

EXABEAM DATALAKE INTEGRATION WITH SECURITY VALIDATION

This integration provides the following benefits:

- Validates that security tools are writing log events to Exabeam Datalake to ensure compliance with security policies and regulations
- Collects events generated by security tools that write to Exabeam Datalake to test the efficacy and configuration of security controls using Security Validation jobs

Use this document to configure the integration using one of the following methods:

- MSI (Supported and recommended for new integration configurations)
- Legacy (Supported for existing integration configurations)

Configure MSI Integration



This method is the recommended approach for configuring new integrations in Security Validation.

To configure an integration, this document walks you through the following high-level steps:

1. Configure the third-party technology
2. Configure Security Validation
3. Verify connectivity

Configure Elasticsearch

API Calls

API	Usage
<code>/_search</code>	Query Elasticsearch for events
<code>/_search/scroll</code>	Collect additional pages of events returned by a query

Supported Versions

Elasticsearch 7.2

To configure this integration, you need:

- The hostname of your Elasticsearch instance
- A valid username and password for a user with permissions to use the API endpoints described in the previous step

Configure Security Validation

1. Go to **Settings > Integrations**.
2. From the Integrations table, click **Add Integration > Elasticsearch**.




You can add this as either a Direct or Remote Integration.

3. Enter a meaningful **Integration Name**.
4. Optional: From the **Proxy** drop-down, choose a proxy profile if one is available. If one isn't available and all outbound connections go through a proxy, first, set up a **Proxy Rule** (<https://docs.mandiant.com/home/msv-proxy-rules>).
5. Optional: Change the **HttpProtocols** value to determine what protocol is used for requests (**Http** or **Https**).
6. Enter a **Host** value (hostname or IP address) for the Elasticsearch instance.

7. Enter a **Port** value. The default is **443**.
8. Enter the **Username** and **Password** that you configured with permissions to use the API endpoints.
9. Optional: Check **Verify Ssl** if you want this verification done for requests to an upstream server.
10. Optional: Change the **Timeout** value if you want a different frequency of requests to an upstream server. The default is **30** (seconds).
11. Optional: For **Target**, enter a comma separated list of data streams, indices, and index aliases to search.


 Wildcard (*) expressions are supported. Examples: "index_one", "index_one,index_two", "index_*", "cluster_one:index_one", "cluster_*:*"


12. Optional: Add or remove **Queries**, as needed. Defaults are provided: **IP Query**, **Hostname Query**, **DNS Query**, and **Email Query**.
13. Optional: Change the default **Alert Query** and add an **Alert Target**, if needed.
14. Optional: Modify the **Field Map** values, as necessary.




- o Each field map box can hold a JSON-formatted comma-separated list of columns returned by the API to be considered for each field when translating into the normalized event object format. Example: description could be configured to be 'msg_s' or 'SyslogMessage' in some environments. The field map tries both if set to: ['msg_s','SyslogMessage'] and whichever matches first is the column that is used.
- o When configuring an integration in Security Validation, you can assign additional host values in the Field Map settings. If none of the assigned fields return a valid host name, Network Actions may miss matched events from the third-party technology. Additional hosts values helps ensure the likelihood of a match between the two environments.

15. Optional: Change the values for **Page Size** (default **500**) or **Max Pages** (default **100**), if needed.
16. Optional: Expand **Advanced options** and update the information as necessary.
 - a. Update **Query Time** and **Delay Time**.

 The **Query time** is the amount of time (minutes) before and after the query runs that the platform looks for events, while the **Delay time** is the amount of time (minutes) that the platform waits to run the first query after a Job Action starts. For example, you configure your integration with the following values: **Query time** = 5, **Query interval** = 30 seconds, and **Delay time** = 0. When a Job Actions starts at 12:00:00, the first time the query runs, the platform looks for events from 11:55:00 to 12:00:00. Then 30 seconds later, it looks for events from 11:55:30 to 12:00:30. This interval continues, with the last query looking from 12:00:00 to 12:05:00. If you instead configured the **Delay time** to equal 10, it would run the same query, but it wouldn't start that query until 12:10:00.

 If your monitors are set to run more frequently than the query time, this configuration impacts the pass/fail results for AEDA monitors.

- b. Update **Query Interval** (seconds).
- c. Select **Correlation Query Enabled** and fill in the **Correlation Query**.
- d. Modify the **Correlation Query Interval**, if necessary (minutes).
- e. Select **Discover network devices automatically**, the default and recommended option.


 If unselected, reported events won't include product information for any matching network security technology.

- f. Select **Save Suspicious Events**.

- g. Modify the **Event Time Adjustment** (seconds). The default is **0**.
- h. Modify the **Limit** value if you need to prevent a flood of results. This value is set to **10000** by default. This limit applies to both events and alerts individually, so if you set it to **10**, you can still see a maximum of 10 events and 10 alerts.

17. Click **Save**.

Verify connectivity

1. Go to **Settings > Integrations**.
2. From the Direct Integrations table, click  **> Test** to verify that:
 - The Director can communicate with the integration host on the port and protocol specified.
 - The integration credentials are valid and working.

For more information on setting up queries, see [Manage Integrations \(https://docs.mandiant.com/home/msv-managing-integrations\)](https://docs.mandiant.com/home/msv-managing-integrations).

Configure Legacy Integration

This document applies to Classic/Legacy Integrations. You may continue to use these integration configurations. While no active development is happening for these integrations, we continue to provide Classic/Legacy Integrations in the product. You do not have to move to MSI Integrations. If your support engineer or TSC recommends or you choose to move to MSI Integrations, you can take advantage of the latest features and functionality. For more information, see the MSI Integration documentation in the [Integrations Overview \(https://docs.mandiant.com/home/msv-integrations-overview\)](https://docs.mandiant.com/home/msv-integrations-overview).

This integration uses the Elasticsearch Scroll API to support querying large data sets.



This integration is remote capable.

Update Elasticsearch

Identify or create the credentials to access Elasticsearch, if applicable.

- Elasticsearch does not provide authentication by default.
- Authentication can be added with Elastic X-Pack, a third-party plug-in, or by using a reverse proxy like nginx.

Update the Validation Platform

Prerequisites

Information to gather before you start:

1. Identify the IP address used to access Elasticsearch. This could be direct access to an Elasticsearch node, Primary node, or something such as an nginx reverse proxy.
2. Identify the port for Elasticsearch communication (default is 9200).
3. Identify whether the protocol is HTTP or HTTPS for connections to the Elasticsearch port.
4. Identify the field name mappings for the following (there could be multiple of each, depending on log sources and configuration):
 - Source IP
 - Destination IP
 - Source Port

- Destination Port
- Event Start Time (timestamp)
- URL/Domain
- Email Sender
- Email Recipient
- Email Subject
- User
- Event Unique ID
- Event Signature ID
- Event Description
- Event Source Host

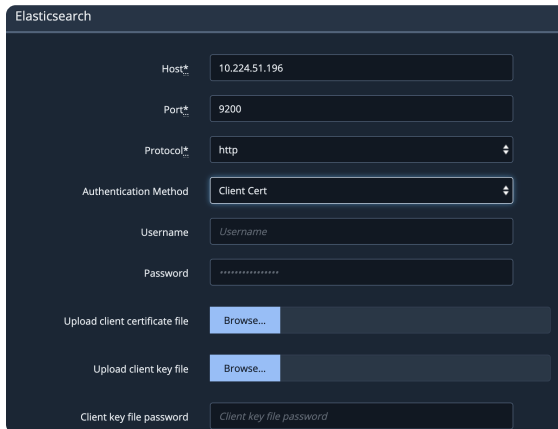


Organizations can create a query index for the integration. With this configuration, searches will query all indexes, which is less efficient than running against a specified index.

Configuration

TO ADD THE ELASTICSEARCH INTEGRATION

1. Go to **Settings > Integrations**.
2. Click **Add Integration > Elasticsearch**.



(<https://dyzz9obi78pm5.cloudfront.net/app/image/id/629a12d5c9cba0017c2f7dc3/n/elasticsearch.png>)

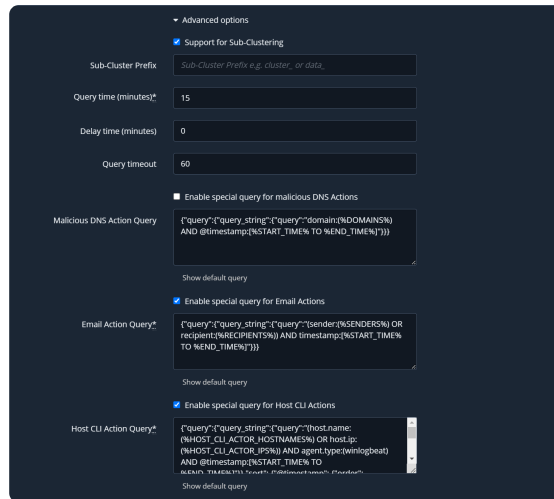
Elasticsearch Integration

3. Enter information for the **Host, Port, and Protocol**.
4. Select your Authentication method.
 - **None**: Skip the Username, Password, and Client fields.
 - **Basic**: Enter a Username and password.
 - **Client Cert**: Upload the client certification file, the client key file, and add the client key file password.



The remote Elasticsearch integration doesn't support client-cert authentication.

5. Update the Query, if needed.
6. Expand **Advanced options**.



Elasticsearch Integration (Advanced Options)

- (Optional) Select **Support for sub-clustering**. When you select this option, a **Sub-Cluster Prefix** input field displays below the option, which allows you to enter a custom sub-cluster prefix (i.e., `cluster_` or `data_`, etc.).



If a Director had any existing ElasticSearch integrations with the **Support for sub-clustering** box checked prior to an upgrade, after the upgrade that field will be set with `cluster_` by default. However, you can edit the field to be any custom sub-cluster prefix you want.

- (Optional) Update **Query time** and **Delay time**.



The **Query time** is the amount of time (minutes) before and after the query runs that the looks for events, while the **Delay time** is the amount of time (minutes) that the platform waits to run the query after a Job Action. For example, you configure your integration with the following values: **Query time** = 5, **Query interval** = 30 seconds, and **Delay time** = 0. When a Job Actions starts at 12:00:00, the first time the query runs, the platform looks for events from 11:55:00 to 12:00:00. Then 30 seconds later, it looks for events from 11:55:30 to 12:00:30. This interval continues, with the last query looking from 12:00:00 to 12:05:00. If you instead configured the **Delay time** to equal 10, it would run the same query, but it wouldn't start that query until 12:10:00.



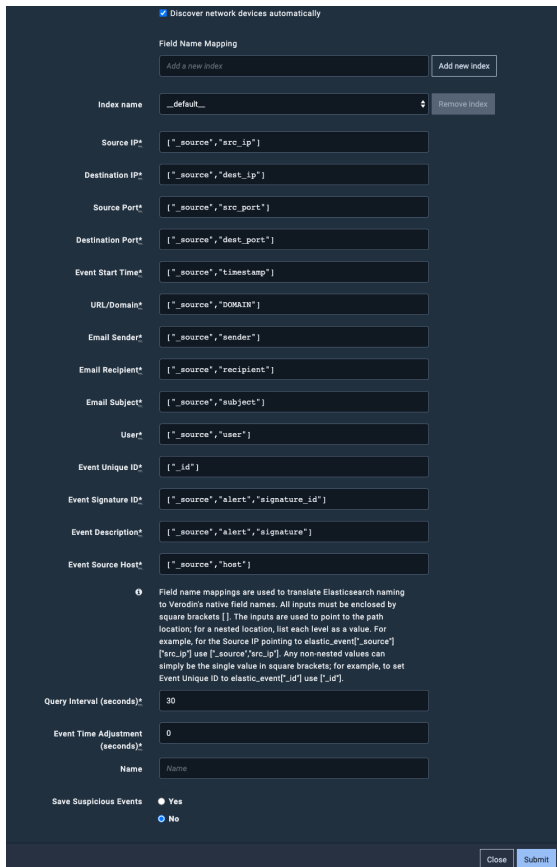
If your monitors are set to run more frequently than the query time, it will impact the pass/fail results for AEDA monitors.

- (Optional) Select **Enable query for Malicious DNS Actions** and configure the **Query**. This query will only be used when you run Malicious DNS Actions or Captive DNS Actions.
- (Optional) Select **Enable query for Email Actions** and configure the **Query**. This query will only be used when you run Email Actions.
- (Optional) Select **Enable query for Host CLI Actions** and configure the **Query**. This query will only be used when you run Host CLI Actions.



If you enable the Host CLI Actions query and use the `%HOST_CLI_ACTOR_HOSTNAMES%` variable, the platform will substitute the plain hostname and the information from the Alternate Hostname field on the Actor configuration page.

12. (Optional) Select **Discover network devices automatically**.
13. Review the field name mappings for the `__default__` index; update as necessary.
 - a. You can use standard UNIX wildcards in the Index name, allowing you to match several index files (for example, `snort-*` matches `snort-123` and `snort-abc`).
 - b. Inputs are enclosed by square brackets `[]`.
 - c. Inputs point to the path location (`["_id"]`).
 - d. Nested locations should be enclosed in one set of brackets, encompassed in quotes, and separated by commas (`["_source","src_ip"]`).



Discover network devices automatically

Field Name Mapping

Index name	Remove index
<code>__default__</code>	<input type="button" value="Remove index"/>

Source IP*

Destination IP*

Source Port*

Destination Port*

Event Start Time*

URL/Domain*

Email Sender*

Email Recipient*

Email Subject*

User*

Event Unique ID*

Event Signature ID*

Event Description*

Event Source Host*

Field name mappings are used to translate Elasticsearch naming to Verodin's native field names. All inputs must be enclosed by square brackets []. The inputs are used to point to the path location; for a nested location, list each level as a value. For example, for the Source IP pointing to elastic_event['_source']['src_ip'] use ['_source','src_ip']. Any non-nested values can simply be the single value in square brackets; for example, to set Event Unique ID to elastic_event['_id'] use ['_id'].

Query Interval (seconds)*

Event Time Adjustment (seconds)*

Name

Save Suspicious Events Yes No

Elasticsearch Integration (Advanced Options)

14. (Optional) Add a new **Index** and configure those fields.



You can delete any Index except the `__default__` by selecting it and clicking **Remove index**.

15. Modify the **Query Interval** and **Event Time Adjustment**, if necessary.
16. (Optional) Assign a **Name**.
17. (Optional) Choose **Yes** to save suspicious events.
18. Click **Submit**.



A message notifies you if there are errors in the Indexes. You must resolve the errors before you can save the integration.

Set up Proxy Assignment

If all outbound connections go through a proxy, you may want to set up a proxy definition and assignment for your integration. For information on setting up your proxy rules, see [Proxy Rules \(https://docs.mandiant.com/home/msv-proxy-rules\)](https://docs.mandiant.com/home/msv-proxy-rules).

Verify connectivity

TO VERIFY CONNECTIVITY TO ELASTICSEARCH

Click **Test** to verify that:

- The Director can communicate with Elasticsearch host IP address on the port specified.
- The Elasticsearch credentials are authorized to perform queries on the index or indexes with relevant data.

Run a Malicious or Captive DNS Action and then review the last run query to verify:

- The custom DNS query works as expected (if configured).