exabeam

Exabeam Advanced Analytics Release Notes

Exabeam Security Management Platform - Version SMP 2019.1 (I48)

Publication date August 5, 2020

Exabeam
2 Waters Park Dr. Suite 200
San Mateo, CA 94403

650.209.8599

Have feedback on this guide? We'd love to hear from you!
Email us at docs@exabeam.com

Disclaimer: Please ensure you are viewing the most up-to-date version of this guide by visiting the Exabeam Community.

Table of Contents

1. What's New

1.1. Google Pub/Sub Log Source Support

If your Advanced Analytics deployment is running on Google Cloud Platform and aggregates all of your logs into Google Pub/Sub, you can now configure Google Cloud Pub/Sub as a log source for Advanced Analytics.

For more information on adding log sources in Advanced Analytics, please refer to the *Advanced Analytics Admin Guide*.

1.2. Support for Granular Rule Reprocessing

When adding new or managing existing Exabeam rule on the **Exabeam Rules** page, you can choose to reload individual or all rules. You can now choose to apply rule changes and reprocess historic data. When applying and reprocessing rule changes to historic data, the reprocess is done in parallel with active, live processing. It does not impede or stop any real-time analysis.

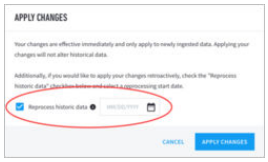

1.3. Auditing User Activities within Exabeam

Advanced Analytics now logs specific activities related to administrators and users of the product, including activities within the UI as well as configuration and server changes, too. This is primarily needed to monitor administrator and analyst activities from a compliance perspective. Advanced Analytics activity data is collected and sent to the syslog destination of your choosing.

For more information on accessing Advanced AnalyticsData Lake log activity, please refer to the Audit Actions within Exabeam section within the Advanced AnalyticsData Lake *Admin Guide*.

1.4. Changes to Exabeam Analytics Engine

We've made several changes to improve the usability of the Exabeam Analytics Engine.

You can now cancel in-progress reprocessing jobs. This is especially helpful if a particularly large reprocessing job slows the entire system, if you accidentally initiated a reprocessing job, or if you simply want to cancel the job for any other reason.

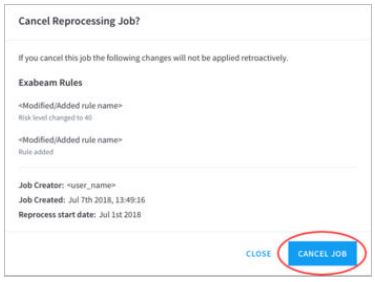

Additionally, we've updated the **Reprocessing Jobs** table to provide more details and control of your complete, pending, in progress, canceled, or failed reprocessing jobs.

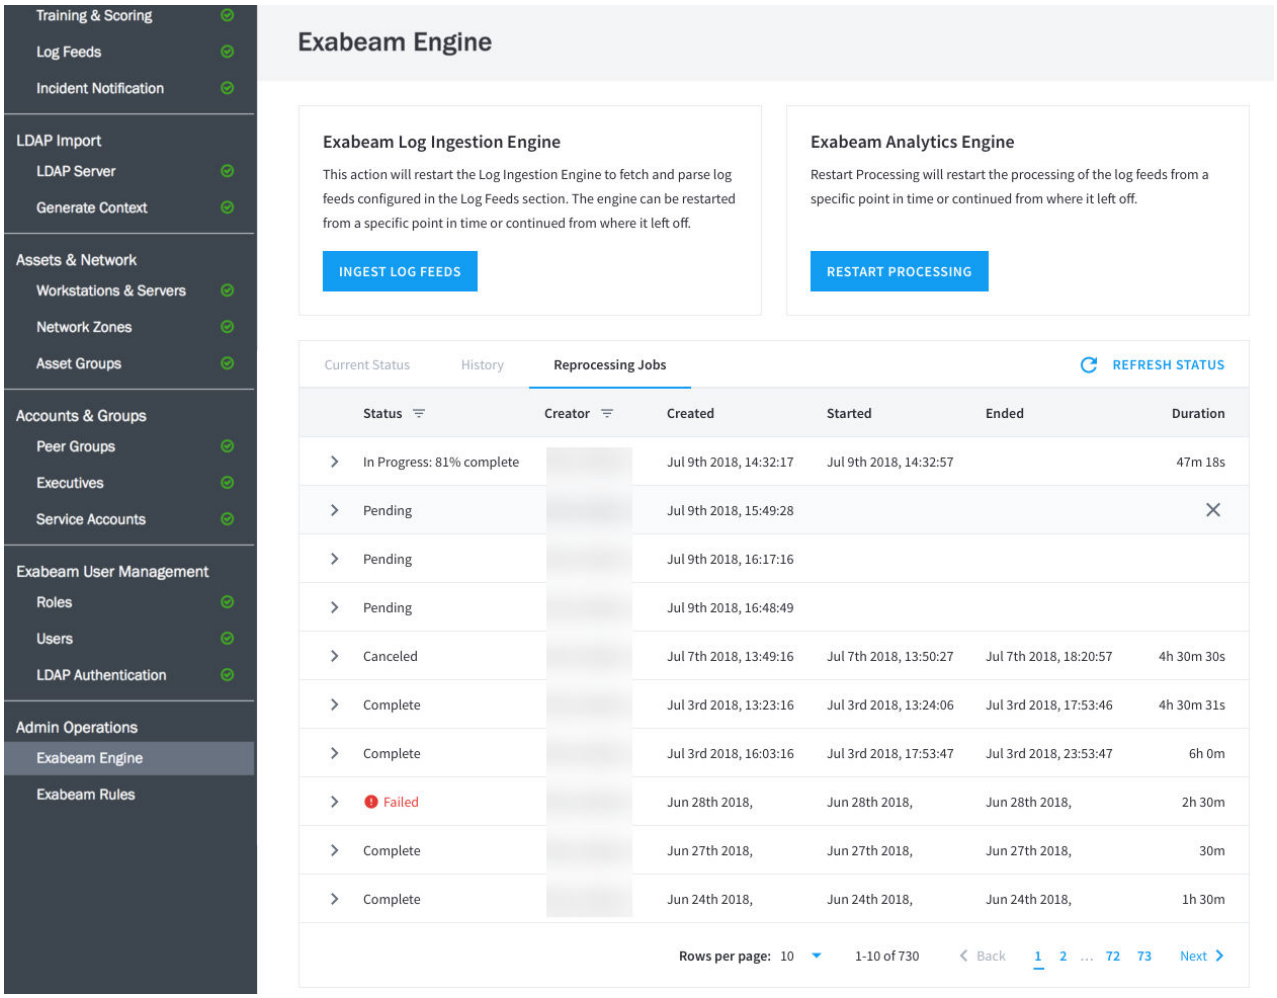

For more information on the Analytics Engine in Advanced Analytics, please refer to the *Advanced Analytics Admin Guide*.

1.5. Updated Exabeam Analytics Engine Reprocessing Table

We've improved the **Exabeam Analytics** page to provide better detail and control on reprocessing jobs. You can now view the status of jobs (for example, completed, in-progress, pending, and canceled), view specific changes and other details regarding a job, and cancel a pending or in-progress job.

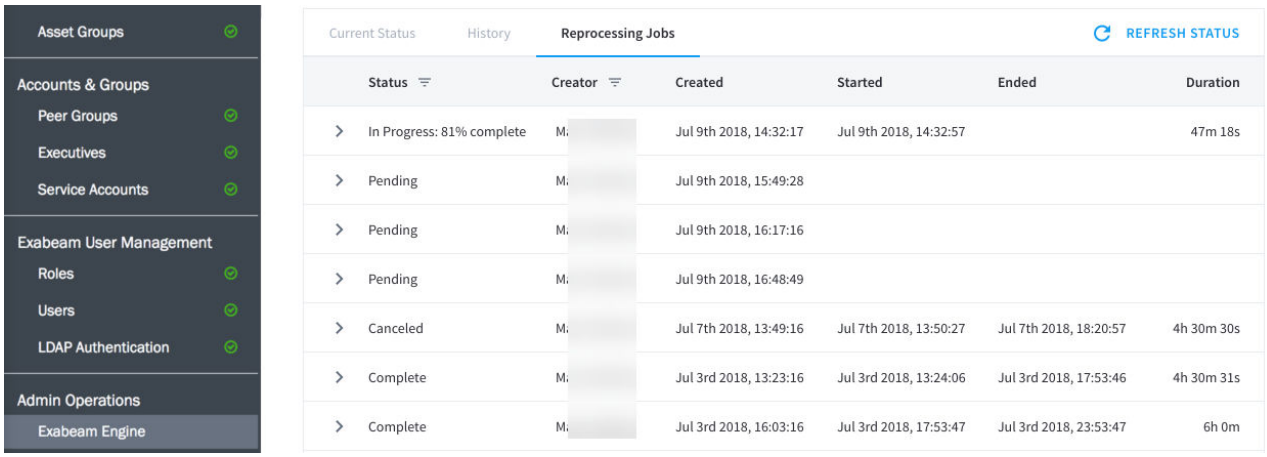

For more information on reprocessing with the Exabeam Analytics Engine, please refer to *Restart and Reprocess* section of the *Advanced Analytics Admin Guide*.

1.6. Model Sizing Improvements

Exabeam helps manage the sizes of your models so that they can no longer cause Advanced Analytics to run out of memory. This is done by setting a max number of bins for categorical models. The new limit is 10 million bins, although some models, such as ones where the feature is "Country" have a lower limit. We have put many guardrails in place to make sure models do not consume excessive memory and impact overall system health and performance. These include setting a maximum limit on bins, enabling aging for models, and verifying data which goes into models to make sure it is valid. If a model is still consuming excessive amounts of memory then we will proceed to disable that model.

Disabled models are displayed on the **System Optimization** tab of the **System Health** page. There are two disabled models tables — the first contains a list of models that have been disabled for all entities (Global Models), while the second contains a list of models that have been disabled for specific entities (Model Instances) within Advanced Analytics.

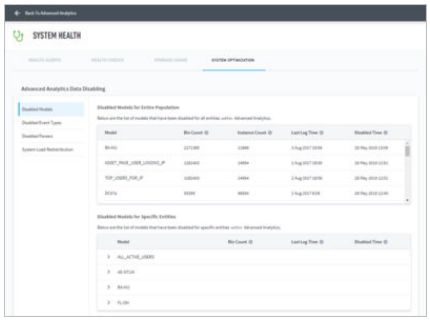

You are also shown an indicator on the User or Asset Page when a model has been disabled for that profile.

In addition, model aging is now configurable and enabled by default, with a window of 16 weeks. For some models that contain more sensitive or rare data, e.g. models that track executives or privileged users, the cycle is 32 weeks. Model aging considers data samples taken from a certain number of weeks instead of all points since the beginning of time. This process enhances system performance by cleaning out unused or underutilized models.

For more information on the Model Sizing Improvements, please refer to *Model Aging* in the *Configuring Advanced Analytics* section and *Disabled Models* in the *Health Status* section of the *Advanced Analytics Admin Guide*.

1.7. Top Users and Assets Improvements

As you continue to use Advanced Analytics, you may find that certain high volume users or assets within your organization amass a large number of events of certain event types. When this happens, Exabeam now takes preventive measures to protect the performance of Advanced Analytics by disabling these event types for a specific user or device.

Exabeam helps manage the event type volume by identifying and blacklisting top users and assets. When a specific activity type such as "web" for a certain user or device exceeds 10 million events and a specific event type within that activity type makes up 70% or more of the events in that total sum, then that event type for that user or device will be automatically disabled. If no single event type accounts for over 70% of the total event count in that activity type, then that entity is disabled. These thresholds are configurable.

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.

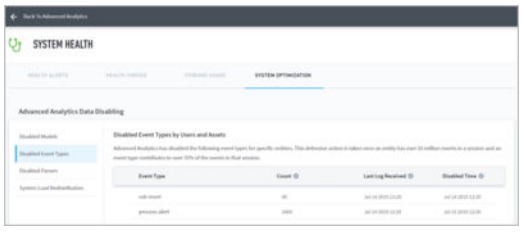

You are also shown an indicator on the **User** or **Asset** Page when an event type has been disabled for that profile. The affected User/Asset Risk Trend and Timeline would account for the disabled event type by displaying statistics only for the remaining events.

For more information on the Top Users and Assets Improvements, please refer to *System Optimization* in the *Health Status* section of the *Advanced Analytics Admin Guide*.

1.8. Parser Defensiveness

Advanced Analytics automatically identifies poor parser performance and disables such parsers in order to preserve the system health.

We determine the average parse time for each parser in a five minute period. We compare that to a configurable threshold variable in lime.conf. Then we divide each by the total time taken by all parsers in the same five minute period and compare the values to a configurable threshold variable in lime.conf. If the parsers average parse time exceeds the threshold and it exceeds the second threshold of being over a certain percentage of the overall parse time by all parsers then it becomes a candidate for disabling. We perform the same check during a second five minute period and if the same holds true then we proceed to disable the parser.

A slow parser is disabled if its average parsing time is above the parsing time threshold and makes up 50% or more of the total parsing time of all parsers.

Disabled parsers are displayed on the **System Optimization** tab of the **System Health** page. You can see a list of all parsers that have been disabled.

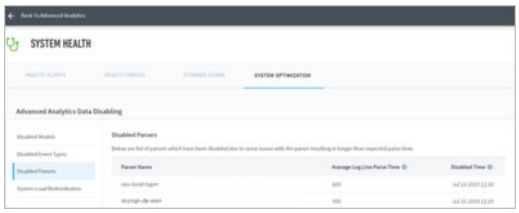

You are also shown an indicator when Advanced Analytics determines that a parser is problematic and disables it.

For more information on parser defensiveness, please refer to *System Optimization* in the Health Status section of the *Advanced Analytics Admin Guide*.

1.9. System Load Redistribution

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

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.

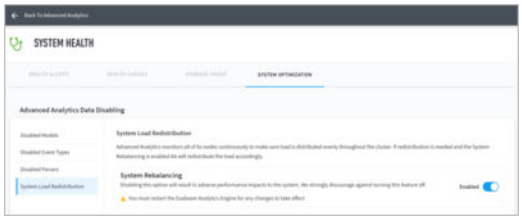

You are also shown an indicator in the UI when a redistribution of load is needed, is taking place, or has completed.

For more information on System Load Redistribution, please refer to *System Optimization* in the *Health Status* section of the *Advanced Analytics Admin Guide*.

1.10. Hadoop Distributed File System (HDFS) Namenode Storage Redundancy

A safeguard has been introduced in the HDFS NameNode (master node), storage to prevent data loss in the case of data corruption. Redundancy is automatically set up for you when you install or upgrade Advanced Analytics and include at least three nodes.

With this feature enabled in the case of the Master NameNode failing the system can still move forward without data loss. In such cases, you can use this redundancy to fix the state of Hadoop (such as installing a new SSD if there was an SSD failure) and successfully restart it.

For more information on HDFS redundancy in Advanced Analytics, please refer to the “HDFS Namenode Storage Redundancy” section within the *Advanced Analytics Admin Guide*.

1.11. Custom Configuration Validation

Any edits you make to your Exabeam custom configuration files are now validated before you are able to restart the analytics engine to apply them to your system. This will help to prevent Advanced Analytics system failures due to inadvertent errors introduced to the config files.

The system validates Human-Optimized Configuration Object Notation (HOCON) syntax, for example, missing a quotes or wrong caps ("SCOREMANAGER" instead of "ScoreManager"). The validation also checks for dependencies such as extended rules in custom config files that are missing dependencies within default config files.

If found, errors are listed by file name during the analytics engine restart attempt.

In addition to helping you troubleshoot your custom config edits, Advanced Analytics also saves the last known working config files. Every time the system successfully restarts, a backup is made and stored for you. Therefore, you can choose to rollback to the last backup if you run into configuration errors that you are unable to fix.

For more information on Custom Configuration Validation in Advanced Analytics, please refer to the “Restart the Analytics Engine” and “Custom Configuration Validation” sections within the*Advanced Analytics Admin Guide*.

1.12. Calico Implementation

The underlying native docker overlay network technology has been replaced with Calico in Data LakeAdvanced Analytics.

Compared to the native docker overlay technology, Calico brings the following benefits to all Exabeam deployments:

- **Simplicity** – Clusters on Docker overlay depend on a Linux kernel technology known as VXLAN which is subject to continuous updates. Calico removes this Linux dependency and uses layer 3 routing, just like non-container networks. Routes are shared using BGP, which is the de facto routing protocol of the Internet (used everywhere in autonomous systems like ISPs for example).
- **Stability** – Docker overlay makes use of tunneling over VXLAN to fool the nodes into thinking they are on the same network. Untunneling and translating headers using this approach adds significant latency and unreliability to overlay, especially at scale. Calico removes this by having the first three BGP peers advertise container routes to all other peers creating a redundant routing framework. When these other peers receive the route information, they will update their routing tables in real time for stability.
- **Performance** – Use of docker overlay requires encapsulating and decapsulating packets over UDP. This adds overhead to the network stack. It can be reduced by hardware acceleration but it cannot be fully removed. This overhead is completely avoided by Calico as it does not employ any encapsulation. Also, since Calico does not use tunneling, the network performance is equivalent to a native network stack.
- **Security** – Calico offers built-in support for network policies that can be controlled by simple metadata on the primary host. To implement the equivalent security in docker overlay (particularly for port restrictions from external services attempting to reach internal containers), complicated firewalled rules must be created, maintained, and terraformed.

Adopting Calico resolves various overlay and network flapping issues with cluster containers that resulted in intermittent network connection breakdowns and UI restarts in Advanced AnalyticsData Lake deployments.

You must meet the following requirements before installing or upgrading to this release:

- AWS deployments: All nodes MUST have src/dest (source/destination) checks turned off.
- GCP deployments: Network open to IP protocol 4 (IP in IP) traffic within the cluster.
- Nodes allow traffic to and from security group to itself.
- Use a load balancer in front of your cluster and use TCP (not UDP) as a transmission protocol between the load balancer and the Data LakeAdvanced Analytics hosts. A load balancer is required (customer-provided) in front of Data LakeAdvanced Analytics in order to have no downtime for Syslog ingestion during the upgrade.

If you have questions about the prerequisites, please create a support ticket at Exabeam Community to connect with a technical representative who can assist you.

1.13. Disaster Recovery Enhancements

Previously, Advanced Analytics disaster recovery only supported a two-cluster (active and standby) deployment scenario. Now, you can configure a multi-cluster disaster recovery deployment involving three or more clusters. In this scenario, you can configure a source (primary) cluster and one or more destination clusters.

The primary cluster is responsible for fetching the logs from SIEM or receiving the logs via Syslog. The secondary cluster(s) are responsible for replicating data from the primary cluster.

For more information on configuring and managing disaster recovery, please refer to the *Disaster Recovery* section of the *Advanced Analytics Admin Guide*.

1.14. Additional Settings Link

Previously, you could not conveniently access certain tabs of the Advanced Analytics **Admin Operations** page. Now, we've added an **Additional Settings** link to the Advanced Analytics settings page to quickly and conveniently access the **Admin Operations** page.

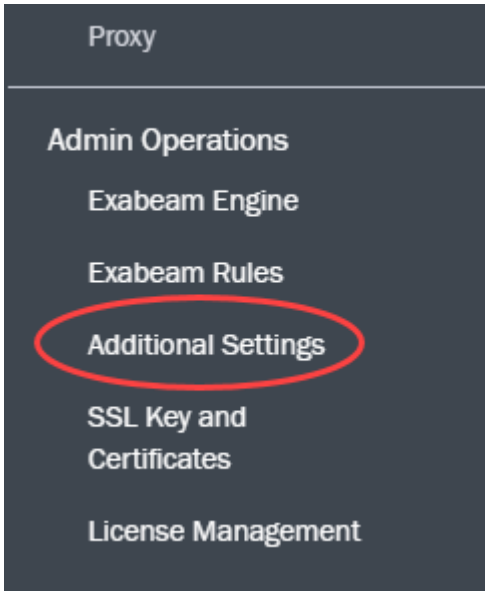

For more information on additional settings in Advanced Analytics, please refer to the *Advanced Analytics Admin Guide*.

1.15. Prevent Searches on Masked Fields

Previously, customers could search for and view non-masked fields, whether or not data masking was enabled for their deployment.

For example, users could successfully search for a non-masked username or account name in Threat Hunter. Additionally, users could copy the URL of a masked username (such as `https://<aa_IP>/uba/#user/<obfuscated_name>/...`) and simply change the URL to be a valid account name (such as, `https://<aa_IP>/uba/#user/<valid_account_name>/...`).

Now, users can no longer search for and view non-masked fields by any means if data masking is enabled for their deployment.

For more information on data masking in Advanced Analytics, please refer to the *Advanced Analytics Admin Guide*.

1.16. Reprocessing Job Notifications

You can now configure email and syslog notifications for certain reprocessing job status changes, including start, end, and failure.

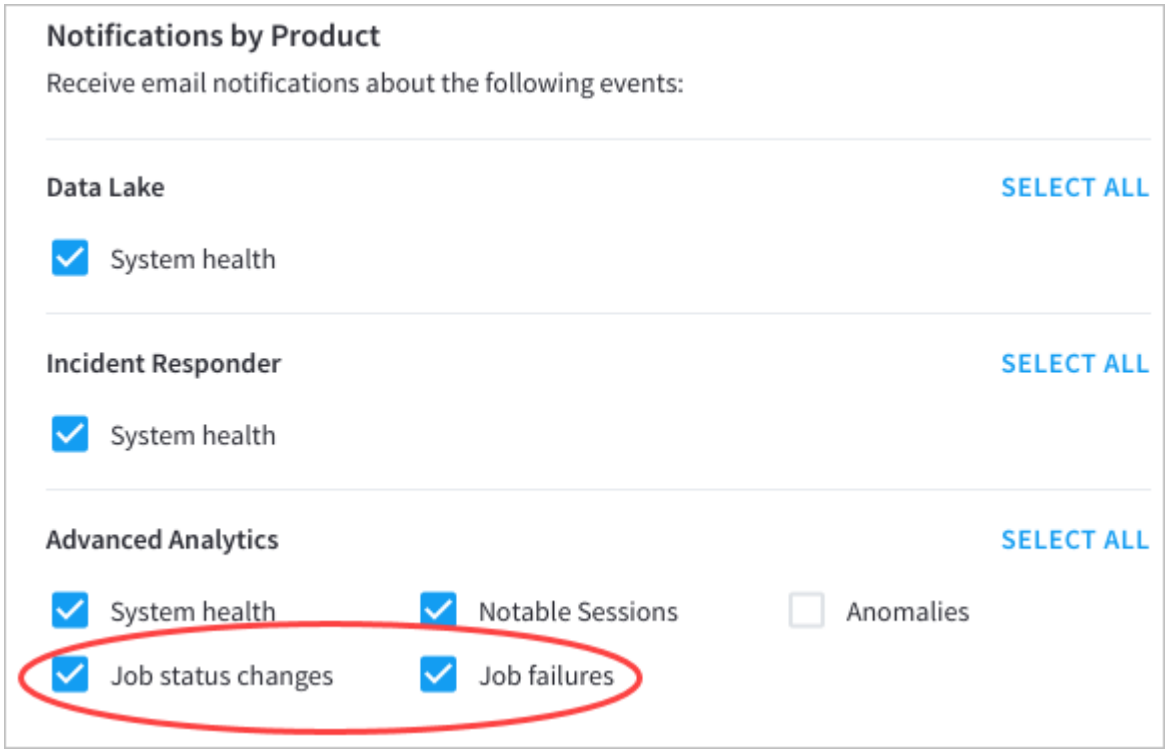

For more information on configuring job notifications, please refer to the *Restart and Reprocess* section of the *Advanced Analytics Admin Guide*.

1.17. EDS Memory Enhancements

We have streamlined and optimized memory management in EDS by avoiding duplicates for enhanced performance in context lookup. This improves performance when performing searches, generating reports, and performing other tasks within the Data LakeAdvanced Analytics UI.

1.18. Exabeam Threat Intelligence Service Enhancements

We've added a new settings page to provide better control over Exabeam Threat Intelligence Service feeds in your Data LakeAdvanced Analytics deployment. Additionally, you can now easily assign or unassign threat intelligence feeds to/from individual or multiple context tables or create new context tables directly from the page.

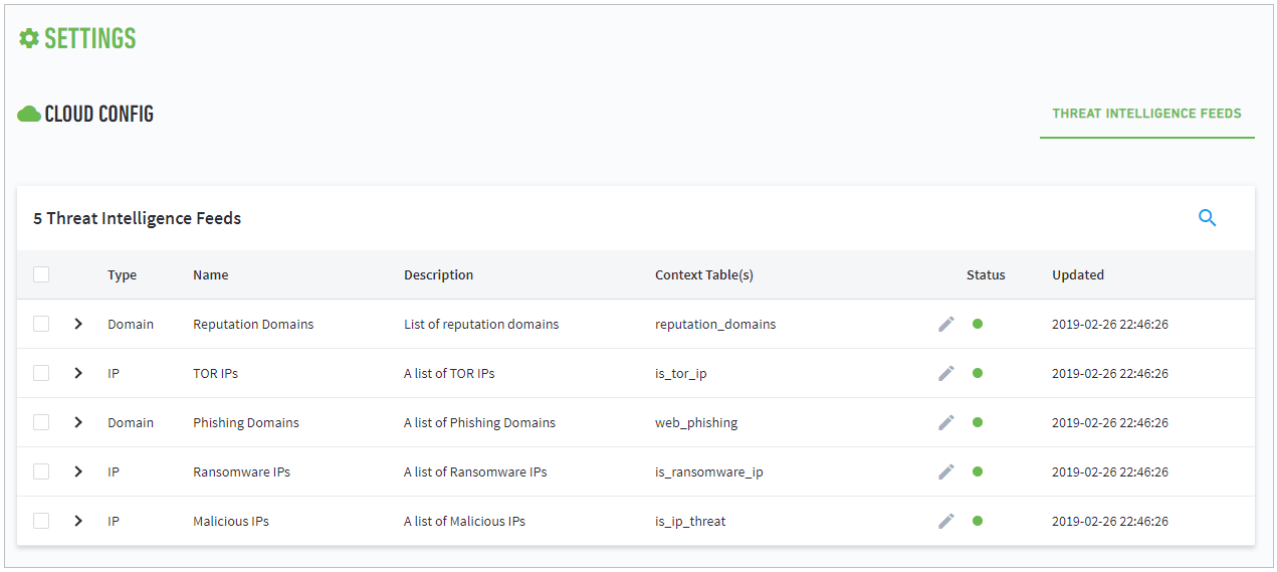

For more information on configuring and managing Exabeam Threat Intelligence Service, please refer to the *Exabeam Threat Intelligence Service Overview* section of the *Advanced AnalyticsData Lake Admin Guide*.

1.19. Exabeam Cloud Telemetry Service

Telemetry data such as events, metrics, and environment data is collected from Data Lake and Advanced Analytics deployments and sent to Exabeam Cloud Platform to provide visibility into the overall system health, reduce system health false positives, and enable Exabeam to gain insight into common system issues, such as processing downtime (for example, processing delays and storage issues) and UI/application downtime.

The Exabeam Telemetry Service is enabled by default, following the installation of this version.

> **NOTE**
> If you do not wish to send any data to the Exabeam Cloud, please follow the opt-out instructions listed in the *Disabling Telemetry Service* in the *Data LakeAdvanced Analytics Admin Guide* before installing this version. You can also choose to opt-out at any future time in the future if you choose to.

For more information, please refer to *Data LakeAdvanced Analytics Admin Guide* > *Exabeam Cloud Telemetry Service Overview*.

1.20. Entity Analytics UI Performance Optimization

We have optimized the following UI aspects of Entity Analytics so that they load significantly faster:

- Notable Assets and asset-based watchlists on the Homepage
- Security Alerts under the asset profile icon
- Entity Profile
- Entity Timeline

1.21. Numerical Clustering Improvements

We've improved the speed of numerical model calculation, with the exception of calculations involving the time of the week, which should improve the performance of Advanced Analytics. The new algorithm for numerical histogram anomaly detection not only performs much faster, but it also has better accuracy than the old algorithm.

The old algorithm was highly computing intensive, especially since it used clustering to find the abnormal area of the historical data distribution. The new algorithm uses gamma distribution to first estimate where the majority of the data points lie, and then calculates how far the new event is from the boundary of the normal area. The further the new event is from the boundary, the riskier the event is.

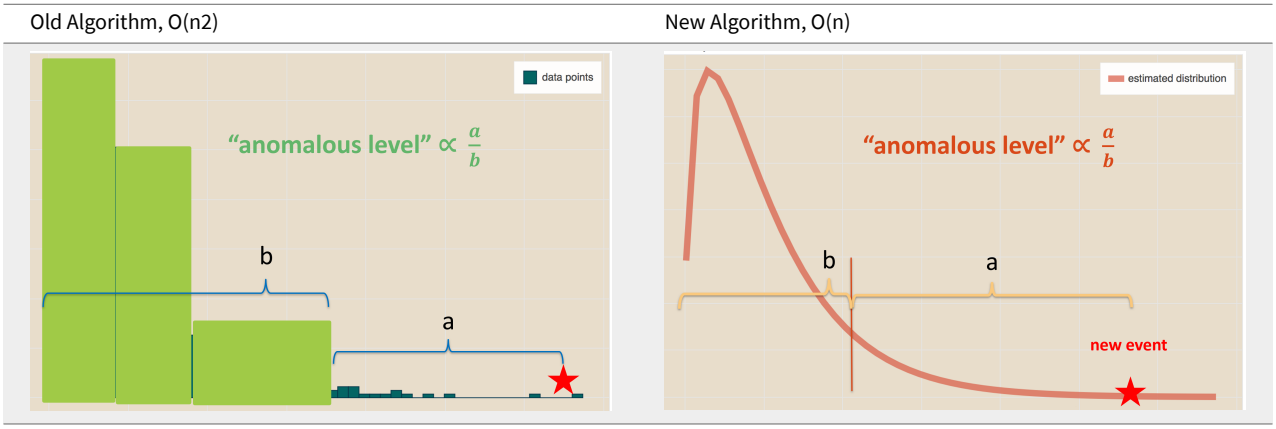

As previously mentioned, the old algorithm used hierarchical clustering, O(n2) to cluster every point into bins. It then used sorting, O(nlog(n)), to identify the normal bin boundary, **b** as seen in the above graph example. On the other hand, the new algorithm simply fits the historical points into gamma distribution, O(n), and then calculates the normal area boundary, **b**, using O(1). The resulting abnormal area is identified as **a** in both graph examples.

1.22. SAML Configuration Settings Link

Previously, the **Configure SAML** link was located under the **Admin Operations** tile of the settings page. Now, the **Configure SAML** link is located under the **User Management** tile of the settings page.

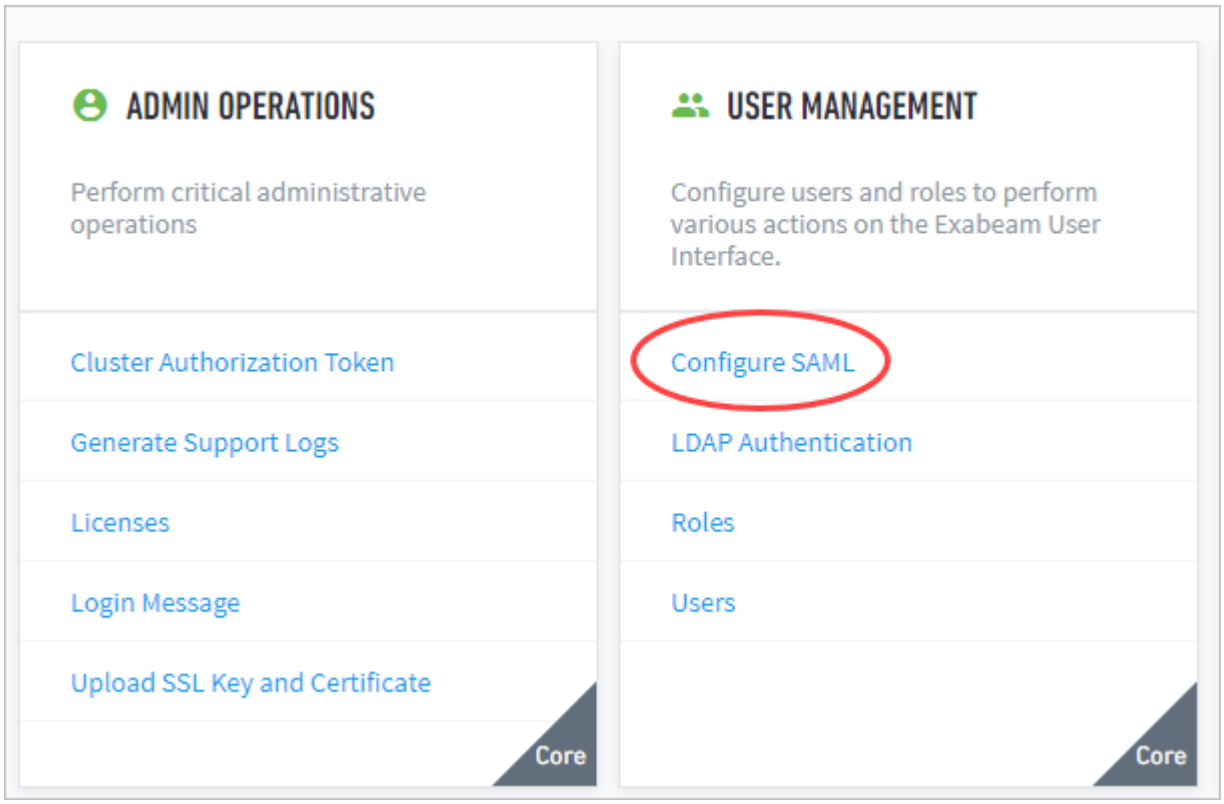

For more information on configuring SAML, please refer to *Advanced AnalyticsData Lake Admin Guide* > *Configuring SAML*.

1.23. Added Additional Option for Watchlists Timeframe Filter

We've added Last 2 days to the timeframe filter for Notable Users, Notable Assets, Account Lockouts, and other Watchlists on the homepage.

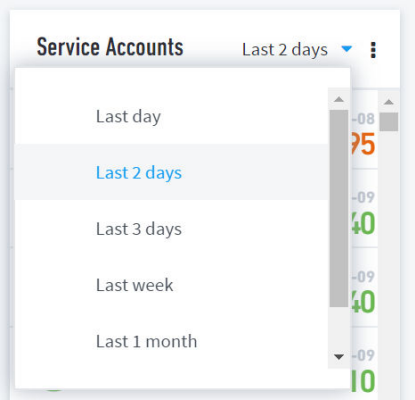

You can also configure the default timeframe filter value in the `application_default.conf` file.

For more information on the configurations for the watchlists timeframe filters, please review the*Advanced Analytics Admin Guide*.

For more information on the configurations for the watchlists timeframe filters, please refer to the *Watchlists Timeframe Filter* section of the *Advanced Analytics Admin Guide*.

1.24. MongoDB Retention and Usage Improvements

You can now monitor and configure your Advanced Analytics MongoDB data retention and usage along with your HDFS storage. This information is visible in new panels on the **Storage Usage** tab within **System Health**. The **MongoDB Data Retention** and **Usage** panels include storage usage of events in the database.

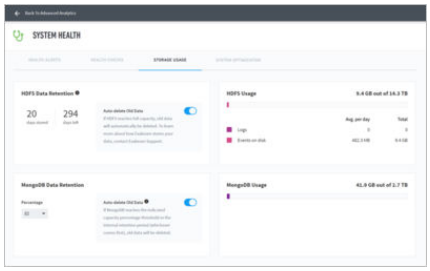

MongoDB data retention is enabled by default. You can disable/enable data retention and set the capacity used percentage threshold of MongoDB collections to help improve the performance and storage usage of Advanced Analytics. Exabeam does not recommend that you disable this feature.

In addition to the capacity used percentage, Advanced Analytics keeps six months of event data in MongoDB by default.

> **NOTE**
> Advanced Analytics maintains the previous default retention period for existing customers. Customized retention values are also retained.

For more information on the Mongo Retention Improvements, please refer to *Storage Usage and Retention* in the *System Health Page* section of the *Advanced Analytics Admin Guide* or the *User Guide*.

1.25. Retention Limits for Triggered Rules and Sessions Collections

You can now monitor and configure the data retention and usage of your triggered rules and sessions collections within Advanced Analytics MongoDB. This information is visible in an updated panel on the **Storage Usage** tab within **System Health**. The **MongoDB Data Retention** and **Usage** panels include storage usage of triggered rules and sessions and events in the database.

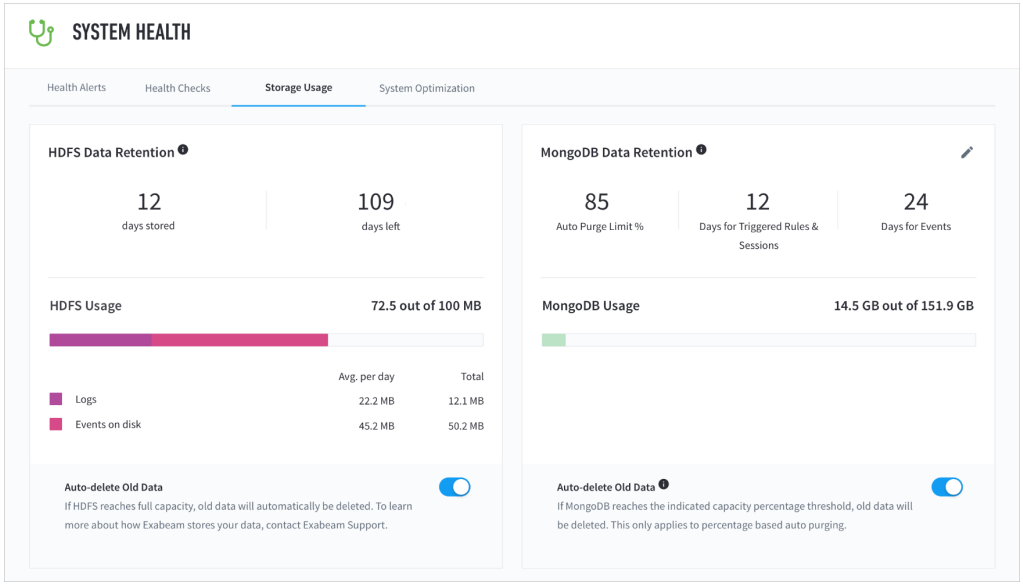

MongoDB data retention is enabled by default. You can disable/enable data retention and set the capacity used percentage threshold of MongoDB collections to help improve the performance and storage usage of Advanced Analytics. The default capacity is 85%. Exabeam does not recommend that you disable this feature.

In addition to the capacity used percentage, Advanced Analytics keeps 365 days of triggered rules and containers data and 180 days of event data in MongoDB by default. This is also configurable through the UI.

NOTE
Advanced Analytics maintains the previous default or customized event retention period for existing customers or 1095 days, whichever is smaller.

For more information on the Retention Limits for Triggered Rules and Sessions Collections, please refer to “View Storage Usage and Retention Settings” in the “Health Status Page” section of the *Advanced Analytics Admin Guide*.

1.26. Improved Martini and Lime Coordination

We have improved how our Ingestion and Processing Engines coordinate together. If Advanced Analytics detects the need to have the Ingestion Engine be restarted along with the Processing Engine restart then the system will prompt the user accordingly.

Performing a coordinated restart will help correctly apply all log feeds and ensure that both systems are running the same configurations. This bundled restart will also help prevent Analytics Engine and Log Ingestion Engine processing failures.

If a Log Ingestion Engine restart is required when you attempt to restart the Analytics Engine, you will be prompted with a dialog box to also restart the Log Ingestion Engine. The below dialog will appear prompting the user to start the Log Ingestion Engine if and only if Advanced Analytics detects a need to

have the Log Ingestion Engine restarted. You can choose to decline the restart if you would like the Log Ingestion Engine to finish its current process, but this will cancel the Analytics Engine restart procedure.

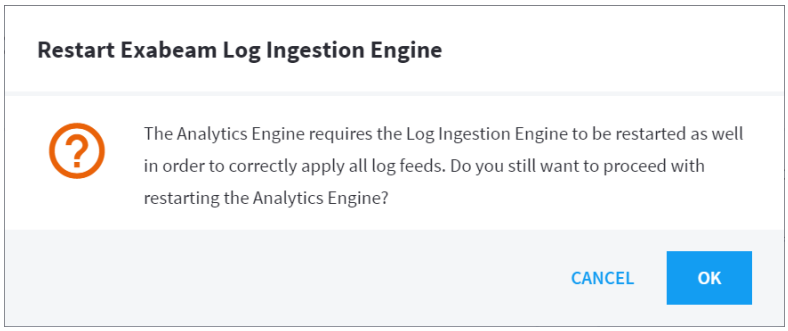

For more information on restarting the Exabeam Analytics Engine, please refer to “Restart and Reprocess” section of the *Advanced Analytics Admin Guide*.

1.27. Draft/Published Modes for Log Feeds

We have improved and added a new functionality to the log feed creation capability. Now when you create a new log feed and complete the workflow, you will be asked if you would like to publish the feed. Publishing the feed lets the Analytics Processing Engine know that the respected feed is ready for consumption.

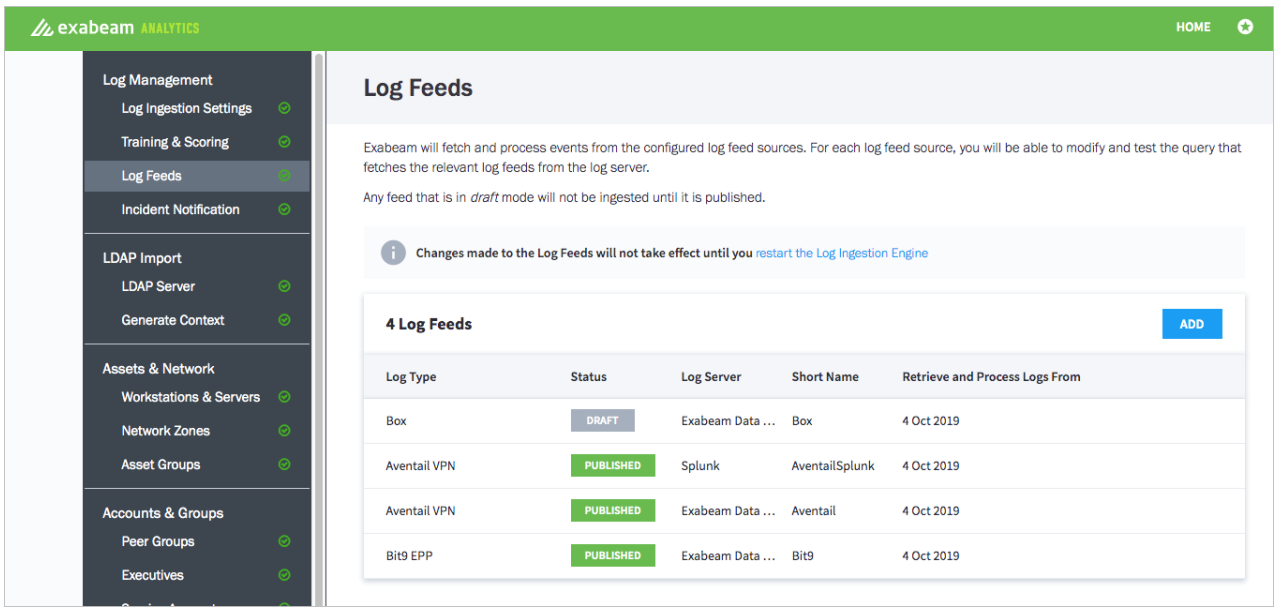

If you choose to not publish the feed then it will be left in draft mode and will not be picked up by the processing engine. You can always publish a feed that is in draft mode at a later time.

For more information on publishing draft log feeds, please refer to “Log Feeds” section of the *Advanced Analytics Admin Guide*.

1.28. 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
  - Tabbing into a text box

Keep in mind, Exabeam user engagement analytics does not:

- Record the screen
- Perform “screen scraping”
- Capture anything entered into any input field

For more information on the user engagement analytics policy—including instructions to opt out of user engagement analytics—please refer to the Advanced AnalyticsData Lake Administration Guide > Configuring Advanced AnalyticsData Lake > User Engagement Analytics Policy.

1.28.1. OPT OUT OF USER ENGAGEMENT ANALYTICS

> **NOTE**
> For customers with a Federal license, we disable user analytics by default.

To prevent Exabeam SMP from sending your data to our user analytics:

1. Access the config file at

   ```
   /opt/exabeam/config/common/web/custom/application.conf
   ```

2. Add the following code snippet to the file:

```
webcommon {
    app.tracker {
      appTrackerEnabled = false
      apiKey = ""
    }
}
```

3. Run the following command to restart Web Common and apply the changes:

```
. /opt/exabeam/bin/shell-environment.bash web-common-restart
```

1.29. Host/IP and Host/Account Mapping Improvements

Advanced Analytics has improved host-to-IP and host-to-account mappings for mapping configuration changes in parallel. This optimization results in lower CPU usage while tying hosts to specific addresses or accounts.

1.30. Syslog Output Improvements

The incidents and alerts created by Advanced Analytics and sent out via syslog previously lacked parsed and enriched event details. These syslog notifications have now been enhanced to include additional event fields and reason rule templates.

The notifications include useful event fields that have values associated with them. This is unique to each event type. If the field does not have a value then it is not included in the syslog output.

Example syslog output after enhancement:

```
2019-12-12T20.Internal.syslog.log:<86> 2019-12-12T20:18:07.973569+00:00
2019-12-12T20:18:07.972Z exabeam-analytics-master Exabeam
timestamp="2019-12-12T20:18:04.396Z" id="user745-20140528185224" score="10"
user="user745" src_ip="46.117.121.157" event_time="2014-05-28 18:52:24"
event_type="vpn-login" host="10.2.0.111" is_session_first="true"
rawlog_time="1401313944000" time="1401303144000" source="VPN" vendor="Juniper
VPN" lockout_id="NA" session_id="user745-20140528185224" getvalue('isp',
src_ip)="013 NetVision Ltd" getvalue('country_code', src_ip)="IL"
session_order="1" account="user745" src_network_type="WAN" event_code="Agent
login succeeded" rule_id="AE-UA-F-VPN" rule_name="First VPN connection for
user" rule_description="First VPN connection for this user" rule_reason="First
VPN connection for user745"<86> 2019-10-11T23:05:04.002809+00:00
2019-10-11T23:05:03.998Z exabeam-analytics-master Exabeam
timestamp="2019-10-11T23:04:30.556Z" domain="Unspecified" host="host36"
score="5" rule_description="This rule is used to identify a user without a
significant event history. A user without a history introduces some risk as
they have not established a baseline for comparing activities" rule_id="NEW-
USER-F" dest_ip="10.2.16.17" dest_host="host37" id="user38-20140504070000"
event_time="2014-05-04 07:00:00" event_type="ntlm-logon" rule_name="User with
no event history" user="user38"
```

For more information on syslog output improvements in Advanced Analytics, please refer to the “Set Up Incident Notification” section within the *Advanced Analytics Admin Guide.*

2. Security Patch

2.1. ZombieLoad Vulnerability Patch

This Data LakeAdvanced Analytics build includes the patch to address the recently discovered ZombieLoad vulnerability affecting Intel CPUs.

3. Content Updates

The following is a high-level list of content changes available with Advanced Analytics v2019.1 (i48).

> **NOTE**
> For a detailed list of specific content changes, please submit a request on the Exabeam Community.

Description	Area
Ensure security vendors are referred to in the content according to Gartner naming.	Feature Support
Added new Entity Analytics rules to use with Exabeam Threat Intelligence Service lists.	Feature Support
Added user-based rules that use Exabeam Threat Intelligence Service lists.	Feature Support
Added content targeting the DCShadow attack to address the Mitre Att@ck technique.	New Detection Logic
Added content targeting the DCSync attack to address the Mitre Att@ck technique.	New Detection Logic
Added content targeting the Pass the Hash attack to address the Mitre Att@ck technique.	New Detection Logic
Introduced a new event type: process-created-failed.	New Detection Logic
Added content around file shares (Windows events 5140/5145) and web repositories (Box, Dropbox, etc.) to address the Mitre Att@ck technique.	New Detection Logic
Added content identifying brute force access attempts to address the Mitre Att@ck technique.	New Detection Logic
Added content targeting admin shares in Windows 5140/5145 events to address the Mitre Att@ck technique.	New Detection Logic
Added content targeting pass the ticket attack to address the Mitre Att@ck technique.	New Detection Logic
Added Entity Analytics rules to score assets involved in Pass the Hash attacks as well as add score when a security alert is triggered for a critical asset.	New Detection Logic
Added rule to identify pst/ost file when copied outside of the Outlook directory path.	New Detection Logic
Added support of Windows 5156/7 events as network-connection-successful/network-connection-failed events.	New Detection Logic
Added support for Data Lake correlation rule event as a security alert in Advanced Analytics.	New Detection Logic
Added content targeting kerberoasting attack to address the Mitre Att@ck technique.	New Detection Logic
Added rules to identify exfiltration to personal email domains using web activity events.	New Detection Logic
Added Entity Analytics rules and models for endpoint events.	New Detection Logic
Added Entity Analytics content around IDS (network-alert) events.	New Detection Logic
Added Entity Analytics Firewall and DNS content utilizing rule dependency and other Entity Analytics features.	New Detection Logic
Added parsers, rules, and models to address netflow events.	Improved Detection Logic
Updated the top 1 million domain list based on Cisco Umbrella. This also includes changes to how the content interacts with the list.	Improved Detection Logic
Added informational models showing users per alert name and assets per alert name.	Improved Detection Logic
Added content for web activities in which TOR paths appear.	Improved Detection Logic
Added additional scores for new users in the environment.	Improved Detection Logic
Added an additional score for users browsing to an IP address instead of a domain name.	Improved Detection Logic
Introduced the local_asset field and triggering once per it instead of once per src_host.	Improved Detection Logic

4. Known Issues

EXA-28840	When the system is under heavy Syslog load, ingestion may start to lag if it cannot parse all the incoming logs in time. To solve this issue, when the lag reaches a certain point (a maximum of one hour), the system will initiate another task to handle the extra Syslog traffic. You may notice a jump in terms of Lime throughput. Conversely, you may also notice a dip in the throughput after the extra task finishes. We are also working on additional optimizations that will reduce the syslog import lag in the next releases. If your `/opt/exabeam/data/input/rsyslog` has more than three hours of rsyslog files backed up, please contact Customer Success by opening a case in the Exabeam Community.

INF-4585	Ingestion does not start after a restart.
	
	This issue is caused by some Hadoop data nodes starting up before the Hadoop master node.
	
	To solve this issue, identify the Hadoop bug:
	
	1. To confirm that the hadoop-master and all hadoop-data nodes are running, run `docker ps` across the cluster:
	
	`ansible all -i /opt/exabeam_installer/inventory -ba "docker ps" | grep hadoop | sort -Vk7`
	
	2. Find which Hadoop data nodes are not reporting to the master:
	
	`docker exec -it hadoop-master hdfs dfsadmin -report | grep Hostname | sort -k2V`
	
	Take note of any nodes that don't report. In this example, data noes a, b, and c are not reporting:
	
	`Hostname: hadoop-data-host1-d`
	`Hostname: hadoop-data-host1-e`
	`Hostname: hadoop-data-host1-f`
	`Hostname: hadoop-data-host2-d`
	`Hostname: hadoop-data-host2-e`
	`Hostname: hadoop-data-host2-f`
	`Hostname: hadoop-data-host3-d`
	`Hostname: hadoop-data-host3-e`
	`Hostname: hadoop-data-host3-f`
	
	3. To fix the data nodes that don't report, run a `sudo systemctl restart hadoop-data-N` command on any data node that doesn't show up:
	
	`for _node in a b c; do systemctl restart hadoop-data-"${_node}"; done`
	
	`for _host in 2 3; do ssh "${_host}"; done`
	
	4. After restarting Hadoop, make sure that an `hdfs -report` comes back clean, and Log Ingestion and Messaging Engine (LIME) and the Analytics Engine are running. If not, restart them.

PLT-9180	Unable to generate verbose logs while generating a support package. Generating verbose logs will cause an error and you will be unable to complete generating the support package. To solve this issue, unselect the **verbose** option when generating support logs in the admin operations settings.

5. Fixed Issues

Advanced Analytics I48.6

The table below is a list of customer-reported issues that have been resolved in Advanced Analytics SMP 2019.1 (AA I48.6).

EXA-29870	When an Advanced Analytics session became notable and created an incident in Incident Responder, the system executed playbooks before the entire incident information had been populated. Therefore, the playbook operated on an incomplete subset of incident fields/entities/artifacts.

Advanced Analytics I48.5

The table below is a list of customer-reported issues that have been resolved in Advanced Analytics SMP 2019.1 (AA I48.5).

EXA-28668	Rule creation dependency expression search did not show all custom rules.

Advanced Analytics I48.4

The table below is a list of customer-reported issues that have been resolved in Advanced Analytics SMP 2019.1 (AA I48.4).

EXA-28498	Standalone parser with especially slow parsing would not provide .msg.gz files.
EXA-28438	Risk score in the UI did not match actual risk score totals.
EXA-28198	Duplicate models were created and interfered with specific scoring after restarting the Advanced Analytics engine.
PLT-7900	Exabeam cloud connection service was not healthy.
INF-4038	Older Docker version caused all services on the master node and the system to stop.
INF-3917	Upgrading from Advanced Analytics I38 to SMP 2019.1 (AA I48) resulted in a pre-check warning regarding the lack of CPU cores.
INF-3898	After upgrading to Advanced Analytics SMP 2019.1 (AA I48.2) logs showed HDFS errors.

5.1. Issues Fixed in Advanced Analytics I48.3_44

EXA-28118	Fixed issue where Analytics Engine always processed the previous hour events.
EXA-27902	Fixed issue where the UI did not load watchlists or user pages due to unresponsive worker node.
EXA-27864	Fixed issue where Windows logs did not properly show up in the UI after upgrading from i41 to i46.2.
EXA-25630, EXA-27726	Fixed issue where the “!WasRuleFired” expression caused rules to trigger multiple times in the same session.
EXA-27676	Fixed a timing issue that caused the system to miss loading entries affected by Bayesian scoring.
EXA-27280	Fixed issue where Threat Hunter search query of a username would timeout.
EXA-27626	Fixed issue where Threat Hunter searches with HostBehaviorClassifier Asset Labels resulted in an error.
EXA-28471	Fixed issue where old notable user email alerts were being sent.
EXA-28494	Fixed issue where unusually large differences in numerical clustering values caused Advanced Analytics Restful Web Services to run out of memory.
EXA-28481	Fixed issue where rules fired on models that did not yet meet the cutoff value.
EXA-28473	Fixed issue where some numerical histogram values displayed unknown as the expected value.

PLT-7655	Fixed issue where it was not possible to apply static Asset Labels via the Context Management UI.
PLT-7597	Fixed issue where users assigned to multiple Okta groups were assigned only permissions to one group.

5.2. Issues Fixed in Advanced Analytics I48.3_47

EXA-28557	Fixed issue where log entries did not parse due to the log.gz.ongoing file being modified while LogParser was still parsing.
EXA-28690	Fixed issue where upgrade to I48.3 failed due to extended index creation time in container_db.

5.3. Issues Fixed in Advanced Analytics I48.4_50

EXA-28724	Fixed issue where Network Zones could not be added to Assets & Network.
EXA-28498	Fixed issue where standalone parser with especially slow parsing would not provide .msg.gz files.
EXA-28438	Fixed issue where the risk score in the UI did not match actual risk score totals.
EXA-28198	Fixed issue where duplicate models were created and interfered with specific scoring after restarting the Advanced Analytics engine.
EXA-27984	Fixed issue where the event navigation for database events failed.
EXA-27821	Fixed issue where Entity Analytics email notifications did not show entities properly.
PLT-8669	Fixed issue where the Cloud Connectivity Service would be interrupted.
PLT-7900	Fixed issue where the ExaCloud connection service was not healthy.

5.4. Issues Fixed in Advanced Analytics I48.4_56

EXA-25630	Fixed issue where rule including !WasRuleFired was triggered twice in the same session.
EXA-28987	Fixed issue where sequence rules did not add to risk score.

5.5. Issues Fixed in Advanced Analytics I48.4_58

EXA-28996	Upgraded Advanced Analytics to the latest JDK to address the JVM bug which caused Entity Analytics worker nodes to get stuck.

5.6. Issues Fixed in Advanced Analytics I48.4_62

EXA-29225	Enhanced model memory guards with new scenarios.

5.7. Issues Fixed in Advanced Analytics I48.5_48

EXA-29228	Fixed issue where the Log Ingestion Engine did not close ongoing files.
DS-90	Fixed issue where the personal email detection feature on the DS-server did not reload updated collection tables.

5.8. Issues Fixed in Advanced Analytics I48.5_60

EXA-29793	Fixed issue where the upgrade to i48.5 would get stuck on the JDK.
EXA-29462	Fixed issue where the worker nodes would get stuck on rule aggregation.

5.9. Issues Fixed in Advanced Analytics I48.6_59

EXA-3008	Asset data may not display after upgrading to Advanced Analytics I48.6_51.
EXA-30093	DLP alerts and VPN logout events dropped before they could start a new user session.

5.10. Issues Fixed in Advanced Analytics i48.6_74

EXA-30249	Failing over to the disaster recovery cluster didn't bring up the UI and other services.
EXA-30141	The UI allowed you to reload a rule, even if it had syntax errors in its definition.
PLT-9362	Service Accounts settings wouldn't save.
PLT-9913	Generate Context settings didn't provide correct information because context tables weren't updated.

5.11. Issues Fixed in Advanced Analytics i48.6_75

EXA-30657	The Log Ingestion and Message Extraction (LIME) engine stopped parsing Google Pub/Sub data.
INF-6367	You couldn't upgrade Advanced Analytics because the /home partition didn't have enough space.

5.12. Issues Fixed in Advanced Analytics i48.6.4

PLT-10310	Fix for internal issue.

5.13. Issues Fixed in Advanced Analytics i48.6.5

EXA-31029	Rules didn't trigger appropriately because models treated expected values as abnormal.