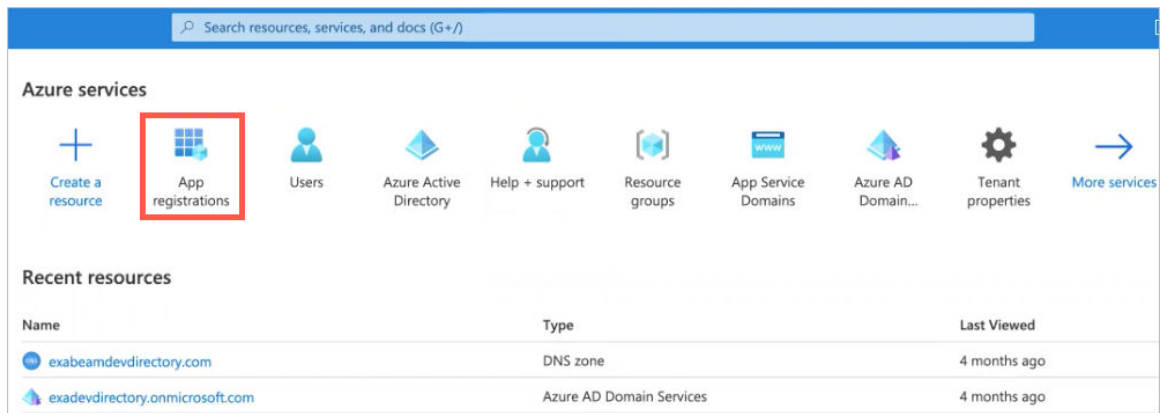
CANCEL SAVE

4. Provide the appropriate values for the following fields:

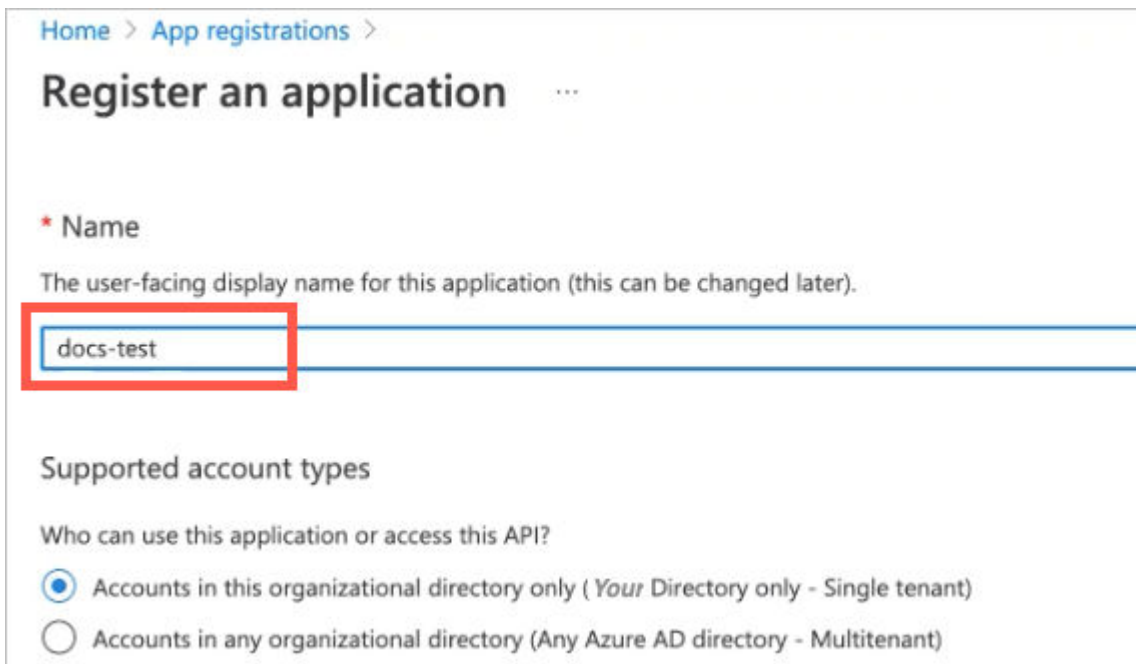
- Application Client ID
- Application Client Secret
- Tenant ID

To generate the appropriate values for these fields, do the following:

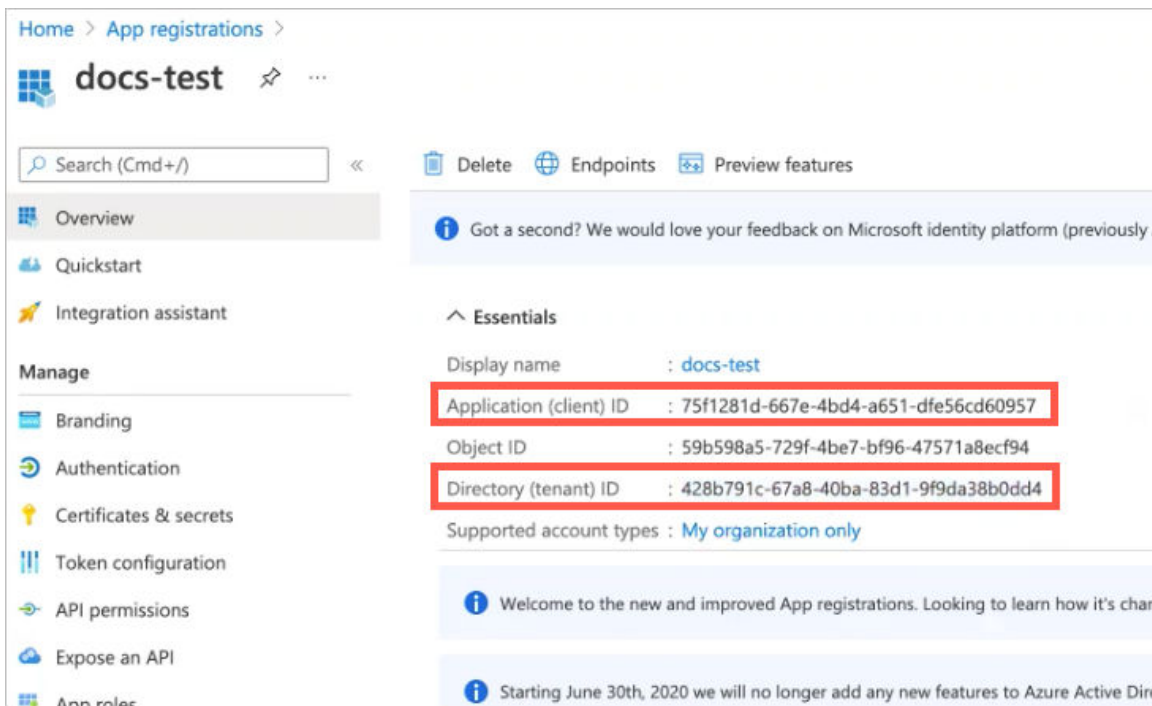
- a. Log in to Microsoft Azure.
- b. Under Azure services, click **App registrations**.



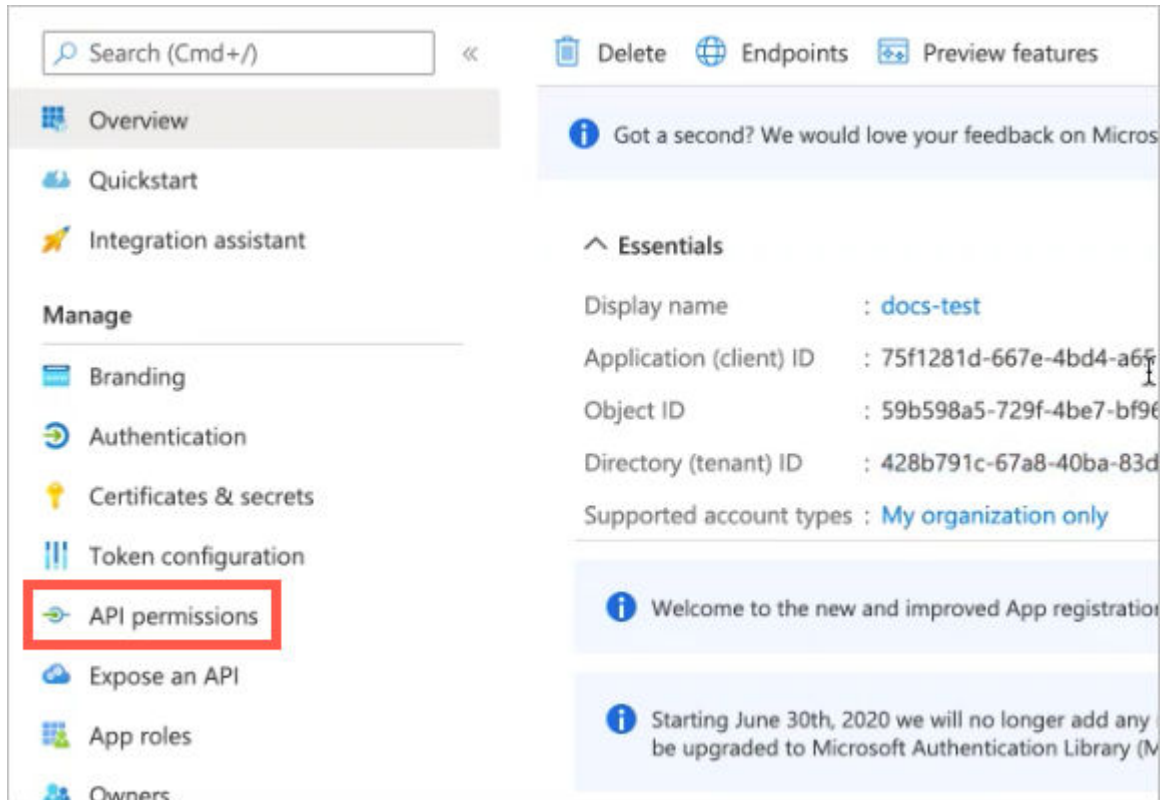
- c. Click **New registrations**.
- d. In the **Name** field, type a name for the app.



- e. Under supported account types, ensure that the following setting is selected: **Accounts in this organizational directory only (Your Directory only - Single tenant)**.
- f. At the bottom of the page, click **Register**.  
The Overview page for your new app appears.
- g. Copy the Application (client) ID and paste it into the **Application Client ID** field in Exabeam; copy the Directory (tenant) ID and paste it into the **Tenant ID** field.

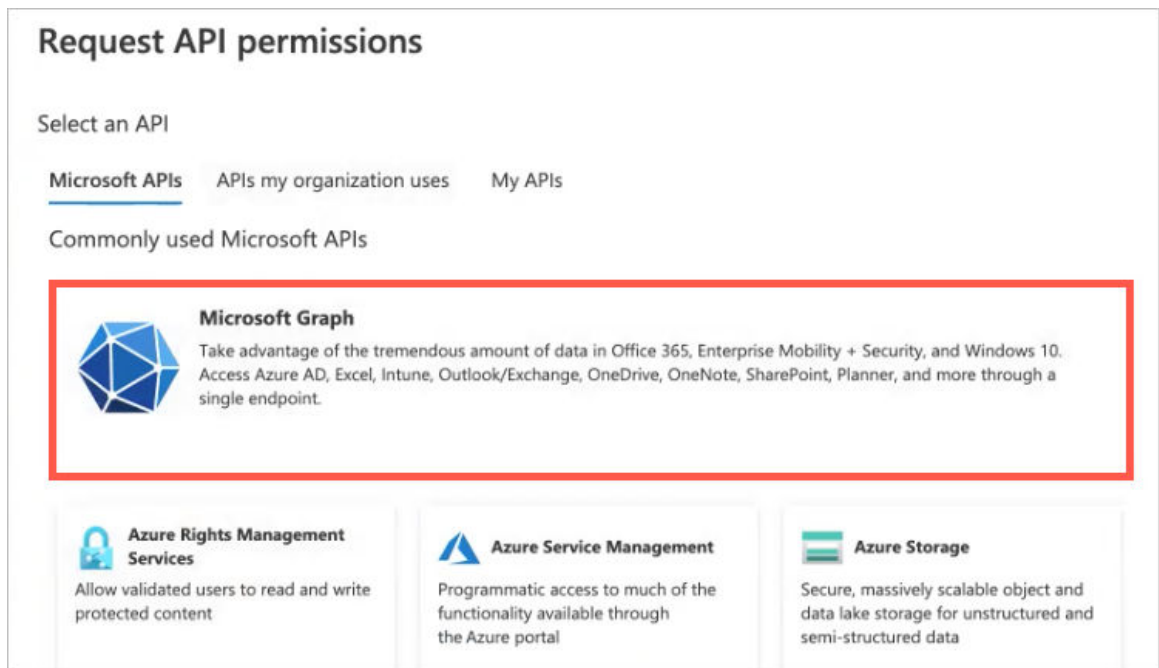


- h. In the Manage menu, click **API permissions**.

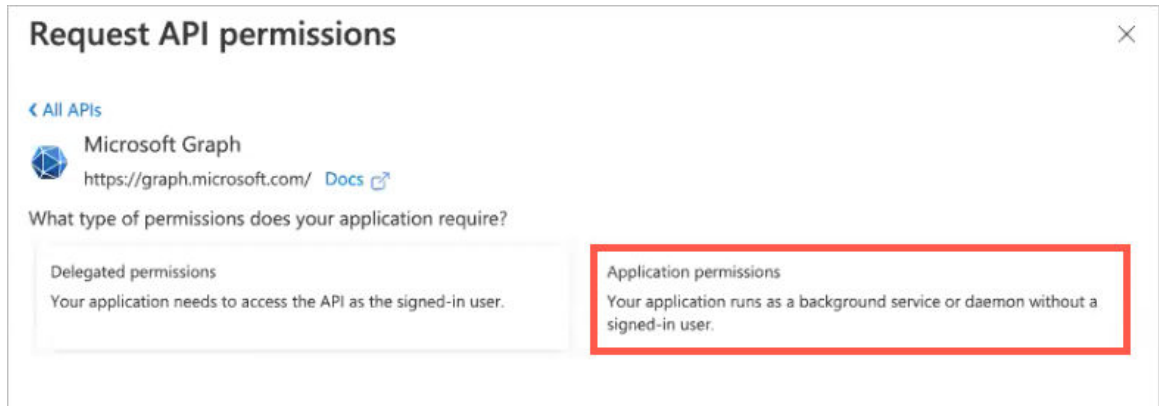


The API permissions page opens.

- i. Click **Add a permission**.  
The Request API permissions panel opens on the right.
- j. Click the **Microsoft Graph** box.



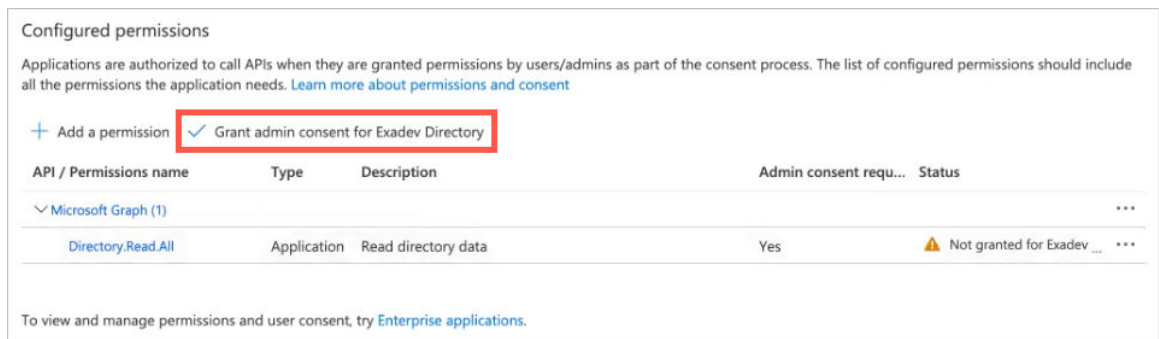
- k. Click the **Application permissions** box.



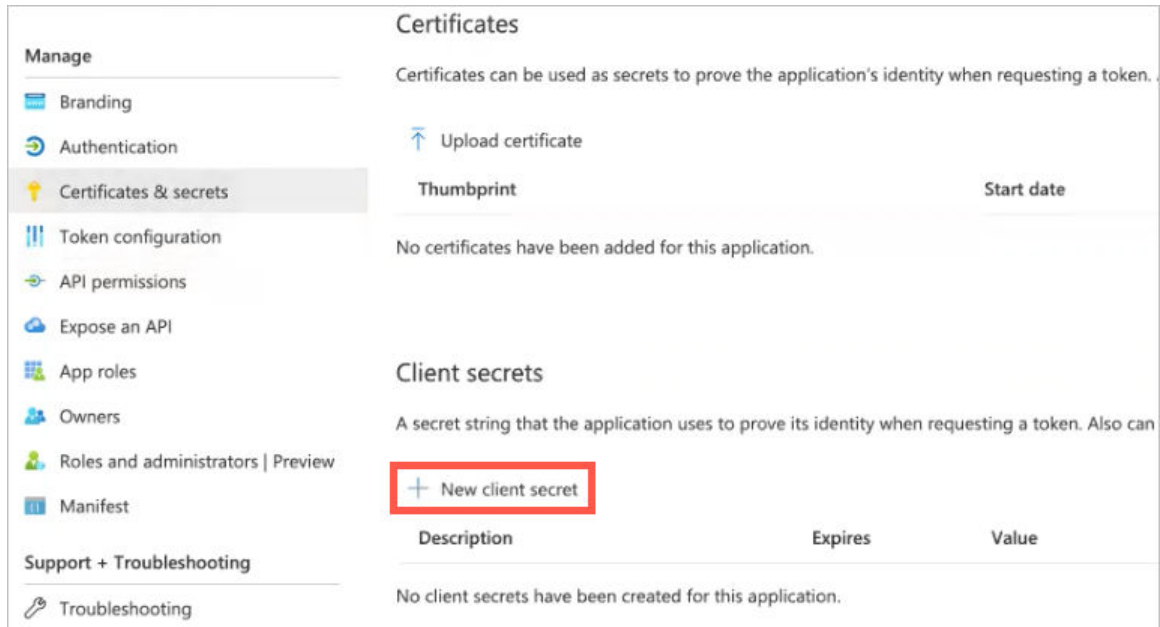
- l. In the **Select permissions** text filter, type **directory**.
- m. Click the **Directory** drop-down arrow, and then select **Directory.Read.All**.



- n. At the bottom of the panel, click **Add permissions**.  
The panel closes and the added permission appears under **Configured permissions**.
- o. Click **Grant admin consent for Exadev Directory**, and then click **Grant admin consent confirmation**.



- p. In the *Manage* menu on the left, click **Certificates & secrets**.  
The Certificates & secrets page opens.
- q. Click **New client secret**.



The Add a client secret panel opens on the right.

- r. In the **Description** field, provide a description of the secret (such as what the secret is being used for).
- s. From the **Expires** drop-down menu, select a time frame for when you want the secret to expire.
- t. At the bottom of the panel, click **Add**.  
The panel closes and the added secret appears in the **Client secrets** list.
- u. Click the copy-to-clipboard icon for the secret **Value**, and then paste the value into the **Application Client Secret** field in Exabeam.



5. To test the connection with Azure AD, click **Validate Connection**.  
A message displays to indicate whether the connection is successful.
6. If the connection is successful, click **Save** to complete the setup.  
Azure AD is added to the list of data sources on the Context Management page.

**SETTINGS**

**CONTEXT MANAGEMENT** [ADD CONTEXT SOURCE](#) [GENERATE CONTEXT](#) [CONTEXT TABLES](#)

Exabeam enriches log events with contextual information regarding your organization from data sources such as CMDB, HRMS, ITAM and IT/OT/IoT applications. Exabeam populates users, assets, accounts, peer groups, executive users, device type, manufacturer, OS version and more. Exabeam updates the context daily with any changes in the environment.

[Generate Context for all](#)

Server Type	Primary	Secondary	Updated On	
Azure	exadevdirectory.onmicrosoft.com	-	Thu May 27 2021 21:08:44 -0600	<a href="#">Refresh</a> <a href="#">Edit</a> <a href="#">Delete</a> <a href="#">Dropdown</a>

[+ Add Context Source](#)

## Set Up Context Management

Logs tell Exabeam what the users and entities are doing while context tells us who the users and entities are. These are data sources that typically come from identity services such as Active Directory. They enrich the logs to help with the anomaly detection process or are used directly by the risk engine layer for fact-based rules. Regardless of where these external feeds are used, they all go through the anomaly detection layer as part of an event. Examples of context information potentially used by the anomaly detection layer are the location for a given IP address, ISP name for an IP address, and department for a user.

Administrators are able to view and edit Exabeam's out-of-the-box context tables as well as create their own custom tables. They can select a specific table, such as Executive Users, Service Accounts, etc. and see the details of the table and all of the objects within the table. Edits can be performed on objects individually or through CSV uploads.

### Out-of-the-Box Context Tables

Context Table	Source	Available Actions
email_user	LDAP	This table is automatically populated when administrators integrate their LDAP system with Exabeam.  Administrators cannot add, edit, or delete the entries in this context table.
fullname_user	LDAP	This table is automatically populated when administrators integrate their LDAP system with Exabeam.  Administrators cannot add, edit, or delete the entries in this context table.
user_account	LDAP	This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.  Administrators can add entries manually via CSV or AD filters. Where Administrators have manually added users, they can also edit or delete entries.



Context Table	Source	Available Actions
user_department	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_division	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_manager	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Advanced Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_department_number	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_country	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_location	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_title	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.</p> <p>Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.</p>
user_fullname	LDAP	<p>This table is automatically populated when administrators integrate their LDAP system with Exabeam.</p> <p>Administrators cannot add, edit, or delete the entries in this context table.</p>

Context Table	Source	Available Actions
user_phone_cell	LDAP	This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.  Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.
user_phone_office	LDAP	This table is automatically populated when administrators integrate their LDAP system with Exabeam and add regular expression through the Advanced Analytics tab.  Administrators can add entries manually via CSV or Active Directory filters. Where Administrators have manually added users, they can also edit or delete entries.
user_is_privileged	Administrators	Administrators can add entries manually, via CSV, or Active Directory. Entries can also be edited or deleted.

## Azure Active Directory Context Tables

Context Table	Description
ID	User's globally unique identifier (GUID).
userType	Indicates whether the user is a member or guest.
userPrincipalName	User principal name (UPN) of the user.
mailNickname	Mail alias for the user.
onPremisesSamAccountName	User's samAccountName in the on-prem LDAP, which is synced to Azure AD.
displayName	Display name for the user.
mail	User's email address from the Azure user profile.

## Threat Intelligence Service Context Tables

The table below shows the description of each available threat intelligence feed to a context table in Advanced Analytics:

Context Table	Description
is_ip_threat	IP addresses identified as a threat.
is_ransomware_ip	IP addresses associated with ransomware traffic.
is_tor_ip	Known Tor IP addresses.
reputation_domains	Domains associated with malware traffic
web_phishing	Domains associated with phishing attacks.

For more information on Exabeam threat intelligence service, please see the section *Threat Intelligence Service Overview*.

## Custom Context Tables

Exabeam provides several filters and lookups to get your security deployment running immediately. However, there may be assets and users within your organization that need particular attention and cannot be fully addressed out of the box. Custom context tables allow you the flexibility to create watchlists or reference lists for assets, threat intelligence indicators, and users/groups that do not fit in the typical deployment categories. Custom context tables let you put

parts of your organization under extra monitoring or special scrutiny, such as financial servers, privileged insiders, and high-level departed employees.

Within Advanced Analytics, you can create watchlists using context tables. When creating the table, the Label attribute allows you to attach tags to records that match entries in your context table. This provides quick access to query your results and/or focus your tracking using a global characteristic.

You can also build rules based on entries in your context tables. Set up alerts, actions, or playbooks to trigger when conditions match records, such as access to devices in a special asset group.

### Prepare Context Data

You can upload data as CSV files with either key and value columns or key-only column. All context tables include a Label to tag matching records into groups during parsing and filtering.

**Key-value CSV** – Two-field data file with a header row. This lookup lists correlations between the two fields, such as:

Key Fieldname	Value Fieldname
AC1Group	Accounts Receivable
AC2Group	Accounts Payable

**Key-only CSV** – Single-field data file with no header row. Items on this list are compared to as being present or not during data filtering. For example, a watchlist context table, SpecialGroup, consists of user groups of special interest:

“Accounts Receivable”

“Accounts Payable”

“Accounting Database Admin”

You can create a correlation rule that sends an alert when the monitoring data contains a user having the group name that matches any in the Special Group table.

**Label** – The named tag associated with a record. This allows you to filter groups of records during parsing or filtering. You can also use labels to assemble watchlists based on groupings rather than by individual asset or user record.

#### NOTE

You can opt not to use labels by selecting **No Label** during table creation. Otherwise, labels are associated with tables and its records. For key-value context tables, the Label is drawn from the value field of the matching context table entry. For key-only context tables, the **Label** is the table attribute you enter in the **Manual Assignment** field during table creation and is used to tag all matching records.

## Set Up Authentication and Access Control

**NEW CONTEXT TABLE**

Exabeam allows you to bring in custom context tables via CSV uploads or by adding objects one at a time. As an example, you can bring in your list of privileged users, or critical assets into Exabeam as context tables.

Title\*

Object Type

Users  Assets  Miscellaneous

Type

Key Only

Label Assignment

No Label  Manual Assignment

CANCEL SAVE

### Create Custom Lookups

You must first create a table object to add contextual data to. Create the table with key-only or key-value field and whether labels will be used based on the needs of your organization. Use the various methods to add content into your table depending on your data source.

### Create a Context Table

To introduce context data into your environment, create a table object to contain your data and reference it in queries and lookups.

1. Navigate to **Settings > Accounts & Groups > Context Tables**.
2. At the top right of the UI, click the blue + to open the **New Context Table** dialogue box.



3. Fill in the details of the type of context table that this will be.

**NEW CONTEXT TABLE**

Exabeam allows you to bring in custom context tables via CSV uploads or by adding objects one at a time. As an example, you can bring in your list of privileged users, or critical assets into Exabeam as context tables.

Name\*

Object Type

Users  Assets  Miscellaneous

Type

Key Value

Label Assignment

No Label  Automatic Assignment from value

CANCEL SAVE

Fill in table attribute fields:

**Name** – A unique name identifying the table in queries and in the context of your organization.

**Object Type** – The type gives the table additional tagging (with information on the potential data source, such as LDAP for users or user groups).

- **Users** – This object type is associated with users and user group context tables. LDAP data sources can be used to fill its content.
- **Assets** – These are itemizable objects of value to your organization. These can be devices, files, or workstations/servers.
- **Miscellaneous** – These are reference objects of interest, such as tags for groups of objects within a department or network zones.

**Type** – Select the field structure in the table as *Key Value* or *Key Only*. See *Prepare Context Data* for more information. If you are creating a correlation context table, use *Key Only*.

**Label Assignment** – Click the text source for creating the label or use no label. See *Prepare Context Data* for more information.

4. Click **Save** to advance to the table details UI for the newly created context table.

Your table is ready to store data. The following sections describe ways to add data to your table. Each method is dependent on the data source and intended use of the table.

### Import Data into a Context Table Using CSV

This is the most flexible method to create custom context tables as the CSV file can contain any category or type of data that you wish to monitor.

1. From the **Settings** menu, navigate to **Analytics > Accounts & Groups > Context Tables**, and then click the **Context Tables** tab.

The list of existing context tables appears.

The screenshot shows the 'CONTEXT MANAGEMENT' section of the Exabeam settings. It includes a table with the following data:

Name	Object Type	Table Type	Label Value	Last Config Change	Connections
user_azure_display_name	Users	Default		09/17/2021	Manual, Azure AD
user_title	Users	Default		09/17/2021	Manual, All LDAP Servers
user_is_executive	Users	Default	executive	09/17/2021	Manual
user_azure_mail	Users	Default		09/17/2021	Manual, Azure AD
is_competition	Users	Default		09/17/2021	Manual
user_phone_office	Users	Default		09/17/2021	Manual, All LDAP Servers

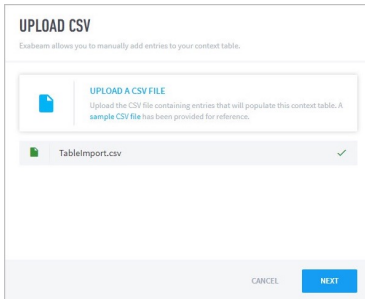
2. Click the name of the table into which you want to import the CSV file. The set up page for the context table appears.

3. Click the **Upload Table** icon.

## Set Up Authentication and Access Control



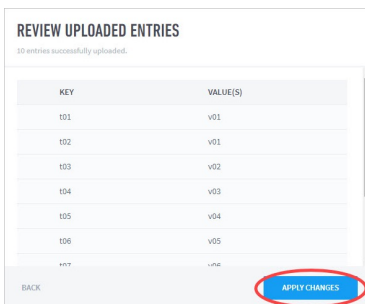
4. Click **Upload CSV**. From your file system, select the CSV file you wish to import, then select **Next**.



### NOTE

Key and value (2 fields) tables require a header first row. Do not include a header for keys-only CSV files (1 field). Table names may be alpha-numeric with no blank spaces. (Underscore is acceptable.)

5. Inspect the contents that will be added to your table. Select **Apply Changes**, when you are done.



Once context has been integrated, it is displayed in the table. You can use the lookup tables in rules as required.

**SETTINGS**

**CONTEXT MANAGEMENT** ADD CONTEXT SOURCE GENERATE CONTEXT CONTEXT TABLES

**CONTEXT TABLES** > **USERACCOUNTCONTROL\_USER** 🔍 📄 +

**DETAILS**

Name: useraccountcontrol\_user  
Object Type: Users  
Type: Key Value  
Label: Assignment From Value

[Edit](#)

**CONNECTIONS**

Connect this table to an LDAP Server to populate it with a filtered set of entries from that directory.

[+ Add Connection](#)

	Source	Key	Value(s)	Created Time
<input type="checkbox"/>	Manual	test_dl	Director	2021-09-17 15:50
<input type="checkbox"/>	Manual	ceo	ceo	2021-09-17 15:50
<input type="checkbox"/>	Manual	manager_a	Manager A	2021-09-17 15:50
<input type="checkbox"/>	Manual	manager_b	Manager	2021-09-17 15:50
<input type="checkbox"/>	Manual	manager_c	Manager	2021-09-17 15:50
<input type="checkbox"/>	Manual	swatid	titled1	2021-09-17 15:50
<input type="checkbox"/>	Manual	revathiwarale	testqa	2021-09-17 15:50
<input type="checkbox"/>	Manual	Test5	user_new	2021-09-17 15:50

**EXPORT AS CSV**

**NOTE**

The Created Time column displays the time that the CSV file was uploaded. Context values may change over time. For example, a user's role may change within an organization, in which case the value for the user's job title would depend on when the context was uploaded. The Created Time field helps to explain such changes in values.

For assistance in creating custom context tables, contact Exabeam Customer Success by opening a case at [Exabeam Community](#).

### Import Data into a Context Table Using an LDAP Connection

This section details the steps required to create context tables to customize your lookups. In this example, we are creating a lookup table with two fields: the `userAccountControl` field and the `User ID` field. This allows the event enricher to map one to the other. For example, let's say you have a log that does not include the username, but instead included the `userAccountControl` field. This would map the two together. A similar use case would be badge logs: you could create a lookup table that maps the badge ID to the actual username, assuming the badge ID is contained in LDAP.

1. Navigate to the **Settings > Analytics > Accounts & Groups > Context Tables**.
2. Click the **+** icon to add a new table.

## Set Up Authentication and Access Control

**SETTINGS**

**CONTEXT MANAGEMENT** ADD CONTEXT SOURCE GENERATE CONTEXT CONTEXT TABLES

Exabeam allows you to bring in custom context tables via CSV uploads or by adding objects one at a time. As an example, you can bring in your list of privileged users, or critical assets into Exabeam as context tables.

Object Type: Users 🔍 +

Name	Object Type	Table Type	Label Value	Last Config Change	Connections
<a href="#">user_azure_display_name</a>	Users	Default		09/17/2021	Manual, Azure AD
<a href="#">user_title</a>	Users	Default		09/17/2021	Manual, All LDAP Servers
<a href="#">user_is_executive</a>	Users	Default	executive	09/17/2021	Manual
<a href="#">user_azure_mail</a>	Users	Default		09/17/2021	Manual, Azure AD
<a href="#">is_competition</a>	Users	Default		09/17/2021	Manual
<a href="#">user_phone_office</a>	Users	Default		09/17/2021	Manual, All LDAP Servers

- Complete the New Context Table dialog box as needed for your context table.

### Example 1.

**NEW CONTEXT TABLE**

Exabeam allows you to bring in custom context tables via CSV uploads or by adding objects one at a time. As an example, you can bring in your list of privileged users, or critical assets into Exabeam as context tables.

**Name\***

**Object Type**

Users  Assets  Miscellaneous ?

**Type** ?

Key Value ▼

**Label Assignment** ?

No Label

Automatic Assignment from value

CANCEL SAVE



### NOTE

If you do not want to add a label to matching records during parsing or filtering, click **No Label**.

- Click **Save**.  
The set up page for the new context table appears.
- Click **+ Add Connection** to connect the context table to an LDAP domain server.



## Set Up Authentication and Access Control

**SETTINGS**

**CONTEXT MANAGEMENT** ADD CONTEXT SOURCE GENERATE CONTEXT CONT

**CONTEXT TABLES > USERACCOUNTCONTROL\_USER** 🔍 📄

**DETAILS**

Name: useraccountcontrol\_user  
Object Type: Users  
Type: Key Value  
Label: Assignment From Value

[Edit](#)

**CONNECTIONS**

Connect this table to an LDAP Server to populate it with a filtered set of entries from that directory.

[+ Add Connection](#)

Source	Key	Value(s)	Created T
No entries added to this context table. Please upload CSV or add entries manually.			

6. Select the **LDAP Server(s)**, **Key**, and **Value** to populate the context table. Optionally, filter the attribute source with conditions by clicking **ADD CONDITION**.

**NEW CONNECTION**

LDAP Server(s)

Attribute Source  
Select attributes to populate your table. To filter attributes add conditions below.

Key: Access Status      Value: User ID

Conditions

No conditions added to this connection. Add conditions to filter attributes.

[ADD CONDITION](#)

[CANCEL](#) [TEST CONNECTION](#)

7. Click **TEST CONNECTION** to view and validate the test results, and then click **SAVE**.

## Set Up Authentication and Access Control

**TEST CONNECTION RESULTS**  
78 results found. All of the results are listed.

Key	Value(s)
Account is active	Guest
Account is active	lisa492
Account is active	judy
Account is active	Administrator
Account is active	alice
Account is active	lisa514_po
Account is active	rianna123
Account is active	rianna
Account is active	riannatest
Account is active	Administrator
Account is active	Guest

BACK SAVE

Once context has been integrated, it is displayed in the table. You can use the lookup table in rules as required.

**SETTINGS**

**CONTEXT MANAGEMENT** ADD CONTEXT SOURCE GENERATE CONTEXT **CONTEXT TABLES**

**CONTEXT TABLES > USERACCOUNTCONTROL\_USER** SEARCH REFRESH + EXPORT AS CSV

**DETAILS**

**Name:** useraccountcontrol\_user  
**Object Type:** Users  
**Type:** Key Value  
**Label:** Assignment From Value

**CONNECTIONS**

Connect this table to an LDAP Server to populate it with a filtered set of entries from that directory.

**LDAP Server(s)**  
All LDAP Servers

Source	Key	Value(s)	Created Time
...	ir-auto-user2	ir-auto-user2	2021-08-29 03:44
...	lin	lin	2021-08-27 14:00
...	Aegon Targayan	aegon	2021-08-27 14:00
...	plt auto	plt_auto_ad	2021-08-27 14:00
...	xudana	xudana11	2021-08-27 14:00
...	David	david	2021-08-27 14:00

### NOTE

The Created Time column displays the time that the context was processed, as context values may change over time. For example, a user's role may change within an organization, in which case the value for the user's job title would depend on when the context was processed. The Created Time field helps to explain such changes in values.

For assistance in creating custom context tables, contact Exabeam Customer Success by opening a case at [Exabeam Community](#)

## Mask Data Within the Advanced Analytics UI



### NOTE

To enable or disable and configure data masking, contact your Exabeam technical representative.



### NOTE

Data masking is not supported in Case Management or Incident Responder modules.

Data masking within the UI ensures that personal data cannot be read, copied, modified, or removed without authorization during processing or use. With data masking enabled, the only user able to see a user's personal information will be users assigned to the permission "View Clear Text Data". The default role "Data Privacy Officer" is assigned this permission out of the box. Data masking is a configurable setting and is turned off by default.

- To enable data masking in the UI, open `/opt/exabeam/config/tequila/custom/application.conf`, and set `dataMaskingEnabled` to `true`.  
If your `application.conf` is empty, copy the following text and paste it into the file:

```
tequila {
  PII {
    # Globally enable/disable data masking on all the PII configured fields.
    Default value is false.
    dataMaskingEnabled = true
  }
}
```

You're able to fully customize which PII data is masked or shown in your deployment. The following fields are available when configuring PII data masking:

- **Default:** This is the standard list of PII values controlled by Exabeam. If data masking is enabled, all of these fields are encrypted.
- **Custom:** Encrypt additional fields beyond the default list by adding them to this custom list. The default is empty.
- **Excluded:** Do not encrypt these fields. Adds that are in the default list to expose their values in your deployment. The default is empty.

For example, if you want to mask all default fields other than "task name" and also want to mask the "address" field, then you would configure the lists as shown below:

```
PII {
  # Globally enable/disable data masking on all the PII configured fields.
  Default value is false.
  dataMaskingEnabled = true
  dataMaskingSuffix = ":M"
  encryptedFields = {
    #encrypt fields
    event {
      default = [
```

```

        #EventFieldName
        "user",
        "account",
        ...
        "task_name"
    ]
    custom=[ "address" ]
    excluded=[ "task_name" ]
}
...
}
}

```

### Mask Data for Notifications

You can configure Advanced Analytics to mask specific fields when sending notable sessions and/or anomalous rules via email, Splunk, and QRadar. This prevents exposure of sensitive data when viewing alerts sent to external destinations.



#### NOTE

Advanced Analytics activity log data is not masked or obfuscated when sent via Syslog. It is your responsibility to upload the data to a dedicated index which is available only to users with appropriate privileges.

Before proceeding through the steps below, ensure your deployment has:

- Enabled data masking (instructions below)
- Configured a destination for Notable Sessions notifications sent from Advanced Analytics via [Notifications](#)

By default, all fields in a notification are unmasked. To enable data masking for notifications, the `Enabled` field needs to be set to `true`. This is located in the `application.conf` file in the path `/opt/exabeam/config/tequila/custom`.

```

NotificationRouter {
    ...
    Masking {
        Enabled = true
        Types = []
        NotableSessionFields = []
        AnomaliesRulesFields = []
    }
}

```

Use the `Types` field to add the notification destinations (Syslog, Email, QRadar, and/or Splunk). Then, use the `NotableSessionFields` and `AnomaliesRulesFields` to mask specific fields included in a notification.

For example, if you want to mask the user, source host and IP, and destination host and IP for notifications sent via syslog and Splunk, then you would configure the lists as shown below:

## Set Up Authentication and Access Control

```
NotificationRouter {  
    ...  
    Masking {  
        Enabled = true  
        Types = [Syslog, Splunk]  
        NotableSessionFields = ["user", "src_host", "src_ip", "dest_host",  
"dest_ip"]  
    }  
}
```

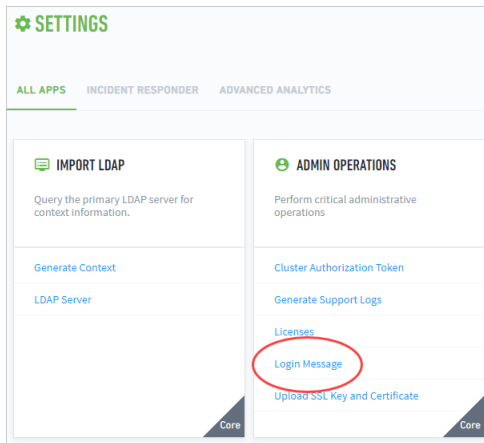
## Additional Configurations

### Display a Custom Login Message


You can create and display a custom login message for your users. The message is displayed to all users before they can proceed to login.

To display a custom login message:

1. On a web browser, log in to your Exabeam web console using an account with administrator privileges.
2. Navigate to **Settings > Core > Admin Operations > Login Message**.



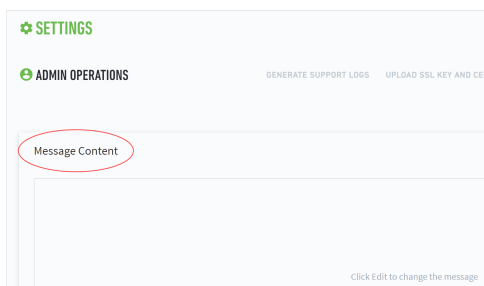
3. Click **EDIT**.



4. Enter a login message in **Message Content**.

**NOTE**

The message content has no character limit and must follow UTF-8 format. It supports empty lines between text. However, it does not support special print types, links, or images.



A common type of message is a warning message. The following example is a sample message:

## Usage Warning

This computer system is for authorized use only. Users have no explicit or implicit expectation of privacy.

Any or all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to an authorized site. By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of the authorized site.

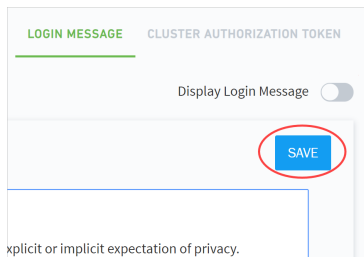
Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties. By continuing to use this system you indicate your awareness of and consent to these terms and conditions of use. LOG OFF IMMEDIATELY if you do not agree to the conditions stated in this warning.



### NOTE

This sample warning message is intended to be used only as an example. Do not use this message in your deployment.

#### 5. Click **SAVE**.

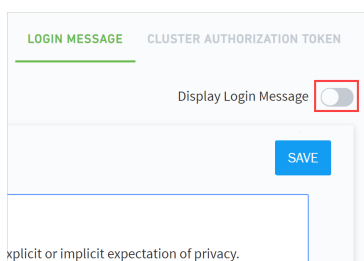


#### 6. Click the **Display Login Message** toggle to enable the message.

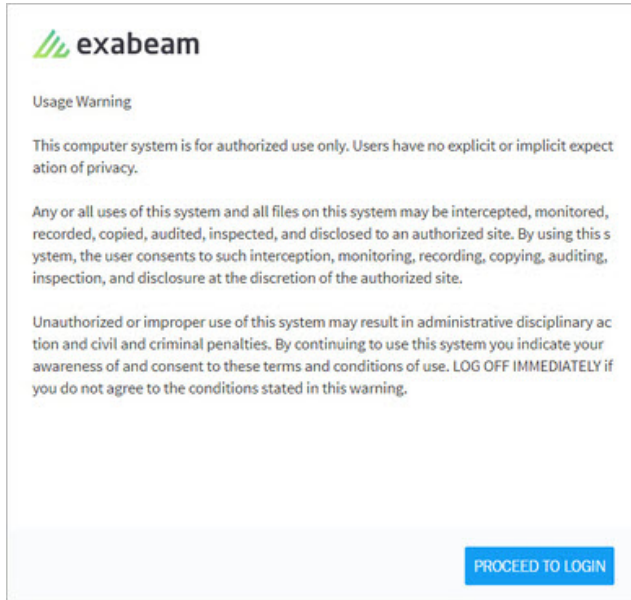


### NOTE

You can hide your message at any time without deleting it by disabling the message content.




Your custom login message is now shared with all users before they proceed to the login screen.



## Reprocess Jobs

Two types of reprocessing are available in Advanced Analytics. You can reparse raw logs in order to generate new events, or you can run the analytics engine to reprocess the log feeds,

To access any of the Exabeam reprocessing options, navigate to the **Exabeam Engine** page in Advanced Analytics:

1. From the left sidebar, click **SETTINGS** , then select **Analytics**.
2. Under **Admin Operations**, select **Exabeam Engine**.
3. Procedures for reprocessing vary depending on which type of reprocessing you want to start. See the appropriate sections below for more information.

### Reparsing Raw Logs to Create New Events

Use this option when you want to re-ingest raw log data to parse new events. The procedure for re-parsing log data varies depending on which version of Advanced Analytics you are using. For more information, see the appropriate section below.

#### ***63 and Later***

In this Advanced Analytics version, you can use the **UIP Log Reprocessing** option located in the **Exabeam Unified Log Ingestion Engine** panel. When you click this option, you are redirected to the cloud-based [Log Stream](#) functionality. Log Stream provides visibility into the unified ingestion pipeline with the following tabs:

- **Re-parsing Jobs** – You can opt to re-parse logs by scheduling a re-parsing job.
- **Live Tail** – View samples of incoming data in real time to ensure proper processing.



### *i60 to i62*

In this Advanced Analytics version, you can use the **Ingest Log Feeds** option located in the **Exabeam Log Ingestion Engine** panel. When you click this option, the LIME engine is restarted so that it ingests logs from log feeds that are defined in the Advanced Analytics **Log Feeds** settings.

To restart the LIME engine:

1. Click **Ingest Log Feeds**, and select specific log feeds for restart.
2. Select a restart option from the following settings:
  - **Restart the engine** – The engine continues processing from where it left off.
  - **Restart from the initial training period** – The engine continues processing from the initial training period.
  - **Restart from a date** – The engine continues processing from a specified date.
3. Click **Ingest feeds** to start the engine.

### Run the Analytics Engine to Reprocess Log Feeds

Reprocess analytics engine jobs when you have made changes that you want to see reflected in events and timelines.

1. In the **Exabeam Analytics Engine** panel, click **Restart Processing** and select a restart option from the following settings:
  - **Restart the engine** – The engine continues processing from where it left off.
  - **Restart from the initial training period** – The engine continues processing from the initial training period.
  - **Restart from a date** – The engine chooses the nearest snapshot available for the specified date and reprocesses from this date.
2. Click **Process** to start the engine. The system validates any changes and checks for errors. If errors are identified, they are listed and the engine does not start processing. If no errors are identified, the engine starts.

To view the status of these reprocessed analytics engine jobs, click the **Reprocessing Jobs** tab in the status table at the bottom of the **Exabeam Engine** page. As shown in the image below, this tab shows the status of each reprocessed job, such as completed, in-progress, pending, and canceled.

Asset Groups	Current Status	History	Reprocessing Jobs				REFRESH STATUS
	Status	Creator	Created	Started	Ended	Duration	
Accounts & Groups	> In Progress: 81% complete	M.	Jul 9th 2018, 14:32:17	Jul 9th 2018, 14:32:57		47m 18s	
Peer Groups	> Pending	M.	Jul 9th 2018, 15:49:28				
Executives	> Pending	M.	Jul 9th 2018, 16:17:16				
Service Accounts	> Pending	M.	Jul 9th 2018, 16:48:49				
Exabeam User Management	> Canceled	M.	Jul 7th 2018, 13:49:16	Jul 7th 2018, 13:50:27	Jul 7th 2018, 18:20:57	4h 30m 30s	
Roles	> Complete	M.	Jul 3rd 2018, 13:23:16	Jul 3rd 2018, 13:24:06	Jul 3rd 2018, 17:53:46	4h 30m 31s	
Users	> Complete	M.	Jul 3rd 2018, 16:03:16	Jul 3rd 2018, 17:53:47	Jul 3rd 2018, 23:53:47	6h 0m	
LDAP Authentication							
Admin Operations							
Exabeam Engine							

To cancel a reprocessing job for any reason, select the job in the **Reprocessing Jobs** table and then click **Cancel Job**.

### Configure Job Status Notifications

You can configure email and Syslog notifications for certain job reprocessing status changes, including start, end, and failure. For information about configuring these notifications, see [Notifications](#). You'll find the job status check boxes listed under **Notifications by Product >Advanced Analytics**.

### User Engagement Analytics Policy

Exabeam uses user engagement analytics to provide in-app walkthroughs and anonymously analyze user behavior, such as page views and clicks in the UI. This data informs user research and improves the overall user experience of the Exabeam Security Management Platform (SMP). Our user engagement analytics sends usage data from the web browser of the user to a cloud-based service called Pendo.

There are three types of data that our user engagement analytics receives from the web browser of the user. This data is sent to a cloud-based service called Pendo:

- **Metadata** – User and account information that is explicitly provided when a user logs in to the Exabeam SMP, such as:
  - User ID or user email
  - Account name
  - IP address
  - Browser name and version
- **Page Load Data** – Information on pages as users navigate to various parts of the Exabeam SMP, such as root paths of URLs and page titles.
- **UI Interactions Data** – Information on how users interact with the Exabeam SMP, such as:
  - Clicking the Search button
  - Clicking inside a text box

## Additional Configurations

- Tabbing into a text box

## Configure Rules

Create and modify rules in Advanced Analytics settings.

From the lower-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations** > **Exabeam Rules**.

From Exabeam Rules settings:

- [View](#) all rules configured in your system.
- [Edit](#) an existing rule.
- Revert to default settings and remove all changes you previously made.
- [Create](#) a fact-based rule.
- Clone and make a copy of a rule.
- [Disable](#) a rule so it doesn't trigger or enable a rule you previously disabled.
- [Reload](#) all rules. You reload rules to apply changes when you create a new rule or modify an existing rule.

### What Is an Exabeam Rule?

So what exactly is a rule anyway? There are two types of Exabeam rules:

- Model-based
- Fact-based

Model-based rules rely on a model to determine if the rule should be applied to an event in a session, while fact based rules do not.

For example, a Fireeye malware alert is fact based and does not require a model in order to be triggered. On the other hand, a rule such as an abnormal volume of data moved to USB is a Model-based rule.

Model-based rules rely on the information modeled in a histogram to determine anomalous activities. A rule is triggered if an event is concluded to be anomalous, and points are allocated towards the user session in which the event occurred. Each individual rule determines the criticality of the event and allocates the relevant number of points to the session associated with that event.

Taken together, the sum of scores from the applied rules is the score for the session. An example of a high-scoring event is the first login to a critical system by a specific user – which allocates a score of 40 to a user's session. Confidence in the model must be above a certain percentage for the information to be used by a rule. This percentage is set in each rule, though most use 80%. When there is enough reliable information for the confidence to be 80% or higher, this is called convergence. If convergence is not reached, the rule cannot be triggered for the event.

## How Exabeam Models Work

Since anomaly-based rules depend on models, it is helpful to have a basic understanding of how Exabeam's models work.

Our anomaly detection relies on statistical profiling of network entity behavior. Our statistical profiling is not only about user-level data. In fact, Exabeam profiles *all* network entities, including hosts and machines, and this extends to applications or processes, as data permits. The statistical profiling is histogram frequency based. To perform the histogram-based profiling, which requires discrete input, we incorporate a variety of methods to transform and to condition the data. Probability distributions are modeled using histograms, which are graphical representations of data. There are three different model types – categorical, numerical clustered, and numerical time-of-week.

**Categorical** is the most common. It models a string with significance: number, host name, username, etc. Where numbers fall into specific categories which cannot be quantified. When you model which host a user logs into, it is a categorical model.

**Numerical Clustered** involves numbers that have meaning – it builds clusters around a user's common activities so you can easily see when the user deviates from this norm. For example, you can model how many hosts a user normally accesses in a session.

**Numerical Time-of-Week** models when users log into their machines in a 24-hour period. It models time as a cycle so that the beginning and end of the period are close together, rather than far apart. For example, if a user logs into a machine Sunday at 11:00 pm, it is closely modeled to Monday at 12:00am.

## Model Aging

Over time, models built in your deployment naturally become outdated. For example, if an employee moves to a different department or accepts a promotion and they do not adhere to the same routines, access points, or other historical regularities.

We automatically clean up and rebuild all models on a regular basis (default is every 16 weeks) to ensure your models are as accurate and up-to-date as possible. This process also enhances system performance by cleaning out unused or underutilized models.

## Rule Naming Convention

Exabeam has an internal Rule ID naming convention that is outlined below. This system is used for Exabeam created rules and models only. When a rule is created or cloned by a customer, the system will automatically create a Rule ID for the new rule that consists of `customer-created`, followed by a random hash. For example, a new rule could be called, `customer-created-4Ef3DDYQsQ` {.

The Exabeam convention for model and rule names is: ET-SF-A/F-Z

ET: The event types that the model or rule addresses. For example,

- RA = remote-access

## Configure Rules

- NKL = NTLM/Kerberos-logon
- RL = remote-logon

SF: Scope and Feature of the model. For example,

- HU = Scope=Host, Feature=User
- OZ = Scope=Organization, Feature=Zone

A/F: For rules only

- A = Abnormal
- F = First

Z : Additional Information (Optional). For example,

- DC: Domain Controller models/rules
- CS: Critical Systems

## View Rules in Advanced Analytics

View all rules configured in your system in Advanced Analytics settings.

From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations** > **Exabeam Rules**

View all rules configured in your system based on your license. To find specific rules, filter all rules or search for the rule.

At a glance, view:

- **Name** – The rule name and rule ID, the use case and scenario the rule detects, and the rule category the rule falls under.
- **Trigger** – How often the rule triggers during user sessions.
- **Updated** – The date and time the rule was updated and by whom.
- **Risk Score** – The risk score assigned to a session when the rule triggers.
- **Status** – Whether the rule is enabled or disabled.
- **Published** – Whether the rule is a **DRAFT** and has been edited but not reloaded, or **PUBLISHED** and all changes are reloaded.

## Filter Rules

In Advanced Analytics settings, filter all rules configured in your system to quickly find a specific rules that fit certain criteria.


To view the existing rules in Advanced Analytics, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations** > **Exabeam Rules**.

To find a specific rule, filter all rules by:

- **Use Cases** – The use case or scenario the rule detects.
- **Rule Category** – The category the rule falls under.
- **Created By** – Who created the rule.
- **Updated By** – Who edited the rule.
- **Status** – Whether the rule is enabled or disabled.
- **Published** – Whether the rule is a **Draft** and has been edited but not reloaded, or **Published** and all changes are reloaded.
- **Trigger Frequency** – Whether the rule was ever triggered.


## Disable or Enable a Rule

Disable a rule so it doesn't trigger or enable a rule you previously disabled in Advanced Analytics settings.

1. From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**
2. Disable or enable an individual rule or multiple rules:
  - To disable or enable an individual rule, click the rule's **Status** toggle. To disable the rule, toggle the **Status** toggle off. To enable the rule, toggle the **Status** toggle on.
  - To disable or enable multiple rules, select the checkbox for the rules, then select **Enable** or **Disable**.

## Create a Fact-Based Rule

Create a [fact-based rule](#) in Advanced Analytics settings.

1. From the lower-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**.
2. Click **Create Rule**.
3. Enter specific information:
  - **Use Case & Scenario** – Click **SELECT SCENARIOS**, then select the [use cases](#) and [scenarios](#) the rule best detects.
  - **Rule Category** – From the list, select which category the rule falls under.
  - **Name** – Name the rule. When the rule triggers, the name is displayed in Advanced Analytics. It's best to be descriptive and indicate the nature of the risky behavior; for example, *Data Exfiltration by a Flight Risk User*.
  - **Description** – Describe the rule and provide additional details that may help your team investigate. To help your team better interpret what happens during a user session, describe why you created the rule and what it detects.
  - **What type of events should the rule evaluate?** – Select the type of events the rule evaluates. For example, if the rule detects user logins, select all event types that reflect those login events.

- **Risk Score** – Enter the risk score added to the session when the rule triggers.
4. Create a Boolean expression the Analytics Engine uses to determine if the rule triggers. Your rule triggers only if the expression is true.
    - a. Under **RULE EXPRESSION**, click **CREATE EXPRESSION**.
    - b. Under **Select Field**, select the event field the Boolean expression evaluates.
    - c. Under **Select Property**, select the property of the event field the Boolean expression evaluates. This differs based on event field.
    - d. Under **Select Function**, select an operator.
    - e. Under **Select Category**, select whether you're evaluating the event field property against another **Field** or a **Value**.
    - f. If you selected **Field**, under **Select Field**, select the event field the rule evaluates the first event field against. Under **Select Property**, select the property of the event field.
    - g. If you selected **Value**, in **Enter Value**, enter a string value.
    - h. To add additional conditions, select a boolean operator: **AND** or **OR**.
    - i. To save the boolean expression, click **DONE**.

5. (Optional) Define what other rules must or must not trigger for your rule to trigger:
  - a. Under **DEPENDENCY**, click **CREATE DEPENDENCY**.

**! IMPORTANT**

Default rules, including both model- and fact-based rules, may be deprecated and consequently disabled in future software updates. If you create a custom rule that includes a dependency on a default rule that becomes deprecated, your rule will be automatically disabled.

This also applies to dependencies on custom rules. If you create a rule that includes a dependency on a custom rule that becomes disabled, your rule will be automatically disabled.

- b. To define a rule that must not trigger for your to rule trigger, toggle **NOT** to the right  NOT . To define a rule that must trigger for your rule to trigger, toggle **NOT** to the left  NOT .
    - c. Under **Search for other rules.**, start typing, then select a rule from the list.
    - d. To add additional rules, select a boolean operator: **AND** or **OR**.
    - e. To save the dependency expression, click **DONE**.
6. Under **How many times should the rule be triggered?**, select how frequently the rule triggers: **Once per session**, **Always**, or **Once per value**.
7. Save the rule:
  - To save your progress without applying the changes, click **SAVE**. Your system validates the rule logic.



- To save the rule and apply the changes, click **SAVE & RELOAD ALL**. Your system validates the rule logic and reloads all rules.

#### Example 2. An Example of Creating a Fact Based Rule

You're creating a fact based rule that adds 15 to a user session's risk score every time a user your Human Resources team considers a flight risk starts a session. You have a context file titled *Flight Risk* containing the IDs of those users.

1. Enter specific information:
  - **Use Case & Scenario** – Click **SELECT SCENARIOS**, navigate to **Malicious Insiders > Abnormal Authentication & Access**, then select **Abnormal User Activity**.
  - **Rule Category** – Select **Asset Logon and Access**
  - **Name** – Enter *Flight Risks*.
  - **Description** – Enter *Users that HR considers flight risks*.
  - **What type of events should the rule evaluate?** – Select **remote-access, remote-logon, local-logon, kerberos-logon, ntlm-logon, account-switch, app-logon, app-activity, and privileged-object-access**.
  - **Risk Score** – Enter *15*.
2. Create a boolean expression:
  - a. Under **Select Field**, select **User**.
  - b. Under **Select Property**, select **User Label**.
  - c. Under **Select Function**, select **Equals**.
  - d. Under **Select Category**, select **Value**.
  - e. In **Enter Value**, enter **Flight Risk**. This is the label in the Flight context table.
  - f. Click **DONE**.
3. Under **How many times should the rule be triggered?**, select **Always**.
4. Click **SAVE & RELOAD**.

## Edit a Rule

Edit a rule using the Advanced Editor or Simple Editor in Advanced Analytics settings.



To edit any default rule or a default rule you cloned, you must use the Advanced Editor. The Advanced Editor is a JSON-style editor that displays the rule's back-end code as it exists in the configuration file `rules.conf`.

To edit a fact-based rule you created, you can use the Advanced Editor or the same interface you used to create the rule, also known as the Simple Editor. If you use the Advanced Editor to edit a fact-based rule, you can't edit the rule using the Simple Editor.

## Edit a Rule Using the Advanced Editor

Use the Advanced Editor to edit any rule. Keep in mind that if you use the Advanced Editor to edit a fact-based rule, you can't edit the rule using the Simple Editor.

You should use the Advanced Editor only if you're familiar with creating or tweaking a machine learning rule and understand the syntax language for expressing a rule. Changing rules can significantly affect the Analytics Engine. If you have questions, contact Exabeam Customer Success on the [Exabeam Community](#).

1. From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**
2. For the rule you're editing, click the More  menu, then select **Advanced Editor**.
3. Edit the [rule attributes](#). The only attribute you can't change in the Advanced Editor is the rule ID.

### **IMPORTANT**



Default rules, including both model- and fact-based rules, may be deprecated and consequently disabled in future software updates. If you edit a rule to include a dependency on a default rule that later becomes deprecated, the edited rule will be automatically disabled.

This also applies to dependencies on custom rules. If you edit a rule to include a dependency on a custom rule that becomes disabled, the edited rule will be automatically disabled.

4. Save the rule:
  - To save your progress without applying the changes, click **SAVE**. Your system validates the rule logic.
  - To save the rule and apply the changes, click **SAVE & RELOAD ALL**. Your system validates the rule logic and reloads all rules.

## Edit a Fact-Based Rule Using the Simple Editor

Use the Simple Editor to edit a fact-based rule you created.

1. From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**
2. For the fact-based rule you're editing, click the More  menu, then select **Simple Editor**.
3. [Edit](#) the rule details.

### **IMPORTANT**



Default rules, including both model- and fact-based rules, may be deprecated and consequently disabled in future software updates. If you edit a rule to include a dependency on a default rule that later becomes deprecated, the edited rule will be automatically disabled.

This also applies to dependencies on custom rules. If you edit a rule to include a dependency on a custom rule that becomes disabled, the edited rule will be automatically disabled.

4. Save the rule:
  - To save your progress without applying the changes, click **SAVE**. Your system validates the rule logic.
  - To save the rule and apply the changes, click **SAVE & RELOAD ALL**. Your system validates the rule logic and reloads all rules.

## Clone a Rule

Copy a rule, save it under a new name, and edit the new rule in Advanced Analytics settings.

1. From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**
2. For the rule you're cloning, click the More  menu, then select **Clone**.
3. Name the rule, then click **CLONE**. The rule appears in the list of rules, and its **Published** status is **Draft**.
4. To change the rule logic, [edit](#) the rule. Keep in mind that to apply any changes, you must reload rules.


## Reprocess a Rule

Reload all rules to publish changes or reprocess old data.

### **IMPORTANT**

When you create a new rule or edit an existing rule, you must reprocess rules to apply changes.

Apply changes starting from the present or retroactively to old data. Advanced Analytics reprocesses old data while actively processing and analyzing new events in real time.



1. From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**
2. Click **PUBLISH ALL RULES**.
3. To the changes in the rules reflected in old data, click **Reprocess historic data**. When you reprocess old data, the data can train models that trigger new rules that weren't on your system at the time. It takes longer for your system to reprocess old data.
4. Click **Apply Changes**.
5. To view the status of the job, navigate to **Admin Operations > Exabeam Engine > [Reprocess Jobs](#)**. When reprocessing is successful, the rule's **Published** status is **PUBLISHED**.

## Revert Out-of-the-Box Rules

Revert out-of-the-box rules you edited to its default configuration in Advanced Analytics settings.

You can only revert out-of-the-box rules you edited. You can't revert custom rules you created.

## Configure Rules

1. From the bottom-left side of the page, click **SETTINGS**  > **Analytics**, and then navigate to **Admin Operations > Exabeam Rules**
2. For the rule you're reverting, click the More  menu, select **Revert to Default**, then click **OK**.
3. To publish this change, [reprocess](#) all rules.

## Exabeam Threat Intelligence Service

The Exabeam Threat Intelligence Service delivers up-to-date threat indicators, on a daily basis, to Advanced Analytics deployments. Threat indicator data is stored in context tables that are associated with each feed. These threat indicators provide enhanced data based on curated threat intelligence.

The table below lists the categories of threat indicators provided by each threat intelligence feed and the rules that leverage each feed. For detailed tables mapping use cases and rules for each corresponding context table, see the Exabeam Community article: [TIS-populated Context Tables Mapped to Rules](#).



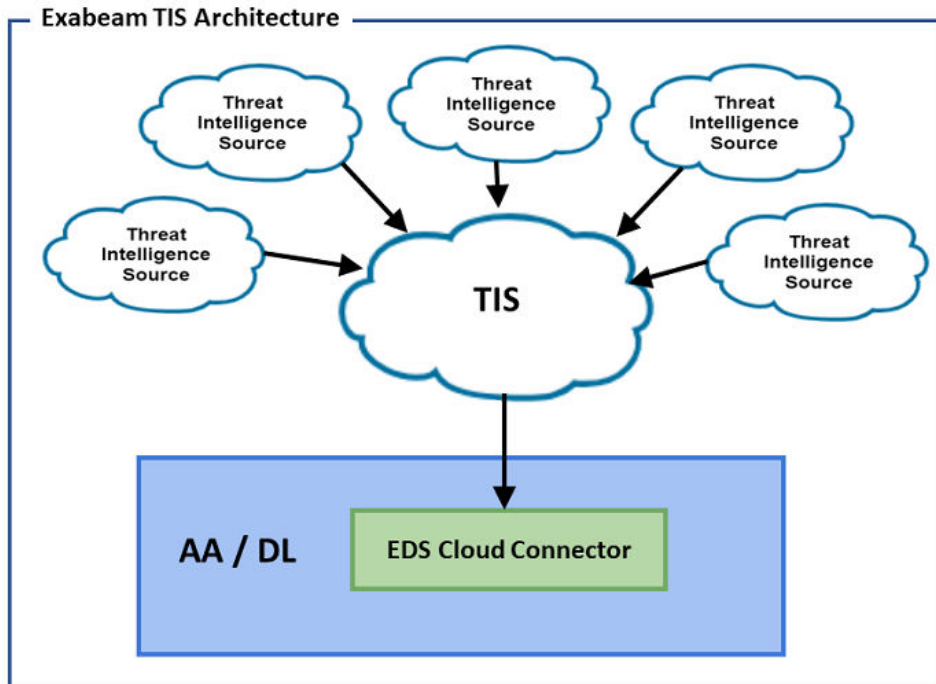
**NOTE**

All of the threat intelligence feeds, except the TOR network category, provide curated threat intelligence from ZeroFox. The TOR network feed is an open source data feed.

IoC Category	Rules
<p><b>Ransomware IP</b></p> <p>IP addresses associated with ransomware attacks</p>	<ul style="list-style-type: none"> <li>Auth-Ransomware-Shost</li> <li>Auth-Ransomware-Shost-Failed</li> <li>A-NET-Ransomware-IP</li> <li>A-NETF-Ransomware-IP</li> <li>WEB-UI-Ransomware</li> </ul>
<p><b>Threat IP</b></p> <p>IP addresses associated with ransomware or malware attacks</p>	<ul style="list-style-type: none"> <li>VPN02</li> <li>Auth-Blacklist-Shost</li> <li>Auth-Blacklist-Shost-Failed</li> <li>EPA-PI-ThreatIp</li> <li>A-NET-TI-IP-Outbound</li> <li>A-NETF-TI-IP-Outbound</li> <li>A-NET-TI-IP-Inbound</li> <li>A-WEB-Reputation-IP</li> <li>EPA-PI-ThreatIp</li> <li>WEB-UI-Reputation</li> </ul>
<p><b>Reputation Domain</b></p> <p>Domain names and URLs associated with sites that often contain malware, drive-by compromises, and more</p>	<ul style="list-style-type: none"> <li>WEB-UD-Reputation</li> <li>A-WEB-Reputation-Domain</li> <li>A-NET-TI-H-Outbound</li> <li>A-NETF-TI-H-Outbound</li> <li>A-NET-TI-H-Inbound</li> <li>A-DNS-MALDOM-QUERY</li> <li>A-DNS-MALDOM-RESPONSE</li> </ul>
<p><b>Web Phishing</b></p> <p>Domain names associated with phishing or ransomware</p>	<ul style="list-style-type: none"> <li>WEB-UD-Phishing</li> </ul>

IoC Category	Rules
<p><b>TOR IP</b></p> <p>IP addresses associated with the TOR network</p>	<ul style="list-style-type: none"> <li>• Auth-Tor-Shost-Failed</li> <li>• Auth-Tor-Shost</li> <li>• EPA-PI-Torlp</li> <li>• WEB-UI-Tor</li> <li>• A-NET-TOR-Outbound</li> <li>• A-NETF-TOR-Outbound</li> <li>• A-NET-TOR-Inbound</li> </ul>

Cloud-delivered deployments of Advanced Analytics and Data Lake connect to the Threat Intelligence Service (TIS) through an Exabeam Data Service (EDS) cloud connector, as shown in the image below. The cloud connector service provides authentication and establishes a secure connection to the Threat Intelligence Service. The cloud connector service collects updated threat indicators from the Threat Intelligence Service Rules and makes them available within Advanced Analytics and Data Lake on a daily basis.



The Threat Intelligence Service does not require a separate license. It is bundled with Advanced Analytics deployments. Additional installation is not required.

For on-premise deployments of Advanced Analytics and Data Lake, threat indicators are downloaded directly from the Threat Intelligence Service on a daily basis.

For more information about the Threat Intelligence Service, contact your technical account manager.

## Threat Intelligence Service Prerequisites

Before configuring Threat Intelligence Service, ensure your deployment meets the following prerequisites:

- At least 5 Mbps Internet connection
- Access to [https://\\*.exabeam.cloud](https://*.exabeam.cloud) over HTTPS port 443
- Access to [https://\\*.cloud.exabeam.com](https://*.cloud.exabeam.com) over HTTPS port 443
- DNS resolution for Internet hostnames (this will only be used to resolve to [https://\\*.cloud.exabeam.com](https://*.cloud.exabeam.com) and [https://\\*.exabeam.cloud](https://*.exabeam.cloud)).

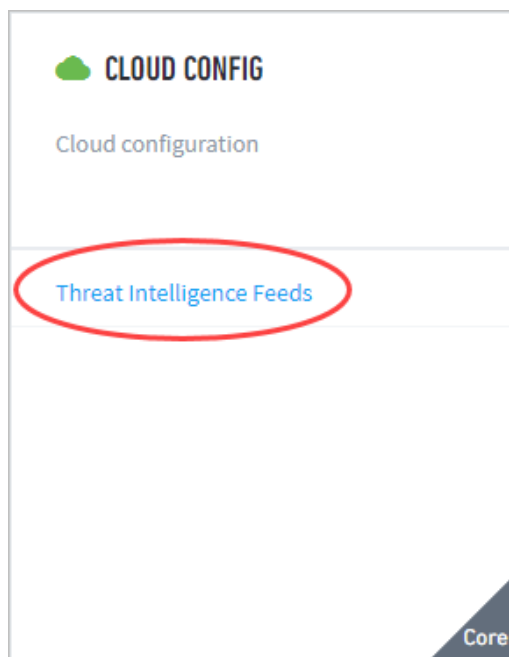


### NOTE

Ensure dynamic access is enabled as the IP address may change. Also, for this reason, firewall rules for static IP and port addresses are not supported.

## View Threat Intelligence Feeds

To view threat intelligence feeds in Advanced Analytics, open the **Settings** page. Navigate to the **Cloud Config** tile and select **Threat Intelligence Feeds**.



The **Threat Intelligence Feeds** page displays a list of the feeds provided by the cloud-based Exabeam Threat Intelligence service. The list includes the following information about each feed:

- **Type:** The type of feed (for example, domain list, IP list, etc.)
- **Name:** The name of the feed (given by the cloud-based service)
- **Description:** A short description of the feed
- **Context Tables:** The context tables associated with the feed

- **Status:** Indicates the availability of the feed in the cloud-based service
- **Updated:** The date and time the feed was last updated from the cloud service

**SETTINGS**

**CLOUD CONFIG** THREAT INTELLIGENCE FEEDS

5 Threat Intelligence Feeds 🔍

<input type="checkbox"/>	Type	Name	Description	Context Table(s)	Status	Updated
<input type="checkbox"/>	> Domain	Reputation Domains	List of reputation domains	reputation_domains	●	2019-02-26 22:46:26
<input type="checkbox"/>	> IP	TOR IPs	A list of TOR IPs	is_tor_ip	●	2019-02-26 22:46:26
<input type="checkbox"/>	> Domain	Phishing Domains	A list of Phishing Domains	web_phishing	●	2019-02-26 22:46:26
<input type="checkbox"/>	> IP	Ransomware IPs	A list of Ransomware IPs	is_ransomware_ip	●	2019-02-26 22:46:26
<input type="checkbox"/>	> IP	Malicious IPs	A list of Malicious IPs	is_ip_threat	●	2019-02-26 22:46:26

To view additional detailed information about a specific feed, click the arrow to the left of the feed. An additional view expands with more information, including **ID**, **Source URL**, **Indicator in Context Tables**, **Retrieved from Source**, and **Feed Indicator Sample**.

<input type="checkbox"/>	<input checked="" type="checkbox"/> Domain	Phishing Domains	A list of Phishing Domains	web_phishing	●	2022-03-21 11:25:00
		<b>Description</b>	<b>Feed Indicator Sample</b>		<b>Context Table(s)</b>	
		A list of Phishing Domains			web_phishing	
		<b>ID</b>	<b>Indicator in Context Tables</b>			
		web_phishing	0			
		<b>Source URL</b>	<b>Retrieved from Source</b>			
		/ti/feeds/web_phishing	2022-03-21 11:25:00			

For information about context tables and how they are related to threat intelligence feeds, see [Threat Intelligence Context Tables](#).

## Threat Intelligence Context Tables

Data provided by threat intelligence feeds is stored in context tables associated with each feed. By default, feeds are initially associated with existing context tables. As a result, when your Advanced Analytics deployment is connected to the Threat Intelligence Service, it immediately begins collecting threat intelligence data.

In Advanced Analytics, the data in context tables can be leveraged by creating rules that match log events to indicators stored in a threat intelligence context table. If the RuleExpression logic finds a match, an event can be identified as malicious without further analysis.

In Data Lake, the data in context tables can help to enrich log event data.



For more information about working with context tables, see the following:

- [View Threat Intelligence Context Tables](#)
- [Assign a Threat Intelligence Feed to a New Context Table](#)
- [Create a New Context Table for a Threat Intelligence Feed](#)

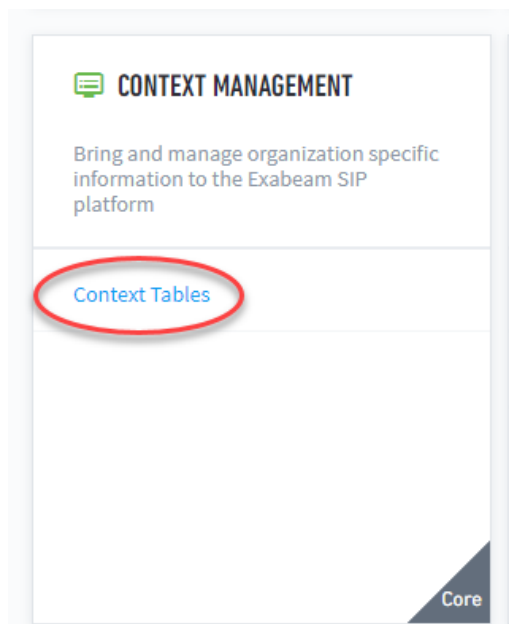


**NOTE**

To view a sample list of Threat Intelligence Service indicator sources see the [Exabeam Community](#).

## View Threat Intelligence Context Tables

To view the current context tables provided by the Threat Intelligence Service, log into your instance of Advanced Analytics and open the **Settings** page. Navigate to the **Context Management** tile and select **Context Tables**.



The **Context Tables** page displays a list of all the context tables currently provided by the Exabeam Threat Intelligence service. To locate a specific context table, scroll through the list or use the search feature 🔍.

**CONTEXT MANAGEMENT** ADD CONTEXT SOURCE   GENERATE CONTEXT   **CONTEXT TABLES**

Exabeam allows you to bring in custom context tables via CSV uploads or by adding objects one at a time. As an example, you can bring in your list of privileged users, or critical assets into Exabeam as context tables.

Object Type: All 🔍 +

Name	Object Type	Table Type	Label Value	Last Config Change	Connections
unix_user_id	Miscellaneous	Default		03/03/2022	Manual
web_malicious_categories	Miscellaneous	Default		03/03/2022	Manual
user_fullname	Users	Default		03/03/2022	Manual, Azure AD, All LDAP Servers
is_tor_proxy	Exabeam Threat Intel	TIS		03/03/2022	Manual, Exabeam Threat Intel service

To view information about keys and values associated with a specific context table, click the table name. A new expanded view of the table is displayed.

**CONTEXT TABLES > WINDOWS\_SERVICE\_TYPE** 🔍   📄   +

**DETAILS**

**Name:** windows\_service\_type  
**Object Type:** Miscellaneous  
**Type:** Key Value  
**Label:** None

---

**CONNECTIONS**

Connect this table to an LDAP Server to populate it with a filtered set of entries from that directory.

+ Add Connection

EXPORT AS CSV

	Source	Key	Value(s)
<input type="checkbox"/>	Manual	0X1	Kernel Driver - A kernel device driver such as a hard disk or other low-level hardware device driver
<input type="checkbox"/>	Manual	0x2	File System Driver - A file system driver, which is also a Kernel device driver
<input type="checkbox"/>	Manual	0X8	Recognizer Driver - A file system driver used during startup to determine the file systems present on the system


## Assign a Threat Intelligence Feed to a New Context Table

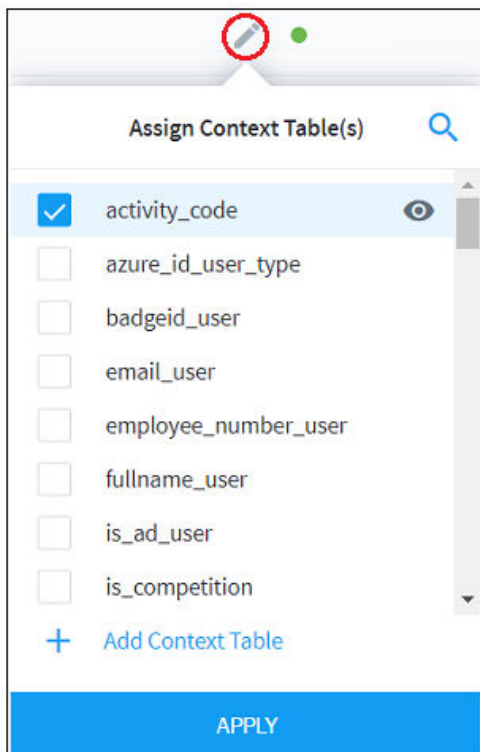
Some threat intelligence feeds are pre-assigned to specific context tables. However, you can easily add, remove, or change feed assignments. You can configure feed assignments in one of two ways, individually or in bulk.


**NOTE**  
 You cannot unassign default context table mappings.

### Individual Feed Assignment

To change the assignment of a single threat intelligence feed to one or more context tables:

1. Navigate to the Threat Intelligence Feeds page, as described in [View Threat Intelligence Feeds](#).
2. Find the feed whose context table assignments you want to change and, in the **Status** column, click **edit** . A list of the available context tables opens.



3. Use the check boxes on the left of each context table to assign or unassign the threat intelligence feed. A single feed can be assigned or unassigned to multiple context tables.
4. To view the existing threat indicators in a specific context table, click view . A new window opens and displays a list of keys and values for the indicators included in the context table. Click **OK** to close the window.
5. When you've finished assigning or unassigning the feed to specific context tables, click **Apply** to save the updated assignments.

### Bulk Feed Assignment

To change the assignment of multiple threat intelligence feeds to one or more context tables:

1. Navigate to the Threat Intelligence Feeds page, as described in [View Threat Intelligence Feeds](#).
2. Use the check boxes on the left of each feed to select multiple feeds whose assignment you want to change.

<input type="checkbox"/>	Type	Name	Description
<input type="checkbox"/>	> Domain	TOR Proxy	A list of TOR Proxy
<input checked="" type="checkbox"/>	> Domain	Dynamic DNS Domains	A list of Dynamic DNS Domains
<input checked="" type="checkbox"/>	> Domain	Reputation Domains	List of reputation domains

- At the top of the feeds list, click **Assign** or **Unassign**, depending on what changes you want to make.
  - Assign:** A list of the available context tables opens in a new window. Use the check boxes on the left to select context tables. To see the indicators included in each table, click **view** . When you've completed your table selections, click **Assign**. All of the specified feeds will be assigned to the selected context tables.
  - Unassign:** All of the specified feeds will be unassigned from their context tables.

## Create a New Context Table from a Threat Intelligence Feed

New context tables can be created from specific threat intelligence feeds. You can create new context tables in one of two ways, from an individual feed or from multiple feeds in bulk.

### Create a Table from a Single Feed

To create a new context table from a single threat intelligence feed:

- Navigate to the Threat Intelligence Feeds page, as described in [View Threat Intelligence Feeds](#).
- Find the feed from which you want to create a new content table and, in the **Status** column, click **edit** . A list of the existing context tables opens.
- At the bottom of the list, select the **Add Context Table** option. A set of options for creating a new context table is displayed.

4. Enter the **Title**, **Object Type**, and **Type** information to define the new context table.
5. Click **Add** to save the new context table.

### Create a Table from Multiple Feeds

To create a new context table from a bulk selection of threat intelligence feeds:

1. Navigate to the Threat Intelligence Feeds page, as described in [View Threat Intelligence Feeds](#).
2. Use the check boxes on the left of each feed to select multiple feeds from which you want to create a new context table.
3. At the top of the feeds list, click **Assign**. A list of the existing context tables opens.

Assign Unassign 2 of 12 Threat Intelligence Feeds			
<input type="checkbox"/>	Type	Name	Description
<input type="checkbox"/>	> Domain	TOR Proxy	A list of TOR Proxy
<input checked="" type="checkbox"/>	> Domain	Dynamic DNS Domains	A list of Dynamic DNS Domains
<input checked="" type="checkbox"/>	> Domain	Reputation Domains	List of reputation domains

- At the bottom of the list, select the **Add Context Table** option. A set of options for creating a new context table is displayed.

### CREATE NEW CONTEXT TABLE

Exabeam allows you to bring in custom context tables via CSV uploads or by adding objects one at a time. As an example, you can bring in your list of privileged users, or critical assets into Exabeam as context tables.

**Title \***

**Object Type:**

Users

Assets

Miscellaneous

**Type**

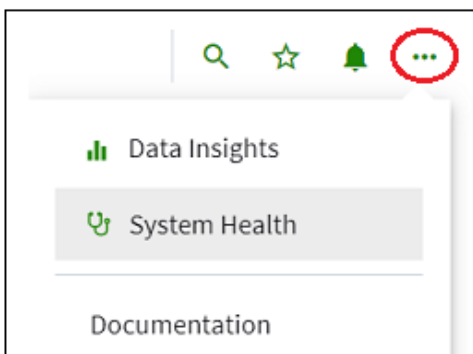
CANCEL
ADD

- Enter the **Title**, **Object Type**, and **Type** information to define the new context table.
- Click **Add** to save the new context table.

## Check ExaCloud Connector Service Health Status

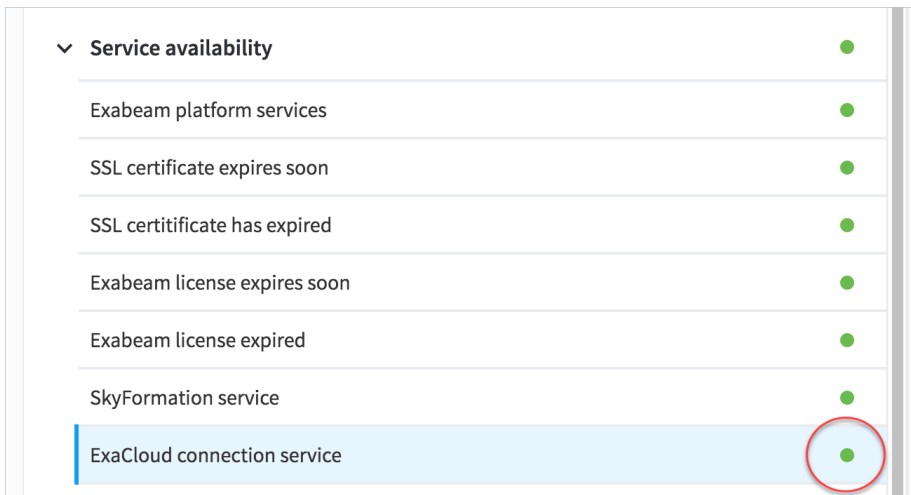
To view the current status of the ExaCloud connector service:

- Log in to your instance of the UI.
- Click the top-right menu icon and select **System Health**.



- Select the **Health Checks** tab.

4. Click **Run Checks**.
5. Expand the **Service Availability** section, and then review the **ExaCloud connection service availability** icon.



The service availability icon shows the current health of the Cloud Connector service that is deployed on your Exabeam product.

- **Green** – The cloud connector service is healthy and running on your on-prem deployment.

**NOTE**

The green icon does not specifically indicate the cloud connector is connecting to the cloud and pulling Threat Intelligence Service data. It only indicates the cloud connector service is up and running.

- **Red** – The cloud connector service has failed. Please contact Exabeam Customer Success by opening a case from [Community.Exabeam.com](https://community.exabeam.com).

## Exabeam Cloud Telemetry Service

The Exabeam telemetry service collects and transmits valuable quality and health metrics to Exabeam Cloud. The transmitted data, which includes system events, metrics, and environment health data, provides insights into system issues and application availability. Examples of system issues include processing downtime such as processing delays and storage issues.

See also [Monitoring](#) in the the Exabeam Security Operations Platform Administration Guide.



## Manage Security Content in Advanced Analytics

Install, get updates on, uninstall, and upload content packages in Advanced Analytics settings.

Manage all your content packages directly in Advanced Analytics settings, under **Admin Operations > Additional Settings > Content Updates**, where you retrieve the latest available content packages from the cloud in real time, including both general Exabeam releases and custom fixes you request.

In these settings, a content package that includes custom fixes you requested is called a *custom package*. A content package from a general Exabeam release is called a *default package*. It's important that you update your content with each release because the release may contain new parsers and event builders, support new log sources and vendors, or include other additions and fixes that keep your system running smoothly.

If you have an environment that can access the internet, you can pull the latest content packages [manually](#) or [automatically](#), select a specific content packages to [install](#), or even [schedule](#) content packages to automatically install on a daily or weekly basis, all from the cloud.

If you have an environment that can't access the internet, you can't connect to the cloud. You must view and download the latest content packages from the [Exabeam Community](#), then [upload](#) them.

You can only install and upload content packages that contain event builders or parsers.

### Manually Install a Content Package


Install a new content package directly from Advanced Analytics settings onto your system.

Select a content package to install from a list of the latest, available content packages. If your environment can't access the internet, you can't install content packages from the cloud. Instead, download the content package from the [Exabeam Community](#) or your case ticket, then manually [upload](#) it.

A content package from a general Exabeam release is called a *default package*. It's important that you update your content with each release because the release may contain new parsers and event builders, support new log sources and vendors, or include other additions and fixes that keep your system running smoothly. You can upload multiple default content packages, but only install one default package at a time.

A content package that includes custom fixes you requested is called a *custom package*. You can upload and install any number of custom packages.

You can only install custom content packages that contain event builders or parsers.

1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **ADMIN OPERATIONS**, select **Content Updates**.


3. To install a default content package, click the **DEFAULT PACKAGES** tab. To install a custom content package, click the **CUSTOM PACKAGES** tab.
4. Click **INSTALL**.  
If the package is a default content package and a newer version of one you previously installed, this newer version replaces the older version. You can no longer view or install the older version.  
If the package is a custom content package and a newer version of one you already installed, ensure that you [uninstall](#) the older version.

## Automatically Install Content Packages

Schedule Advanced Analytics to automatically check for and install new content packages on a daily or weekly basis.

If you have an environment that can't access the internet, you can't install content packages from the cloud. Instead, download a content package from the [Exabeam Community](#) or your case note, then manually [upload](#) it.

Only content packages that contain event builders or parsers are available.



1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **ADMIN OPERATIONS**, select **Content Updates**.
3. Click **Install Schedule**, then toggle **Auto Install** on.
4. After **Install package**, select the day of the week when Advanced Analytics downloads new content.
5. After **at**, select the time when Advanced Analytics downloads new content.
6. Click **SAVE**. If newer versions of custom content packages were installed, ensure that you [uninstall](#) the older version.

## Manually Check for New Content Packages

Manually fetch the latest available content packages. You can also set Advanced Analytics to [automatically](#) check for new packages every 30 minutes.

If you have an environment that can't access the internet, you can't connect to the cloud to view the latest, available content packages. Instead, check the [Exabeam Community](#) for the latest content packages. If you manually refresh the list, Advanced Analytics says you have no new packages.

Only content packages that contain event builders or parsers are available.

1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **ADMIN OPERATIONS**, select **Content Updates**.
3. Click refresh . Advanced Analytics checks for new default and custom content packages and updates both lists.


## Automatically Check for New Content Packages

Set Advanced Analytics to automatically check for new content packages and fetch them every 30 minutes.

This setting automatically checks for new content packages but doesn't install them. To automatically install them, you must [schedule](#) it separately.

If you have an environment that can't access the internet, you can't connect to the cloud to view the latest, available content packages. Instead, check the [Exabeam Community](#) for the latest content packages.

Only content packages that contain event builders or parsers are available.

1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **ADMIN OPERATIONS**, select **Content Updates**.
3. Click **Last Update Checked**, toggle **Auto Updates** on, then click **SAVE**. Advanced Analytics checks for new content packages every 30 minutes and updates the list.

## Upload a Content Package



If you have an environment that can't access the internet, download new content from the [Exabeam Community](#) or your case ticket, then upload it directly in Advanced Analytics settings.

If you can access the internet in your environment, you see a list of the latest available content packages in the cloud. Directly [install](#) a content package from this list instead of uploading packages to Advanced Analytics.

A content package from a general Exabeam release is called a *default package*. It's important that you update your content with each release because the release may contain new parsers and event builders, support new log sources and vendors, or include other additions and fixes that keep your system running smoothly. You can upload multiple default content packages, but only install one default package at a time.

A content package that includes custom fixes you requested is called a *custom package*. You can upload and install any number of custom packages.

You can only upload content packages that contain event builders or parsers.


1. Ensure that you download the content package from the [Exabeam Community](#) or your case ticket.
2. In the sidebar, click **SETTINGS** , then select **Core**.
3. Under **ADMIN OPERATIONS**, select **Content Updates**.
4. To upload a content package from a general Exabeam content release, click **DEFAULT PACKAGES**. To upload a content package you requested, click **CUSTOM PACKAGES**.
5. Click upload .
6. Click **UPLOAD THE PACKAGE**, then select a ZIP file to upload, up to 100 MB.

7. Click **SAVE**. Your content package is uploaded.
8. Locate your content package, then click **INSTALL**. Your content package is installed. If the package is a custom content package and a newer version of one you already installed, ensure that you [uninstall](#) the older version.

## Uninstall a Custom Content Package

Uninstall a custom content package if there's an issue with the package or to remove an older version of a package once you [upload](#) a newer version.

A content package from a general Exabeam release is called a *default package*. A content package that includes custom fixes you requested is called a *custom package*. You can only uninstall a custom content package, not a default content package. To remove a default content package, you must [install](#) another default content package.

1. In the sidebar, click **SETTINGS** , then select **Core**.
2. Under **ADMIN OPERATIONS**, select **Content Updates**.
3. Click the **CUSTOM PACKAGES** tab, then next to the content package, click **UNINSTALL**.

## Health Status Page

The Health Status page offers an on-demand assessment of the Exabeam pipeline. The assessment has three categories:

**General Health:** General health tests that all of the back-end services are running - database storage, log feeds, snapshots, CPU, and memory.

**Connectivity:** Checks that Exabeam is able to connect to external systems, such as LDAP and SIEM.

**Log Feeds:** This section reports on the health of the DC, VPN, Security Alerts, Windows Servers, and session management logs.

In all of the categories, the statuses are color-coded as follows: GREEN = good, YELLOW = warning, and RED = critical.

Located on the homepage are the Proactive Health Checks that alert administrations when:

- Any of the core Exabeam services are not running
- There is insufficient disk storage space
- Exabeam has not been fetching logs from the SIEM for a configurable amount of time


### Proactive and On-Demand System Health Checks

System Health is used to check the status of critical functionality across your system and assists Exabeam engineers with troubleshooting. Exabeam provides visibility on the backend data pipeline via **Health Checks**. Graphs and tables on the page visually represent the health status for all of the key systems, as well as indexes and the appliance, so you are always able to check statuses at a glance and track health over time.

Proactive health checks run automatically and periodically in the background.

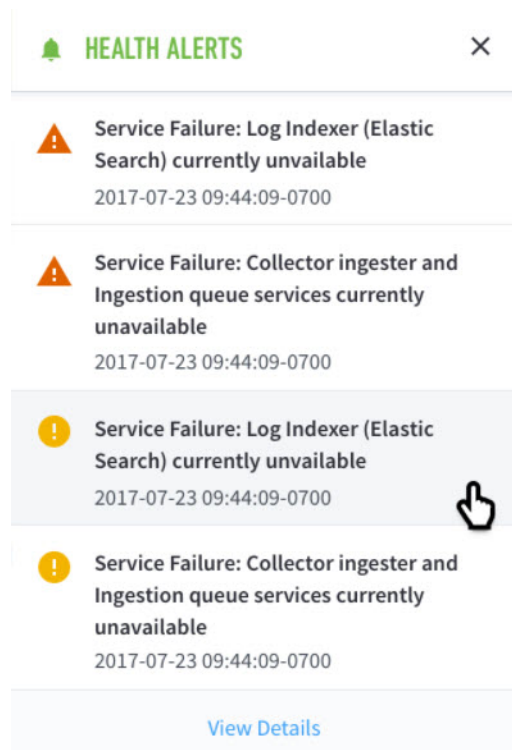
On-demand health checks can be initiated manually and are run immediately. All newly gathered health check statuses and data is updated in the information panes on the page. All proactive and on-demand health checks are listed on the **Health Checks** page. Proactive health checks are visible by any user in your organization. Only users with administrator permission can reach the page.

Figure 1. Exabeam Notification Icon



When a health check is triggered, a notification message is displayed in the upper right corner of the UI. Select the alert icon to open a side panel that lists the alerts and provides additional detail. A panel listing all notifications is expanded.

Figure 2. Health Alerts panel



These alerts are also listed under the Health Alerts tab in the System Health page. In general:

- **Warning**: There is an issue that should be brought to the attention of the user.
- **Critical**: Immediate action is recommended. In all cases, if an alert is raised on your system, please contact Exabeam Customer Success.

To reach the **Health Checks** page, navigate to the **System Health** page from the Settings tab at the top right corner of any page, then select the **Health Checks** tab.

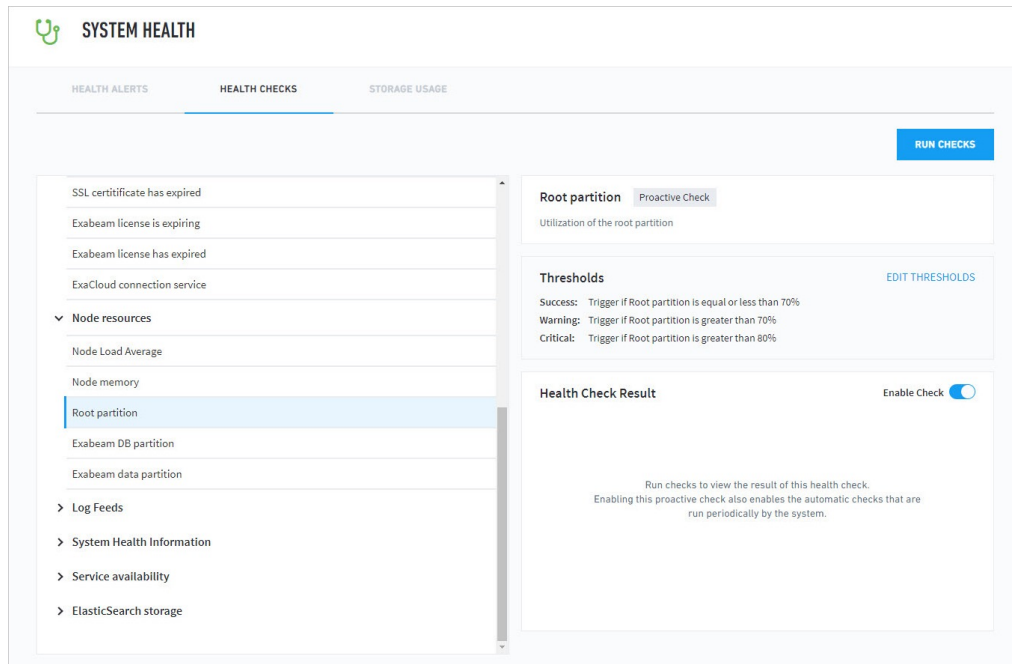
Health check categories are:

- **Service Availability** – License expiration, database, disaster recovery, Web Common application engine, directory service, aggregators, and external connections
- **Node Resources** – Load, performance, and retention capacity
- **Service Availability** (Incident Processors and Repositories) – IR, Hadoop, and Kafka performance metrics

### Advanced Analytics Specific Health Checks

- **Log Feeds** – Session counts, alerts, and metrics
- **System Health Information** – Core data and operations processor metrics
- **Elasticsearch Storage** (Incident Responder) – Elasticsearch capacity and performance metrics

Figure 3. System Health - Advanced Analytics Health Checks page



### Configure Alerts for Worker Node Lag

When processing current or historical logs, an alert will be triggered when the worker node is falling behind the master node. How far behind can be configured in `/opt/config/exabeam/tequila/custom/health.conf`. The parameters are defined below:

- `RTModeTimeLagHours` - During real-time processing the default setting is 6 hours.
- `HistoricalModeTimeLagHours` - During historical processing the default setting is 48 hours.
- `syslogIngestionDelayHour` - If processing syslogs, the default setting is 2 hours.

```

}

slaveMasterLagCheck {
    printFormats = {
        json = "{ \"lagTimeHours\": \"$lagTimeHours\", \"masterRunDate\": \"$masterRunDate\", \"slaveRunDate\": \"$slaveRunDate\", \"isRealTimeMode\": \"$isRealTimeMode\"}"
        plainText = "Worker nodes processing lagging by more than $lagTimeHours hours. Is in real time: $isRealTimeMode "
    }
}

RTModeTimeLagHours = 6
    
```

```

HistoricalModeTimeLagHours = 48
}

limeCheck {
    syslogIngestionDelayHour = 1
}

```

## System Health Checks

**Martini Service Check:** Martini is the name Exabeam has given to its Analytics Engine. In a multi-node environment, Martini will be the Master node.

**Tequila Service Check:** Tequila is the name Exabeam has given to its User Interface layer.

**Lime Service Check:** LIME (Log Ingestion and Message Extraction) is the service within Exabeam that ingests logs from an organization's SIEM, parses and then stores them in HDFS. The main service mode parses message files and creates one message file per log file. This mode is used to create message files that will be consumed by the main node.

**Mongo Service Check:** MongoDB is Exabeam's chosen persistence database. A distributed MongoDB system contains three elements: shards, routers, and configuration servers (configsvr). The shards are where the data is stored; the routers are the piece that distributes and collect the data from the different shards; and the configuration servers which tracks where the various pieces of data are stored in the shards.

**Zookeeper Service Check:** Zookeeper is a centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services. In a distributed multi-node environment, we need the ability to make a value change inside one process on a machine and have that change be seen by a different process on a different machine. Zookeeper provides this service.

**Hadoop Service Check:** Master - Hadoop is Exabeam's distributed file system, where the raw logs and parsed events are stored. These files are available to all nodes.

**Ganglia Service Check:** Ganglia is a distributed monitoring system for computing systems. It allows us to view live or historical statistics for all the machines that are being monitored.

**License Checks:** The status of your Exabeam license will be reported in this section. This is where you will find the expiration date for your current license.

## Alerts for Storage Use

Available on the System Health page, the Storage Usage tab provides details regarding the current data retention settings for your Advanced Analytics deployment. Advanced Analytics sends notifications when available storage capacity dwindles to a critical level. Admins have the option to enable and configure automatic data retention and data purging for both HDFS and MongoDB usage.

For more information on data retention, see [Data Retention in Advanced Analytics](#).



## Default Data Retention Settings

Default Advanced Analytics data retention settings depend on the license you purchased. For licensing information, see [Product Entitlements](#) on the Exabeam Community site.

## System Health Alerts for Low Disk Space

Get notified when your disk is running low on space.

Your Advanced Analytics system may go down for several hours when it upgrades or restarts. During this down time, your log source continues to send logs to a disk, accumulating a backlog of unprocessed logs. Log Ingestion and Messaging Engine (LIME) tries to process this backlog when it starts running again. If your log source sends logs faster than what your system size can handle, LIME may struggle to process these logs, run out of disk space, and stop working correctly.

Your Advanced Analytics system already uses mechanisms, like compressing files, to conserve as much space as possible. If your disk is still running out of space, you receive system health alerts.

If LIME is running, you receive two system health alerts. When the disk has 25 percent capacity remaining, the first health alert notifies you that you're running low on disk space. In the rare case that your disk has 15 percent capacity remaining, a second health alert notifies you that your system has deleted files, starting with the largest one, as a last resort to keep your system running.

If LIME goes down, Advanced Analytics can't send health alerts. When LIME is running again, you may receive a belated health alert. In the rare case that your disk reached 15 percent remaining capacity while LIME was down, this health alert notifies you that your system has deleted files, starting with the largest one, as a last resort to keep your system running.

When you receive these health alerts, consider tuning your system so it ingests less logs or ingests logs more slowly. If you ingest logs from Data Lake, consider [setting](#) a lower log forwarding rate.

## System Health Alerts for Paused Parsers

Get notified automatically when Advanced Analytics pauses a parser.

To protect your system from going down and ensure that it continues to process data in real time, Advanced Analytics [detects](#) when a parser is taking an abnormally long time to parse a log, then pauses it.

When Advanced Analytics pauses a slow parser, you receive a health alert that describes which parser was paused, on which node, and recommended actions you could take. You don't receive a health alert when Advanced Analytics pauses a stuck or failed parser.

When Advanced Analytics pauses additional slow parsers, it resolves the existing health alert and sends a new alert that lists all paused slow parsers. Advanced Analytics also resolves an existing health alert any time you restart the Log Ingestion and Messaging Engine (LIME) and resume a paused parser. For example, if you change the threshold for when parsers get paused, you must restart LIME, which resumes any paused parsers and prompts Advanced Analytics to resolve any existing health alerts. If a resumed parser continues to perform poorly, Advanced Analytics pauses it again and sends another health alert about that parser.

## View Storage Usage and Retention Settings

By default, Advanced Analytics implements data retention settings for logs and messages. This allows the system to automatically delete old data, which reduces chances of performance issues due to low repository space.

Available on the **System Health** page, the **Storage Usage** tab provides details regarding the current data retention settings for your Advanced Analytics deployment, including:

- **HDFS Data Retention** – Number of days the data has been stored, the number of days remaining before it is automatically deleted, and a toggle to enable/disable auto-deletion of old data.



### NOTE

If enabling or disabling **Auto-delete Old Data**, you must also [restart](#) the Log Ingestion Engine before the change goes into effect.

A warning message appears if the average rate of HDFS usage exceeds the intended retention period.

- **HDFS Usage** – Total HDFS usage, including specific breakdowns for logs, events on disk, and events in database .



### NOTE

The Volume field displays the compressed index size on disk for storage planning. This differs from the total consumption of daily ingested logs used for billing.

- **MongoDB Retention** – A dialog to set retention settings and a toggle to enable/disable disk-based deletion of old data.  
You can edit the MongoDB data retention settings by clicking the pencil icon.  
The capacity used percentage threshold is set to 85% by default, with a maximum value of 90%. It helps to prevent MongoDB from reaching capacity before hitting the default retention period threshold. As soon as MongoDB meets the percentage on any node, the system will start purging events until it is back below the capacity used threshold.  
The maximum length of your retention period depends on the license you purchased. For licensing information, see [Product Entitlements](#) on the Exabeam Community site.



### NOTE

The **Days for Triggered Rules & Sessions** value cannot be less than the **Days for Events** value.

- **MongoDB Usage** – Total MongoDB usage.

## Set Up System Optimization

This tab is a single aggregated page for auditing and viewing disabled data types (including models, event types, and parsers) and system load redistribution.

## Disabled Models

When a model takes up too much memory, it is disabled. If you enable these models, the system may suffer performance issues.



### NOTE

Exabeam disables models as a last resort to ensure system performance. We have other defensive measures to prevent models from using all of the system's memory. These measures include enabling model aging and limiting the model bin size count. If these safeguards cannot prevent a model from consuming too much memory, the system prevents itself from shutting down as it runs out of memory.

## Paused Parsers

If your parser takes too long to parse a log and meets certain conditions, Advanced Analytics pauses the parser.

To keep your system running smoothly and processing data in real time, Advanced Analytics detects parsers that are performing poorly, then pauses them. Your parsers may perform poorly because:

- [The parser is slow and is taking too long to parse a log.](#)
- [The parser is stuck. It contains incorrect regular expressions, so it goes into an infinite loop or fails with a non-timeout exception; or, the input data is incorrect and the parser can't parse the data.](#)

In both cases, your system calculates whether the parser exceeds configured thresholds and meets certain conditions, then pauses the parser.

When a parser meets the conditions on a Log Ingestion and Messaging Engine (LIME) node, your system pauses the parser only on that node. If you have multiple LIME nodes, it is not automatically paused on all nodes unless it meets these conditions on every node.

When Advanced Analytics pauses a parser on any node, the parser appears in a [list](#) of paused parsers. You receive a [system health alert](#) only for paused slow parsers, not stuck or failed parsers.

### **Conditions for Pausing Slow Parsers**

To identify a slow parser, your system places the parser in a cache. Every configured period, `OutputParsingTimePeriodInMinutes` (five minutes by default), it calculates how long it takes, on average, for the parser to parse a log. It compares this average to a configurable threshold, `ParserDisableThresholdInMills`, in `lime.conf`.

To calculate a percentage of how much the parser makes up of the total parsing time, your system divides the previously-calculated average by the total time all parsers took to parse an event in the same five minute period. It compares the percentage to another configurable threshold, `ParserDisableTimePercentage`, in `lime.conf`.

Your system conducts a second round of checks if the parser meets certain conditions:

- The average time it takes for the parser to parse a log exceeds a threshold, `ParserDisableThresholdInMills` (seven milliseconds by default).
- The parser constitutes more than a certain percentage, `ParserDisableTimePercentage` (50 percent by default), of the total parsing time of all parsers.

During another five-minute period, your system checks the parser for a second time. If the parser meets the same conditions again, your system pauses the parser.

If you have a cloud-delivered deployment, contact Exabeam Customer Success to configure these variables.

### Conditions for Pausing Stuck Parsers

To identify a stuck parser, your system measures how long it takes for a parser to parse a log. If the time exceeds a threshold, `StuckParserWaitTimeoutMillis` (100 milliseconds by default), the parser fails with a timeout exception. Your system logs the error at a `DEBUG` security level and notes the parser in internal error statistics.

In each configured period, `ParserMaxErrorTimeWindowForChecksMillis` (9000 milliseconds, or 15 minutes, by default), your system checks the internal error statistics for any parsers that have accumulated a certain number of errors. If the errors exceed a threshold, `ParserMaxErrorNumberThreshold` (100 errors by default), the parser is paused and removed from the internal error statistics.

If you have a cloud-delivered deployment, contact Exabeam Customer Success to configure these variables.

### View Paused Parsers

To keep your system running smoothly and processing data in real time, Advanced Analytics detects slow or inefficient parsers, then pauses them. View all paused parsers in Advanced Analytics under **System Health**.

1. In the sidebar, click the menu **☰**, then select **System Health**.
2. Select the **System Optimization** tab.
3. Select **Paused Parsers**.



View a list of paused parsers, sorted alphabetically by parser name, and information about them, including:

- **Parser Name** – The name of the paused parser.
- **Average Log Line Parse Time** – Average time the parser took to parse each event.
- **Paused Time** – Date and time when the parser was paused.

## Disabled Event Types

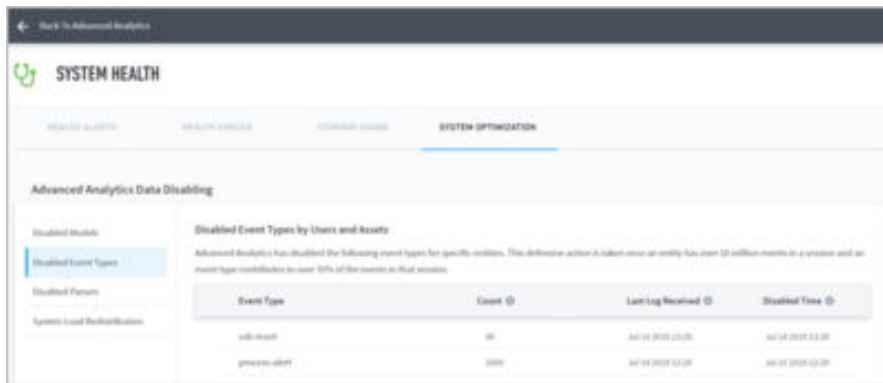
When a high volume user or asset amasses a large number of events of a certain event type, and that event type contributes to a large portion of the overall event count for that user (typically 10M+ events in a session) the event type is automatically disabled and listed here.

 **NOTE**

You are also shown an indicator when Advanced Analytics determines that the event type is problematic and disables it for the entity. The affected User/Asset Risk Trend and Timeline accounts for the disabled event type by displaying statistics only for the remaining events.

Disabled event types are displayed on the **System Optimization** tab of the **System Health** page. You can see a list of all event types that have been disabled, along with the users and assets for which they have been disabled for.

Figure 4. System Health - System Optimization menu



The **Disabled Event Type by Users and Assets** table is sorted first alphabetically by event type, then sorted by latest update timestamp.

The table includes columns with the following categories:

- **Event Type** – The disabled event type.
- **Count** – Last recorded total number of events for this entity.
- **Last Log Received** – Date and time of the event that triggered the disabling of this event type for the specified entity.
- **Disabled Time** – Date and time for when the event type was disabled for this entity.

## Automatically Redistribute System Load

Exabeam can automatically identify overloaded worker nodes, and then take corrective action by evenly redistributing the load across the cluster.

This redistribution is done by measuring and comparing job completion time. If one node finishes slower by 50% or more compared to the rest of the nodes, then a redistribution of load is needed. The load is then scheduled to be rebalanced by event categories.

You can enable automatic system load redistribution on the **System Optimization** tab of the **System Health** page. This option is enabled by default. Doing so allows the system to check the load distribution once a day.



**NOTE**

It is not recommended that you disable the system rebalancing option as it will result in uneven load distribution and adverse performance impacts to the system. However, if you choose to do so, you can configure manual redistribution to avoid such impacts.

You must restart the Exabeam Analytics Engine for any changes to **System Rebalancing** to take effect.

The **System Load Redistribution** tab shows an indicator when a redistribution of load is needed, is taking place, or has completed. The rebalancing process can take up to two hours. During this time you may experience slower processing and some data may not be available. However, the system resumes normal operations once redistribution is complete.

## Automatic Shutdown

If the disk usage on an appliance reaches 98%, then the Analytics Engine and log fetching will shut down automatically. Only when log space has been cleared and usage is below 98% will it be possible to restart.

Users can restart either from the **CLI** or the **Settings** page.