Mandiant

ELASTICSEARCH

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).

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**.

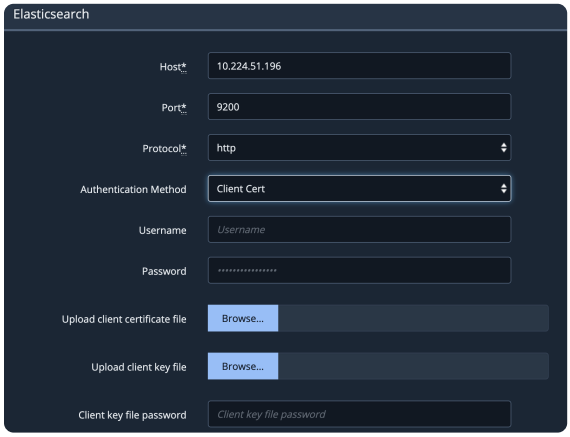

(https://dyzz9obi78pm5.cloudfront.net/app/image/id/629a12d5c9cba0017c2f7dc3/n/elasticsearch.png)

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**.

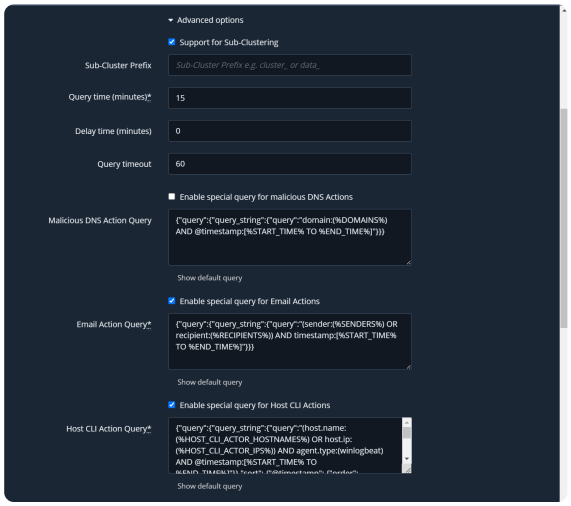

Elasticsearch Integration (Advanced Options)

7. (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.

8. (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.

9. (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.
10. (Optional) Select **Enable query for Email Actions** and configure the **Query**. This query will be only be used when you run Email Actions.
11. (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"]`).

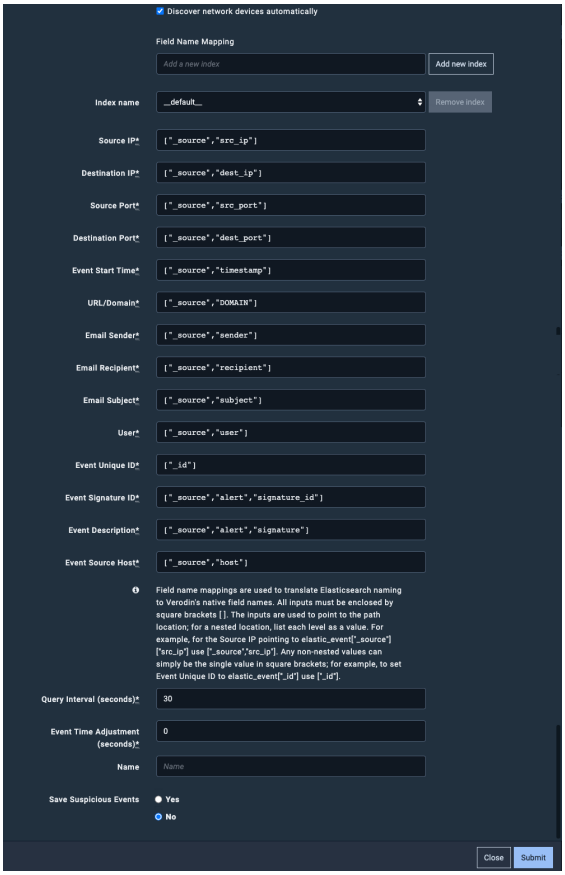

Elasticsearch Integration (Advanced Options)

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)**.

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).