

EclecticIQ Platform user guide

Discover and analyze entities, collaborate, manage users, and configure the platform — 3/4

Last generated: October 20, 2017



©2017 EclecticIQ

All rights reserved. No part of this document may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without written permission from the author, except in the case of a reviewer, who may quote brief passages embodied in critical articles or in a review.

Trademarked names appear throughout this book. Rather than use a trademark symbol with every occurrence of a trademarked name, names are used in an editorial fashion, with no intention of infringement of the respective owner's trademark.

The information in this book is distributed on an "as is" basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author nor the publisher shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this book.

©2017 by EclecticIQ BV. All rights reserved.
Last generated on Oct 20, 2017

Table of contents

Table of contents	2
User guide to EclecticIQ Platform	8
Scope	8
Goal	8
Audience	8
Feedback	8
View the dashboard overview	10
Discover entities	12
Apply discovery filters to view specific entities	12
Apply discovery rules to retrieve specific entities	12
View discovery rules	13
Create discovery rules	13
Save options	14
Edit discovery rules	14
Delete discovery rules	15
Enable and disable discovery rules	15
Manually run discovery rules	16
View entities on the graph	17
Add entities to the graph	17
Open the graph	18
Add entities to an open graph	20
Toggle visualization layouts	24
Act on exposed entities	26
About exposure	26
View exposure	26
Configure exposure	28
Override entity exposure	29
Search for entities	30
Search	30
Search cheatsheet	30
Search query fields	31
Search timeout	32
Analyze entities on the graph	33
About the graph	33
Add entities to the graph	33
Open the graph	35
Analyze entities on the graph	37
Add entities to an open graph	41
Toggle visualization layouts	44
Move around on the graph	45
Undo and redo	46
Filter entities with the histogram	46
Filter entities with the timebar	48
Save and export the graph	49
Group entities in datasets	51
About datasets	51
Static and dynamic datasets	51
Create a dataset	51
Save options	53
Edit a dataset	53
Delete a dataset	53
Add an entity to a dataset	54
Add multiple entities to a dataset	54
View entity details	56
About the entity detail pane	56
Access the entity detail pane	56
Examine the entity overview	56
Manage the entity	59

View entity observables	60
About the entity detail pane	60
Access the entity detail pane	60
Examine the entity observables	60
Apply bulk actions	61
Promote and demote observables	62
Promote an observable	63
Demote an observable	64
Create an indicator from an observable	65
Create a sighting from an observable	65
Add observables to the graph	66
Set observable maliciousness	67
Manually enrich observables	68
Manually add observables	70
Manage the entity	73
Inspect the neighborhood	74
About the entity detail pane	74
Access the entity detail pane	74
Explore the entity neighborhood	74
View related entities and observables	77
Set relationships	80
Set campaign relationships	80
Set course of action relationships	81
Set exploit target relationships	81
Set incident relationships	82
Edit indicator relationships	83
Set report relationships	83
Set sighting relationships	84
Set threat actor relationships	84
Set TTP relationships	85
View related datasets	85
View related workspaces	86
View related tasks	86
Manage the entity	86
View the JSON structure	87
About the entity detail pane	87
Access the entity detail pane	87
Examine the entity JSON structure	87
Manage the entity	87
View entity versions	89
About the entity detail pane	89
Access the entity detail pane	89
View alternative entity versions	89
Manage the entity	90
View entity history	91
About the entity detail pane	91
Access the entity detail pane	91
View the entity history	91
Manage the entity	91
Filter menus	93
Use the quick filters	93
About the Dataset filter	93
Use the context filters	96
Filter by source reliability	98
About source reliability	98
Filter entities by source reliability	99
Filter by tag and taxonomy	101
Search for entities by tag	101
Filter entities by tag on the graph	101
Filter by TLP color	103

About TLP	103
Filter entities by TLP color code	104
Tag and classify entities	106
About taxonomies	106
Predefined taxonomies	106
The Admiralty code	106
The kill chain	108
Create a taxonomy entry	110
Save options	111
Edit a taxonomy entry	111
Delete a taxonomy entry	112
About tags	112
Manually tag entities	113
Auto-tag entities	113
Create tags	113
Delete tags	114
Search for entities by tag	114
Filter entities with the histogram	115
Filter entities with the timebar	117
Entity rules	118
Entity rules	118
Add an entity rule	118
Save options	123
Edit rules	123
Delete rules	124
Filter rules	124
Example	124
Observable rules	129
Observable rules	129
Add an observable rule	129
Save observable rules	133
Derivation and levels	133
Original + level 1	135
Original + level 2	135
Derived + level 1	136
Derived + level 2	136
Edit rules	137
Delete rules	137
Filter rules	138
Example	138
Observables	142
View matching observables	142
Delete matching observables	144
Enrichment rules	145
Automatically enrich entities	145
Manually enrich entities	145
View enrichment rules	149
Add enrichment rules	149
Save options	150
Edit enrichment rules	150
Delete enrichment rules	151
Discovery rules	152
Apply discovery filters to view specific entities	152
Apply discovery rules to retrieve specific entities	152
View discovery rules	153
Create discovery rules	153
Save options	154
Edit discovery rules	154
Delete discovery rules	155
Enable and disable discovery rules	155

Manually run discovery rules	156
Workspaces	157
Workspace types	157
Access workspaces	158
Create a workspace	158
Save options	159
Edit a workspace	159
Archive a workspace	160
Restore a workspace	160
Delete a workspace	160
Toggle between personal and public workspace	161
Add collaborators	161
Create and review tasks	163
Create a task	163
Save options	165
View tasks	165
View tasks created by or assigned to the current user	165
View tasks by status	166
View task details	166
Edit tasks	166
Edit task status	167
Write and review comments	169
Add a comment to a workspace	169
Edit and delete workspace comments	169
Add a comment to a task	169
Edit and delete task comments	170
Work with entities and attachments	171
Get an overview on the workspace dashboard	171
Work with entities, datasets, and attachments	171
Work with saved graphs	172
Review exposure	173
View the workspace history	174
Manage users	175
View users	175
Create users	176
Save options	176
Edit users	177
Edit your user profile	177
Change avatar image	177
Edit your user details	177
Change your password	177
Change user password	178
Revoke user access	178
Manage groups	179
View groups	179
Create groups	180
Save options	180
Edit groups	181
Add users to a group	181
Remove users from a group	181
Edit users in a group	182
Delete groups	182
Manage roles	183
View roles	183
Create roles	184
About permissions	184
Edit roles	185
Delete roles	185
Manage automation users	186
Create an automation group	186

Save options	187
Create an automation role	187
About permissions	187
Create an automation user	188
Get the automation group meta.source ID	189
Step 1 of 2: get the group ID	189
Step 2 of 2: get the group source ID	189
Get the automation group meta.source ID example	190
Get the group ID	190
cURL API request — fetches the group meta.source	190
API response — returns the group meta.source	190
Authentication	191
Auth request	192
Auth response	192
View the help	194
Set host name and timeout	195
Configure general server settings	195
Delete server settings	196
Configure the proxy	197
Configure proxy settings	197
Update proxy settings	198
Delete proxy settings	198
Configure email	200
Configure email settings	200
Test email	201
Delete email settings	201
Register the license	202
Update license information	202
Delete license information	203
Configure STIX	204
Configure STIX settings	204
Delete STIX settings	205
Configure TAXII	207
About TAXII services	207
View TAXII services	208
Add a TAXII service	208
Configure the general options	208
Configure specific options per service	209
Discovery service	209
Collection management	209
Inbox	209
Poll	209
View system jobs	211
View jobs	211
Terminate jobs	212
Audit the system	214
Enable auditing	214
View audit logs	214
View audit logs in the web interface	215
View audit logs in Kibana	215
Search for audit logs in Kibana	217
Adjust the update interval	217
Search by level	217
Search by tag	218
Search with Boolean operators	219
Check system health	220
Check system health via the GUI	220
Check system health via the API	222
API endpoint	222
Status request	222

Status response	223
Monitor system health	226
Tools	226
Core components	226
Monitoring	227
Monitor components with Supervisor	228
Supervisor actions	231
Monitor components with systemd	232
Monitor processes	234
Monitor ingestion queues with Redis	234
Monitor running tasks with Celery	235

User guide to EclecticIQ Platform

This user guide helps you configure the main options of the platform, as well as familiarize with EclecticIQ Platform, so that you can start collecting and analyzing potential threats efficiently.

Scope

The user guide to EclecticIQ Platform aims at providing clear and to-the-point help to get you acquainted with the threat intelligence platform, so that you can configure it as needed, and you can use it to collect and analyze intelligence on potential threats, as well as share it and collaborate with other analysts.

Although it is not a complete reference manual, this guide shows end-users how they can use the platform and its rich feature set to collect data, to analyze and investigate potential threats, and to collaborate and share intelligence with other analysts.

Goal

Learn how to incorporate the platform in your daily workflow as a powerful tool to:

- Automate data ingestion
- View, edit, create, and delete platform entities
- Enrich entities with additional contextual details
- Analyze entities on the graph to identify potential threats and their relationships
- Search, filter, and slice data using rules
- Share your findings and collaborate

Audience

This document targets the following audience:

- Cyber threat intelligence analysts
- Cyber threat intelligence specialists

Feedback

No one reads manuals, ever. We know.

Yet, we strive to give you clear, concise, and complete documentation that helps you get stuff done neatly.

We are committed to crafting good documentation, because life is too short for bad doc.

We appreciate your comments, and we'd love to hear from you: if you have questions or suggestions, drop us a line and share your thoughts with us!

 The Product Team

View the dashboard overview

The dashboard is the default point of entry to the platform after signing in. Go back to the dashboard any time you want to get an overview of the platform status at a glance.

The dashboard is the default point of entry page you land on after successfully signing in to the platform. The dashboard gives you a quick view of the current overall status of the platform.

Depending on the platform configuration, the dashboard view may differ slightly from the description given here.

The dashboard gives you a bird's-eye view of the status of your intelligence within the platform. The dashboard gauges convey core information visually using charts and diagrams. You can assemble any number of modular gauges, among the available ones, to map the platform intelligence landscape as needed.

This is an overview of the default dashboard gauges you can enable and disable:

Gauge	Description
Entity count	Shows a number representing the total amount of entities currently stored in the platform.
Entities per producer	Filters the total amount of ingested entities, based on the corresponding producer. Entity aggregation is based on the <code>data.producer.identity.name</code> field.
Entities per source	Filters the total amount of ingested entities, based on the corresponding data source.
Entities per source in the last 24 hours	Filters the amount of ingested entities in the past 24 hours, based on the corresponding data source.
Entities per source in the last week	Filters the amount of ingested entities in the past week, based on the corresponding data source.
Entities per source in the last month	Filters the amount of ingested entities in the past month, based on the corresponding data source.
Entities per type	Filters the total amount of ingested entities, based on the entity type.
Entities per type in the last 24 hours	Filters the amount of ingested entities in the past 24 hours, based on the entity type.
Entities per type in the last week	Filters the amount of ingested entities in the past week, based on the entity type.
Entities per type in the last month	Filters the amount of ingested entities in the past month, based on the entity type.
Entities per destination	Filters the total amount of ingested entities, based on the corresponding destination. A destination is an outgoing feed where entities are published to.
Entities per destination in the last 24 hours	Filters the amount of ingested entities in the past 24 hours, based on the corresponding destination. A destination is an outgoing feed where entities are published to.
Entities per destination in the last week	Filters the amount of ingested entities in the past week, based on the corresponding destination. A destination is an outgoing feed where entities are published to.

Gauge	Description
Entities per destination in the last month	Filters the amount of ingested entities in the past month, based on the corresponding destination. A destination is an outgoing feed where entities are published to.
Errors over time	Shows how many errors occurred in the system in the past 24 hours.
Logs per component	Shows the number of log messages that have been recorded per each platform component.

Discover entities

Discovery helps you explore ingested intelligence. Use discovery filters and rule-based searches to retrieve specific cyber threat information.

Apply discovery filters to view specific entities

The **Discovery** page returns an overview of selected ingested entities, after applying rule-based search queries. You can refine the results by applying one or more quick filters. They are available on the left-hand navigation sidebar:

- On the top navigation bar click **Discovery**.
- By default, quick filters are switched off.
 - To toggle quick filter visibility click
 - To toggle quick filter visibility click
- On the left-hand navigation sidebar click a filter group name to expand the corresponding sub-nodes:
 - **Entity type**: select one or more checkboxes to include in the filtered results only the specified entity types.
 - **Source**: select one or more checkboxes to include in the filtered results only the specified entity sources.
 - **TLP**: select one or more checkboxes to include in the filtered results only entities flagged with the specified TLP color codes.
 - **Date**: select a time interval to include in the filtered results only entities ingested between the specified start and end dates.
 - **Reliability**: select one or more checkboxes to include in the filtered results only entities with the specified level(s) of reliability.
 - **Discovery rules**: select one or more checkboxes to include in the filtered results only entities matching the specified rule criteria.
 - **Dataset**: select one or more checkboxes to include in the filtered results only entities belonging to the specified datasets.
The **Dataset** filter is not available when the results do not include any entities belonging to at least one dataset.

You can stack and combine filters as you need.

For example, you can create a filter to retrieve only indicators ingested from Hailataxii in the first two weeks of last month, and whose reliability flag is either A (completely reliable) or B (usually reliable).

Apply discovery rules to retrieve specific entities

The **Discovery** service is a rule-based feature looking for cyber threat information that satisfies specific search criteria. You define the search criteria in a search query. The query sets the scope for the discovery rule. If you want, you can further restrict the discovery rule context by selecting one or more workspaces and/or workspace types. Query task execution is capped: the response can return max. 500 matches.

In the platform discovery rules work like configurable, specialized intel fetchers:

- Configurable because you can define discovery rules as necessary.

- Specialized because the rules use search queries to focus on a specific search scope.

When you execute a discovery rule for the first time, it runs incrementally as a provider: the first run returns matching data, up to a maximum of 500 entities, *since the beginning of time*; that is, there is no start time setting to limit the discovery scope to a specific starting point in the past.

Following runs execute the specified query starting from the previous successful run, and they discover only entities added since the previous successful execution of the same rule. Repeated runs return all discovered entities since the previous successful execution of the same query.

If you want to run a discovery task without this temporal constraint, you need to create a new discovery rule.

Editing a rule does not affect this behavior. If you want a discovery query to go through all available data since the beginning of time, you need to create a new rule, and then you need to run it for the first time.

You can also edit a discovery rule, and then click **Save and re-run for all time**.

This option saves any changes, resets the execution time counter, and then it runs the rule task without applying any time constraint.

The run returns matching data for the rule, up to a maximum of 500 results, *since the beginning of time*; that is, there is no start time setting to limit the discovery scope to a specific starting point in the past.



- When a rule is active, it automatically runs every 15 minutes.
- Query task execution is capped: the response can return max. 500 matches.
- Discovery search queries use the **Elasticsearch query syntax** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html>).

View discovery rules

To view a list of all saved discovery rules, do the following:

- On the top navigation bar click **⚙ > Rules > Discovery**.
- **Rules > Discovery** shows an overview of the existing discovery rules.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.

Create discovery rules



Input fields marked with an asterisk are required.

To create a new discovery rule, do the following:

- On the top navigation bar click **⚙ > Rules > Discovery**.
- Click the **+ Rule** button.

- Fill out the **Rules > Discovery > Create** form with the necessary details to create the new rule:
 - **Name:** enter a name to describe the rule. It should be descriptive and easy to remember.
Example: *China or Russia, 1 year till now*
 - **Description:** enter a short description to briefly explain what the rule does, its purpose, and the type of data it looks for.
Example: *Discovers any indicator data types having either “China” or “Russia” as a tag, and whose creation date falls in the range “one year ago until now”.*
 - **Search query:** the search query you want to run when executing the rule. It should do what you explain in the rule description field. Search queries for discovery rules and rules in general use the **Elasticsearch query syntax** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html>).
Example: `data.type:indicator OR entity.tags:China OR entity.tags:Russia AND created_at:[now-1y TO now]`
 - **Correlated workspaces:** you can select one or more workspaces to focus the search only on those entities that are associated with the selected workspaces. To remove a selection from the input field, click the **✕** icon corresponding to the item(s) you want to remove.
Example: *IOCs originating in China and Russia*
 - **Correlated workspaces types:** if you want, you can specify one or more workspace types to focus the search only on those entities that are related to all workspaces of a specific type. To remove a selection from the input field, click the **✕** icon corresponding to the item(s) you want to remove.
Example: *Topic*
 - **Enabled:** select or deselect this checkbox to enable or disable the rule.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new:** saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate:** saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit discovery rules

To edit a rule, do the following:

- On the top navigation bar click **⚙ > Rules > Discovery**.
- On the rule overview, click the row corresponding to the rule you want to modify.
- An overlay slides in from the side of the screen. It displays detailed rule information in a flash-card format.
- On the rule detail view, select **Actions > Edit**.
- On the **Rules > Discovery > Edit** form, you can change the field inputs as appropriate.
- Click **Save** to store your changes, or **Cancel** to discard them.

Alternatively:

- On the top navigation bar click **⚙ > Rules > Discovery**.

- On the rule overview, click the  icon on the row corresponding to the rule you want to modify.
- On the **Rules > Discovery > Edit** form, you can change the field input as appropriate.
- Click **Save** to store your changes, or **Cancel** to discard them.



You can also edit a discovery rule, and then click **Save and re-run for all time**.

This option saves any changes, resets the execution time counter, and then it runs the rule task without applying any time constraint.

The run returns matching data for the rule, up to a maximum of 500 results, *since the beginning of time*; that is, there is no start time setting to limit the discovery scope to a specific starting point in the past.

Delete discovery rules

To delete a rule, do the following:

- On the top navigation bar click  > **Rules > Discovery**.
- On the rule overview, click the  icon on the row corresponding to the rule you want to delete.
- From the pop-up context menu, select **Delete**.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.
- The discovery rule is deleted.

Enable and disable discovery rules

To manually enable and disable an existing rule, do the following:

- On the top navigation bar click  > **Rules > Discovery**.
- On the rule overview, click the row corresponding to the rule you want to run manually.
- An overlay slides in from the side of the screen. It displays detailed rule information in a flash-card format.

On the **Details** tab you can enable and disable the rule.

- If the rule is disabled:
 - Click **Enable**.
 - The button name changes to **Enabled** to notify that the rule is active.
 - A pop-up dialog asks you whether you want to run the rule right away.
 - A notification message confirms enabling the rule.
- If the rule is enabled:
 - Click **Disable**.
 - The button name changes to **Disabled** to notify that the rule is inactive.
 - A notification message confirms disabling the rule.

Manually run discovery rules

You can bypass automatic execution and decide to manually run a rule, for example to test it immediately after creating it.

To manually run a rule, do the following:

- On the top navigation bar click **⚙ > Rules > Discovery**.
- On the rule overview, click the row corresponding to the rule you want to run manually.
- An overlay slides in from the side of the screen. It displays detailed rule information in a flash-card format.
- On the **Details** tab, either click the **Run now** button, or select the **Actions > Run now** menu option.

After completing the run, you can review the outcome on the **Details** tab:

- Under the **Status** column you can check the execution outcome.

 Started	The task run has been initiated, it has been added to the queue, and it is waiting to be executed.
 Success	The task run completed correctly.
 Error	The task run failed. Click the status icon to view an error message and a traceback with more details about the failure. This information can be helpful to troubleshoot the issue.

- Under the **Results** column you can see whether the discovery action yielded any new results matching the rule criteria.

View entities on the graph

Load entities on the graph to analyze them, explore relationships, and manipulate data easily and powerfully.

The graph is a powerful tool to examine entities, and to look for meaningful cues during an analysis. On the graph you can add, remove, and filter entities, examine and manipulate them to gain insights, inspect relationships, and draw a map of the threat landscape under investigation.

The graph is an easy and intuitive way to review complex relationships, and to look at scenarios from multiple angles using different layouts, the histogram filtering options, and the timebar.

Add entities to the graph

You can manually export an entity from almost anywhere in the platform.

To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.

- On the active view, browse to the entity you want to view on the graph.
- Click the  icon corresponding to the entity you want to view on the graph.
- From the drop-down menu select **Add to graph**.

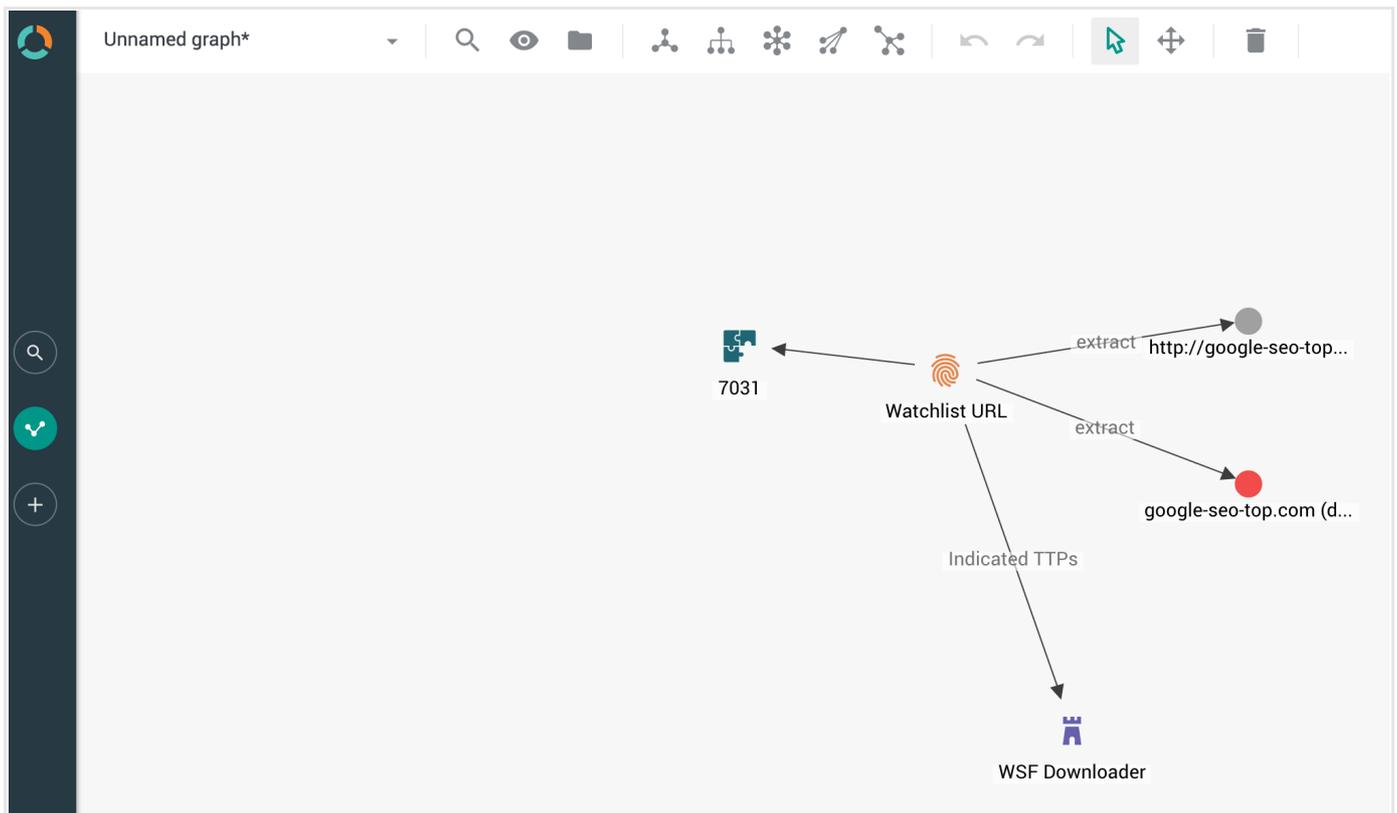
Alternatively:

- On the active view, browse to the entity you want to view on the graph.
- Click anywhere on the row corresponding to the entity you want to view on the graph.
The entity detail pane slides in from the side of the screen.
- On the bottom half of the detail pane, click **Actions**.
- From the pop-up menu select **Add to graph**.



You can load on the graph only published entities and observables.
You cannot view draft entities on the graph.

To quickly view if the entities are loaded on the graph, hover the mouse pointer on the graph icon on the top navigation bar to display a thumbnail view of the current graph canvas:

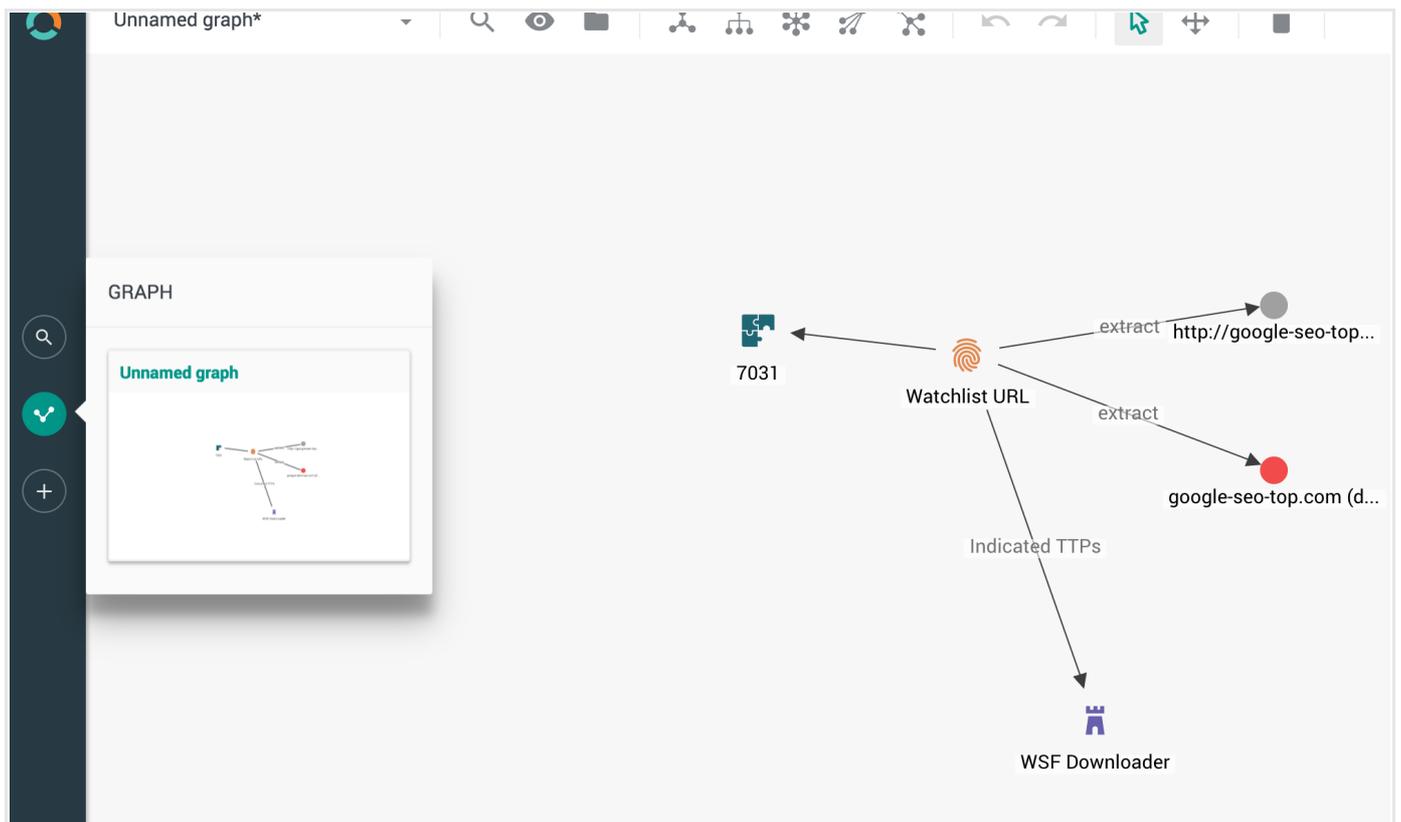


Open the graph

You can open the graph in one of the following ways:

Open the graph by clicking the graph icon on the top navigation bar

- On the top navigation bar click the graph icon.
- By default, the graph loads the most recently open graph session.
- If the loaded graph has never been saved before, the default name is **Unnamed graph***. An asterisk appended to the graph name indicates that the currently loaded data on the graph is not saved yet.
- To save the loaded data as a graph, click the graph menu and select **Save as**.
- To discard the loaded data and start from a clean canvas, click the graph menu and select **New**.



Open the graph from the Saved graphs section in a workspace

- On the top navigation bar click **Workspaces** > <workspace_name> > **Browse** > **Saved graphs**.
- Click the  icon on the graph tile you want to open.
- From the drop-down menu select **Load**.

Open the graph from the Neighborhood tab on the entity detail pane

- On the top navigation bar click **Browse**, **Discovery**, or **Exposure**. Any of these selections directs you to entity overview pages.
- On the selected entity overview page, click anywhere on a row corresponding to the entity you want to load on the graph.
- An overlay slides in from the side of the screen to display the entity detail pane.
- On the entity detail pane, go to the **Neighborhood** tab.

- On the **Neighborhood** tab, click the small graph image, if available, to load the corresponding content onto the larger graph canvas.

The screenshot displays the 'Domain Traffic Blocking COA' interface. At the top, it shows 'Ingested: Today at 7:20 AM' and 'Incoming feed: TAXII Stand Samples' with a 'TLP Red' indicator. The navigation tabs include OVERVIEW, OBSERVABLES, NEIGHBORHOOD (selected), JSON, VERSIONS, and HISTORY. The main area is titled 'Graph' and contains a network diagram with a central node labeled 'Domain' and several surrounding nodes connected by arrows, including 'Block malicious links on web proxies', 'Indicated TTPs', 'Suggested courses of action', 'Super Secret Proprietary Response COA', 'Phishing Target Notification COA', 'Email cleanup COA', 'Email blocking COA', and 'COA Phishing target identification'. Below the graph, the 'Directly related entities' section shows a table with columns for 'Relationship type' and 'Related entity'. One entry is visible: 'suggested_coas' with a left-pointing arrow and the 'US-China' Phishing Indicator icon. An 'EDIT RELATIONSHIPS' button is located below the table. At the bottom, there is an 'Actions' dropdown menu.

Add entities to an open graph

When you are working on an open graph, you can add and remove entities on the fly without leaving the graph canvas. Click the **Add from search**, **Add from discovery**, or **Add from workspaces** buttons on the graph top navigation bar to load more entities on the graph:



When you add entities from the search and the discovery services, as well as from one or more workspaces, you can apply filters to target specific entity types, data sources, and time ranges:

- **Entity type:** select one or more checkboxes to select and load on the graph only the specified entity types.
- **Source:** select one or more checkboxes to select and load on the graph only entities belonging to the specified data sources.
- **Date:** select a time interval to select and load on the graph only entities ingested between the specified start and end dates.
- **Dataset:** select one or more checkboxes to select and load on the graph only entities belonging to the specified datasets.
The **Dataset** filter is not available when the results do not include any entities belonging to at least one dataset.

Add entities from search

You can add entities from the default search page. You can also run platform-wide search queries to look for specific entities to load, as well as filter search results to further narrow down your scope.

To open the search pane from the graph, do the following:

- Click **Q Add from search**.
- Enter search terms, **Elasticsearch queries** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>), or use the drop-down search filters to search for the entities you want to load on the graph.
The search or the filters return all the matches meeting your search/filtering criteria.
- Select the checkboxes corresponding to the entities you want to add to the graph.
- Click **Add to graph**.

Jnnamed graph*

Filter...

Filters: Entity types ▾ Source ▾ Date ▾ 57064 results

<input type="checkbox"/>	Title	Source	Ingestion time	
<input type="checkbox"/>	 untitled	Testing Group	24.07.2017 11:09	
<input type="checkbox"/>	 untitled	Testing Group	24.07.2017 11:09	
<input type="checkbox"/>	 untitled	Testing Group	24.07.2017 11:09	
<input type="checkbox"/>	 untitled	Testing Group	24.07.2017 18:01	
<input type="checkbox"/>	 !! E un fapt bine stabilit că cititorul va fi sustr...	Testing Group	02.08.2017 12:46	
<input type="checkbox"/>	 sid:2405021 "ET CNC Shadowserver Report...	INCTENRICH_System_Default	02.08.2017 14:58	

1 - 6 of 57 064

ADD TO GRAPH

Add entities from the discovery service

To limit your scope to discovered entities only, do the following:

- Click  **Add from discovery**.
- Enter search terms, **Elasticsearch queries** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>), or use the drop-down search filters to search for the entities you want to load on the graph. The search or the filters return all the matches meeting your search/filtering criteria.
- Select the checkboxes corresponding to the entities you want to add to the graph.
- Click **Add to graph**.

Jnnamed graph*

× Default Workspace

Filter...

Filters: Entity types ▾ Source ▾ Date ▾ Datasets ▾ 101 results

<input type="checkbox"/>	Title	Source	Ingestion time	
<input checked="" type="checkbox"/>	 This domain tor.globenet.org has been identified as a TOR ...	tor	22.08.2017 17:15	
<input type="checkbox"/>	 This domain torexit1.duffsdevice.com has been identified a...	tor	22.08.2017 17:14	
<input type="checkbox"/>	 This domain cassel.dogeneral.net has been identified as a ...	tor	22.08.2017 17:14	
<input checked="" type="checkbox"/>	 This domain sing-tor.cryptoligarch.com has been identified ...	tor	22.08.2017 17:18	
<input checked="" type="checkbox"/>	 This domain 55-21-50-84.dyn.estpak.ee has been identified ...	tor	22.08.2017 17:14	
<input type="checkbox"/>	 This ipAddress 202.85.233.34 has been identified as a TOR ...	tor	22.08.2017 17:18	

1 - 6 of 101

⏪ < > ⏩

ADD TO GRAPH

Add entities from workspaces

To load on the graph entities belonging to one or more specific workspaces, do the following:

- Click  **Add from workspaces**.
- From the drop-down menu select one or more workspaces you want to load entities from.
- Enter search terms, **Elasticsearch queries** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>), or use the drop-down search filters to search for the entities you want to load on the graph. The search or the filters return all the matches meeting your search/filtering criteria.
- Select the checkboxes corresponding to the entities you want to add to the graph.
- Click **Add to graph**.

Jnamed graph*

x Default Workspace
x ▼

Filters: Entity types ▼ Source ▼ Date ▼ Datasets ▼ 101 results

<input type="checkbox"/>	Title	Source	Ingestion time	
<input checked="" type="checkbox"/>	This domain tor.globenet.org has been identified as a TOR ...	tor	22.08.2017 17:15	
<input type="checkbox"/>	This domain torexit1.duffsdevice.com has been identified a...	tor	22.08.2017 17:14	
<input type="checkbox"/>	This domain cassel.dogeneral.net has been identified as a ...	tor	22.08.2017 17:14	
<input checked="" type="checkbox"/>	This domain sing-tor.cryptoligarch.com has been identified ...	tor	22.08.2017 17:18	
<input checked="" type="checkbox"/>	This domain 55-21-50-84.dyn.estpak.ee has been identified ...	tor	22.08.2017 17:14	
<input type="checkbox"/>	This ipAddress 202.85.233.34 has been identified as a TOR ...	tor	22.08.2017 17:18	

1 - 6 of 101 |< < > >|

ADD TO GRAPH

Toggle visualization layouts

You can switch among different visualization layouts to analyze entities from different perspectives. For example, you can focus on hierarchical relationships, or you may want to examine how a network evolves over time.

The available graph layouts enable you to approach a scenario from multiple angles, so that you can look for patterns, relationships, and structures providing meaningful context.

On the graph top navigation bar, click the icon corresponding to the desired layout to automatically rearrange the view accordingly.



Layout type	Description
-------------	-------------

Layout type	Description
Standard	In the standard layout, links on the graph are a consistent length. Nodes and edges overlap as little as possible, and they are evenly distributed on the graph surface. It offers a consistent and clean view. It is a good starting point to begin analyzing any kind of data and any dataset size, especially when you are looking for patterns and symmetries.
Hierarchy	It is a tree structure with nodes. It displays child nodes horizontally below the corresponding parents. Connections flow top-down through the chart from the original subject. It is an efficient layout to visualize workflows and processes, impact analysis, and hierarchical relationships.
Radial	The radial layout arranges nodes in concentric circles around the original subject in a radial tree. Each set of nodes becomes a new orbit extending outwards from the original parent. This layout works best with networks with a large volume of child nodes to each parent.
Structural	It is similar to the standard layout. However, in the structural layout nodes with similar attributes are grouped together in fans. This visualization provides a clear overview of the clusters within a network, without focusing on a specific one.
Tweak	The tweak layout shows how networks evolve. The layout automatically adapts as links are created and destroyed, so that you can see where and how the changes occur. It is ideal for visualizing the behavior of dynamic and changing graphs.

Act on exposed entities

Exposure shows you what your organization is doing with the ingested cyber threat intelligence, so that you can assess its usage to define courses of actions and other preventive or reactive procedures within the organization.

About exposure

In the 2004 Pixar movie **The Incredibles** (<http://www.imdb.com/title/tt0317705/>), Helen Parr goes to see Edna Mode, only to find out her husband Bob Parr has resumed superhero work. And he's been gone from home for a few days. When Edna asks Helen *"Do you know where he is?"*, Helen cannot answer. Previously in the movie, she had witnessed some changes in her husband's behavior that should have alerted her, but she disregarded that information.

This is exposure in a nutshell.

When platform entities are flagged as exposed, your organization is not making the most of the available cyber threat intelligence (CTI) to drive effective courses of action. Intelligence is either underutilized, or it is ignored.

Exposure helps you assess how your organization uses and leverages CTI: how is CTI affecting the organization? Is the organization using CTI to drive processes to detect, deter, and defeat attacks and to minimize risk? What is working well, and what can be done to improve intel utilization?

Exposure gives you a comprehensive and user-friendly overview that helps you answer these questions by showing you how your organization uses existing CTI, and what it can do to use CTI more efficiently.

View exposure

Exposed entities are ingested and processed. However, their intelligence value is not leveraged to produce follow-up actions.

For example, triggering a detection event in a malware detection application downstream in the system; or a prevention event such as creating a firewall rule; or a community event such as sending a notification message to inform other parties about the possible threat the entity represents.

The entities hold intelligence value that is not consumed.



You first need to configure **Exposure** to specify the filtering criteria the platform should apply when flagging entities as exposed.

After defining the exposure settings you can view exposed entities, based on your configuration.

To view exposed entities, do the following:

- On the top navigation bar click **Exposure**.
- It shows an overview of all exposed entities.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.

- To toggle quick filter visibility click
- To toggle quick filter visibility click
- On the left-hand navigation sidebar click a filter group name to expand the corresponding sub-nodes:
 - **Entity type**: select one or more checkboxes to view exposure details for the specified entity types.
 - **Date**: select a time interval to view exposure details for the entities ingested between the specified start and end dates.
 - **Dataset**: select one or more checkboxes to view exposure details for the entities belonging to the specified datasets.
The **Dataset** filter is not available when the results do not include any entities belonging to at least one dataset.

You can stack and combine filters as you need.

For example, you can create a filter to view exposure details for indicators belonging to the X, Y, and Z datasets, ingested in the first two weeks of last month.

The **Exposure** view shows the following exposure-specific information:

Entity Type	Exposed	Name	Timestamp	Detection	Prevention	Community	Sighting	Refresh
Indicator (130666)	EXPOSED	Fake indicator d0412409924b #96356	08/25/2017 11:35 AM	●	●	●	!	↻
TTP (29533)	EXPOSED	Fake indicator d0412409924b #94518	08/25/2017 11:35 AM	●	●	●	!	↻
Report (86)	EXPOSED	Fake indicator d0412409924b #98541	08/25/2017 11:35 AM	●	●	●	!	↻
Campaign (9)	EXPOSED	Fake indicator d0412409924b #94270	08/25/2017 11:35 AM	●	●	●	!	↻
Incident (9)	EXPOSED	Fake indicator d0412409924b #95011	08/25/2017 11:35 AM	●	●	●	!	↻
Course of action (8)								

- **Exposed**: the **Exposed** label indicates that the entity is exposed, that is, it is not used in any detection, prevention, or community integrations or processes.
- **Detection**: the entity and the intelligence value it holds are being consumed in an integration with an external system. In this case, with a detection system. If the dot is green, the entity information is used to carry out a follow-up action. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.
- **Prevention**: the entity and the intelligence value it holds are being consumed in an integration with an external system. In this case, with a prevention system. If the dot is green, the entity information is used to carry out a follow-up action. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.
- **Community**: the entity and the intelligence value it holds are being consumed in an integration with an external system. In this case, with an information distribution system. If the dot is green, the entity information is used to carry out a follow-up action. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.
- **Sighting**: a ! flag means that the entity has been seen in a secured domain, and there should be a sighting entity recording the occurrence.
- Click the ↻ icon to refresh and update the view.

- ✓ If an entity has been sighted, it is by default exposed no matter the level of integration with external detection, prevention or information distribution systems.

Configure exposure

You can configure Exposure to be as generic or as specific as you need:

- On the top navigation bar click **Exposure**.
- On the left-hand navigation sidebar click click **Settings**.
- On the **Exposure > Settings** page click **Edit exposure settings** to change exposure behavior.

On the configuration page you can define which entities you want to watch for exposure, as well as set filters to minimize unwanted data noise:

- **Entity types:** from the drop-down menu select Entity types to include one or more entity types in the exposure configuration.
The entity types you add here are tracked to assess their exposure.
- **Observable types:** from the drop-down menu select one or more observable types.
This option filters the selected entities to include in the exposure configuration only entities with at least one observable type matching the selection(s) you specify here.
- **Confidence values:** from the drop-down menu select one or more confidence values.
This option filters the selected observable types to include in the exposure configuration only observables whose maliciousness confidence value matches at least one of the selections you specify here.
Confidence corresponds to the value you set under **Data configuration > Rules > Observable > + (Create rule) > Action > Mark as malicious > Confidence**.
- **Entity age:** it defines a time interval ranging from now, that is, the current time, to a point in the past.
It is an integer and it represents days.
Only entities that fall inside this range and that are not older than the number of days specified here are tracked to assess their exposure.
- **Relevancy threshold:** *Relevancy* is a numerical value based on the current time and the estimated start time of the threat. You can use it to sort and filter entities. *0%* = low relevancy — *100%* = high relevancy. Its value is 100% when the current time (*now*) is included between the threat start and end times. Otherwise, its value is 0. If the estimated end time is not available, relevancy is calculated using the estimated start time and the half-life value.
- **Show enrichment observables:** if you select this checkbox, enrichment observables are included and displayed, when available.
- Click **Save** to store your changes, or **Cancel** to discard them.

After configuring exposure behavior, you should configure which outgoing feeds should share and distribute exposure information to external systems and devices, so that the data can trigger appropriate actions and responses as part of a concerted course of action.

- On the top navigation bar click **Exposure**.
- On the left-hand navigation sidebar click click **Outgoing feeds**.

On the **Exposure > Outgoing feeds** page you can define how to publish the ingested CTI to minimize exposure. For example, if you are publishing an outgoing feed to an external detection system, the feed data stream is used to detect potential threats.

On this page you map outgoing feeds to the purpose they serve in the context of an integration with external tools and systems.

Within exposure an unused outgoing feed, or a wrongly mapped outgoing feed — for example, an outgoing feed marked as **Detect** but used to distribute CTI to a relevant community, instead — is flagged as exposed.

For each outgoing feed in the overview, you can select one or more checkboxes to map feed usage as appropriate:

- **Detect**: the outgoing feed is published to an external detection system. The feed data is used to detect potential threats that have infiltrated your organization.
- **Prevent**: the outgoing feed is published to an external prevention system. The feed data is used to prevent potential threats from attacking your organization.
- **Community**: the outgoing feed is published to an external information distribution system. The feed is used to share CTI with other parties within or outside the organization.
- **N.A.**: the outgoing feed is not published to any external system.

Override entity exposure

You can manually override the configured exposure settings for an entity. The **Override exposure** option allows you to reverse the **Detection**, **Prevention**, and **Sighting** exposure values, and to set them to their opposites.

The entity exposure override history is stored in reverse chronological order, based on the time when the change was applied.

To manually change the exposure state of an entity, do the following:

- On the top navigation bar click **Exposure**.
- On the **Exposure** page click the **i** icon on the row corresponding to the entity whose exposure settings you want to override.
- From the context menu select **Override exposure**.
- On the **Override exposure state** dialog window, select the **Override exposure state to (ON or OFF**, depending on the current exposure value) checkboxes to reverse the current exposure value — from **ON** to **OFF** or the other way around — of **Detection**, **Prevention**, and **Sighting**.
- If you want, you can specify a start date for the override value(s) to become effective: from the drop-down menu select the desired start date.
- When you are done, close the dialog window. Your settings are automatically saved.



After confirming and saving a manual exposure override, the override value persists until new content is generated, and the entity is updated.

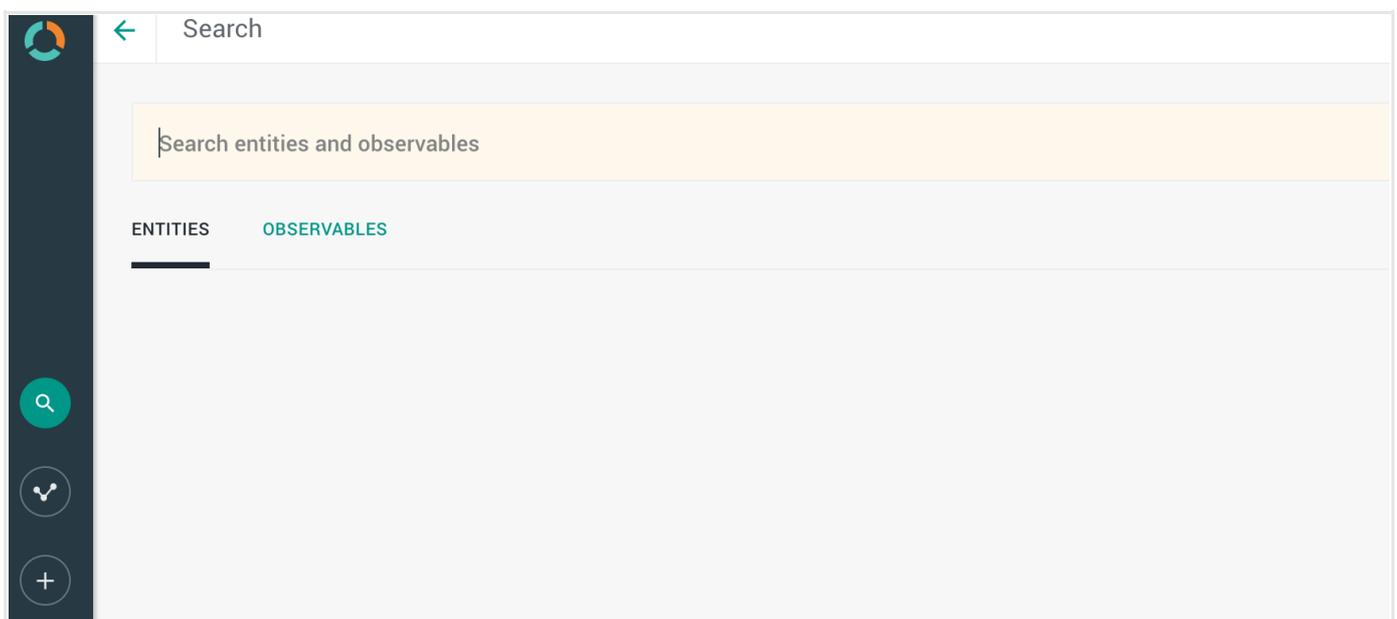
Search for entities

Use the search field to look for entities using literals, wildcards, and regex.

The top bar is your entry point to run platform-wide search and upload operations, and to edit your profile information.

Search

You can find the search box on the top bar:



Enter search terms and search queries, and then press **ENTER** or click the search icon to run the search. Searches you run through this search box are executed platform-wide.



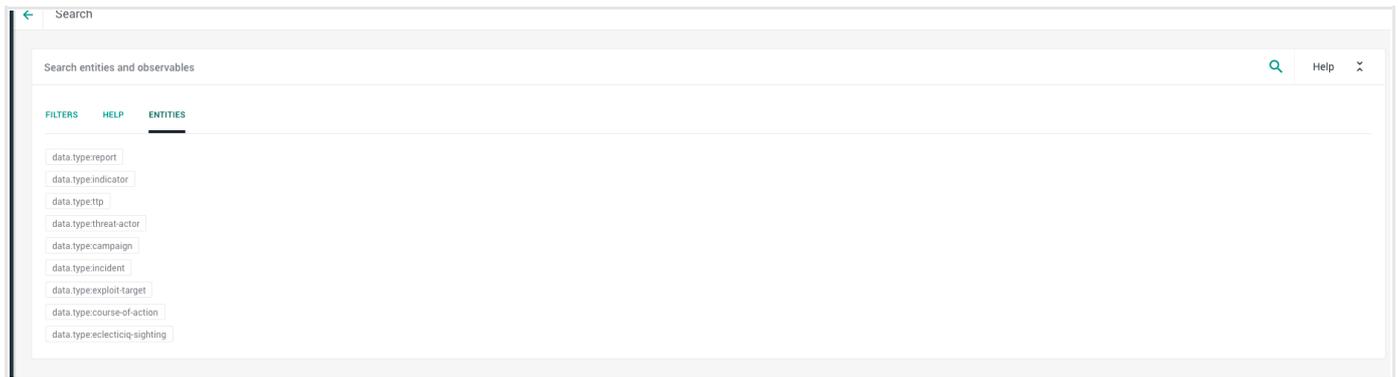
The search functionality uses **Elasticsearch query syntax**

(<https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html>).

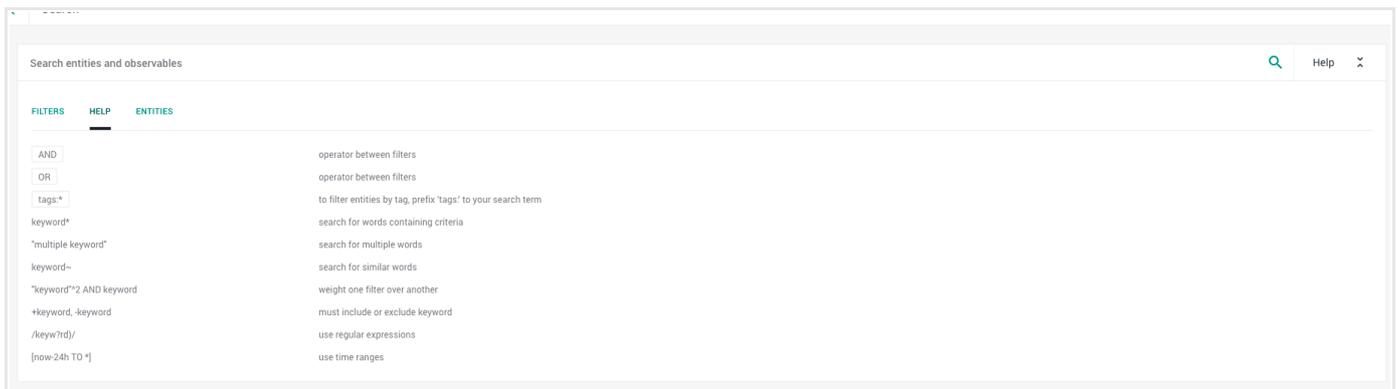
Search cheatsheet

To access a cheatsheet with search examples using entity types, filters, and for help with the search syntax, click **Help** to display thematic drop-down lists with common search queries:

- **Filters:** examples of quick search filters.
- **Help:** examples of regex, Boolean, wildcards, and tag search usage.
- **Entities:** examples of searchable entity types.



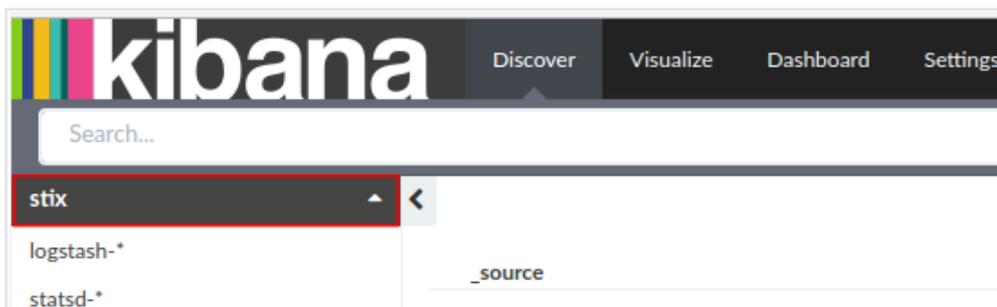
Besides full text search, you can use Boolean operators, wildcards, regex, and you can combine these filtering options to create more refined searches.



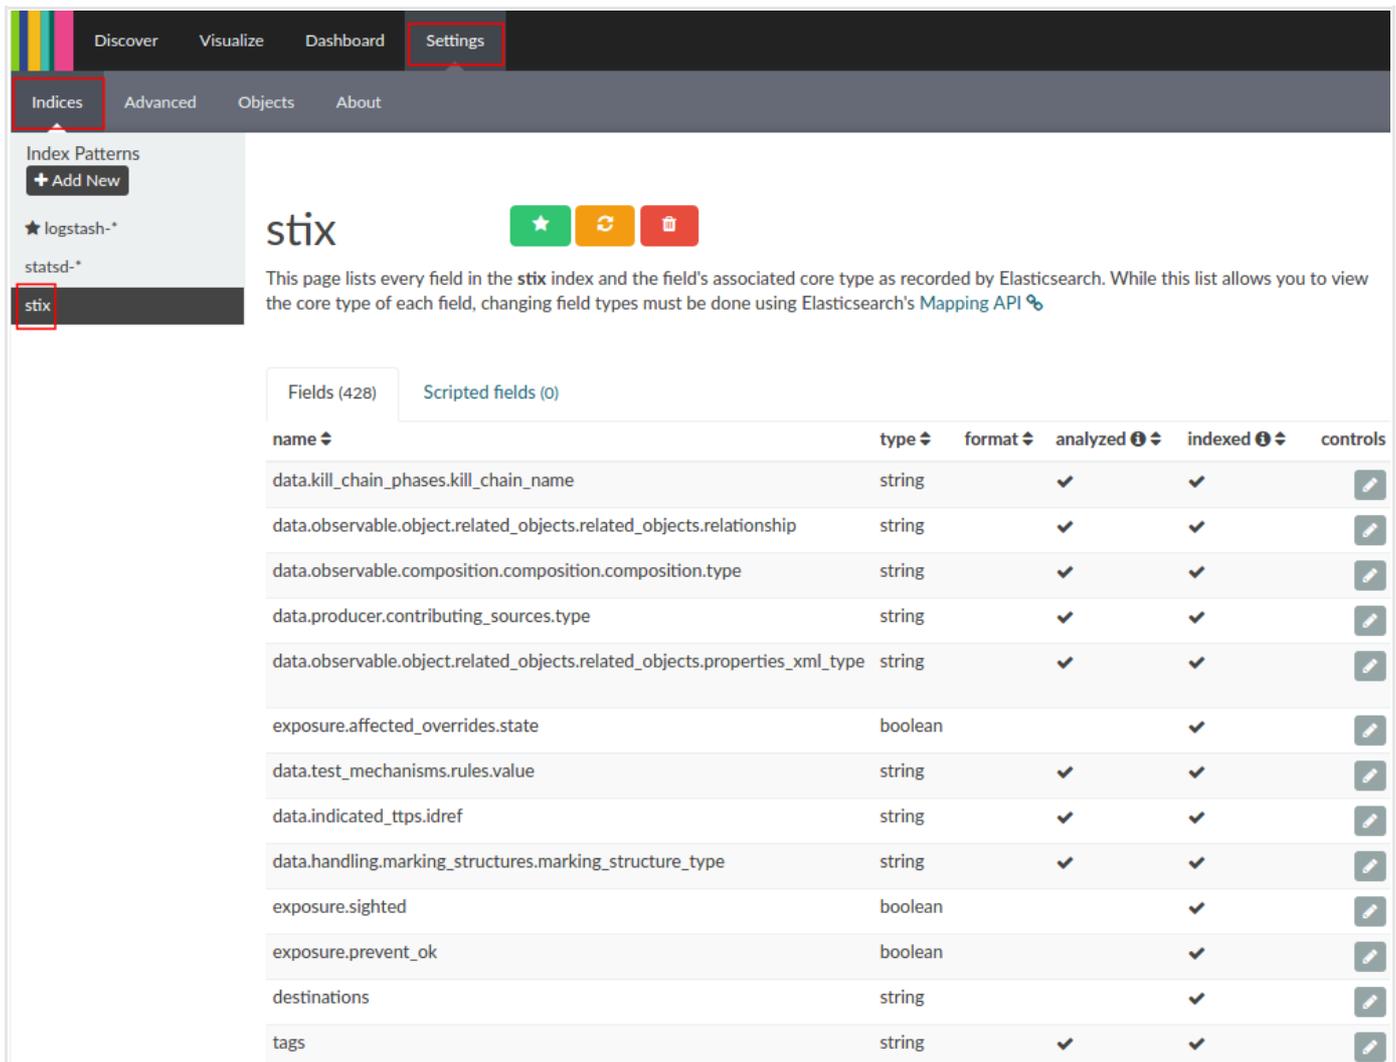
Search query fields

For reference, you can look up a complete list of all available search query fields in Kibana:

- To access Kibana, in the web browser address bar enter a URL with the following format:
https://<platform_host>/api/kibana/app/kibana#/
 Keep the trailing /
 Example: <https://platform.host.com/api/kibana/app/kibana#/>
- Select the **stix** index field:



- On the main menu bar, select **Settings**:



The screenshot shows the Kibana Settings page for the 'stix' index. The navigation bar includes 'Discover', 'Visualize', 'Dashboard', and 'Settings'. The left sidebar shows 'Indices' selected, with 'stix' highlighted. The main content area displays the 'stix' index settings, including a list of fields (428) and scripted fields (0). The table below lists the fields and their associated core types and settings.

name	type	format	analyzed	indexed	controls
data.kill_chain_phases.kill_chain_name	string		✓	✓	
data.observable.object.related_objects.related_objects.relationship	string		✓	✓	
data.observable.composition.composition.composition.type	string		✓	✓	
data.producer.contributing_sources.type	string		✓	✓	
data.observable.object.related_objects.related_objects.properties_xml_type	string		✓	✓	
exposure.affected_overrides.state	boolean			✓	
data.test_mechanisms.rules.value	string		✓	✓	
data.indicated_ttps.idref	string		✓	✓	
data.handling.marking_structures.marking_structure_type	string		✓	✓	
exposure.sighted	boolean			✓	
exposure.prevent_ok	boolean			✓	
destinations	string			✓	
tags	string		✓	✓	

Search timeout

By default, Elasticsearch search queries that do not resolve time out after 20 seconds.

To set a different search timeout value, do the following:

- Open the `/opt/eclecticiq/etc/eclecticiq/platform_settings.py` configuration file.
- Browse to the `ELASTICSEARCH_QUERY_TIMEOUT = '20s'` line.
- Replace the default value with a custom one, for example `ELASTICSEARCH_QUERY_TIMEOUT = '30s'`. The `ELASTICSEARCH_QUERY_TIMEOUT` parameter value represents seconds.
- Save the configuration file.

Analyze entities on the graph

Load entities on the graph to analyze them, explore relationships, and map the context around potential threats.

When you load entities and observables on the graph to analyze them, you are weaving a story to describe the threat scenario under investigation. A good story needs villains. The cyber crime line of business appeals to a wide range of such characters. Villains have motives and goals they want to achieve. To attain them, they engage in actions and behaviors that usually damage a third party. Enter the victim. The outcome of the villain's tactics, techniques, and procedures to produce the intended effects may be beneficial for the villain and detrimental for the victim.

Now you've got the basic elements to build a consistent narrative for a threat scenario:

Threat actors apply TTPs to hit a targeted victim, so that they can achieve their intended effects.

The actors may leverage existing exploit targets to carry out a series of attacks. Their malicious activities may leave some traces behind.

An analyst on the victim's side may pick up on those traces and report them to alert the organization. Following up on the report, another analyst may detect them in a log file.

However, before the victim can react with appropriate measures and procedures, a security breach occurs.

In most real-life cases, the script that builds the narrative of a threat scenario is fragmented and scattered: you have only a few pieces of the puzzle, and a couple of them possibly belong to a different puzzle altogether. The graph canvas is the stage where you analyze, reorganize, restructure, assess, test alternatives, and ultimately position all the pieces in place to produce a factual and consistent narrative that can answer these basic questions:

- What happened?
- When did it happen?
- Where did it happen?
- Why did it happen?
- Who did it?
- Why did they do it?

About the graph

The graph is a powerful tool to examine entities, and to look for meaningful cues during an analysis. On the graph you can add, remove, and filter entities, examine and manipulate them to gain insights, inspect relationships, and draw a map of the threat landscape under investigation.

The graph is an easy and intuitive way to review complex relationships, and to look at scenarios from multiple angles using different layouts, the histogram filtering options, and the timebar.

Add entities to the graph

You can manually export an entity from almost anywhere in the platform.

To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.

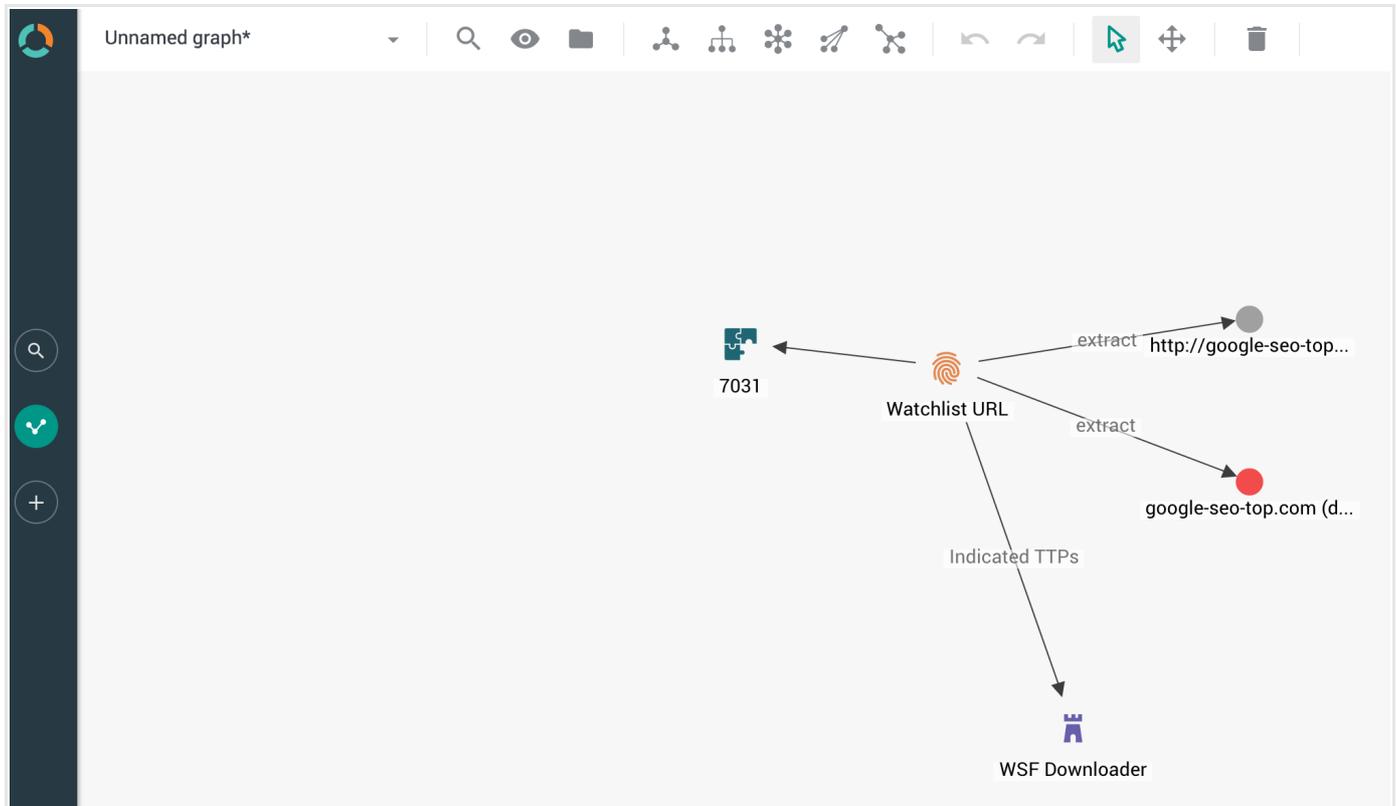
- On the active view, browse to the entity you want to view on the graph.
- Click the  icon corresponding to the entity you want to view on the graph.
- From the drop-down menu select **Add to graph**.

Alternatively:

- On the active view, browse to the entity you want to view on the graph.
- Click anywhere on the row corresponding to the entity you want to view on the graph. The entity detail pane slides in from the side of the screen.
- On the bottom half of the detail pane, click **Actions**.
- From the pop-up menu select **Add to graph**.

 You can load on the graph only published entities and observables.
You cannot view draft entities on the graph.

To quickly view if the entities are loaded on the graph, hover the mouse pointer on the graph icon on the top navigation bar to display a thumbnail view of the current graph canvas:

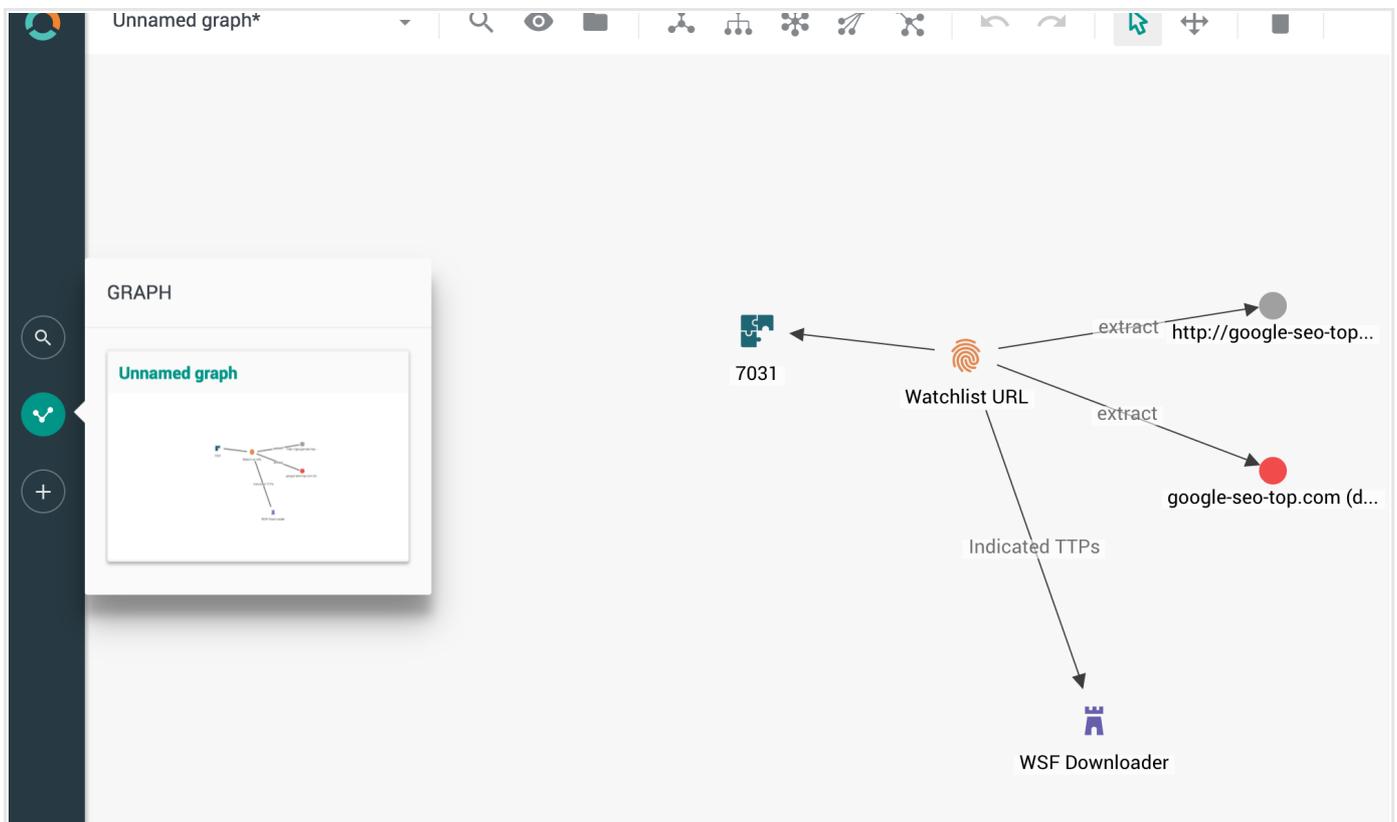


Open the graph

You can open the graph in one of the following ways:

Open the graph by clicking the graph icon on the top navigation bar

- On the top navigation bar click the graph icon.
- By default, the graph loads the most recently open graph session.
- If the loaded graph has never been saved before, the default name is **Unnamed graph***.
An asterisk appended to the graph name indicates that the currently loaded data on the graph is not saved yet.
- To save the loaded data as a graph, click the graph menu and select **Save as**.
- To discard the loaded data and start from a clean canvas, click the graph menu and select **New**.



Open the graph from the Saved graphs section in a workspace

- On the top navigation bar click **Workspaces** > <workspace_name> > **Browse** > **Saved graphs** .
- Click the  icon on the graph tile you want to open.
- From the drop-down menu select select **Load**.

Intelligence C&C servers BNL Dashboard Browse Exposure History Tasks (4) Comments

C&C servers BNL [View public dashboard](#)

Short description

Saved graphs

07/15/2017 6:59 PM

08/21/2017 1:13 PM

LAST TUESDAY AT 10:07 AM

Open the graph from the Neighborhood tab on the entity detail pane

- On the top navigation bar click **Browse**, **Discovery**, or **Exposure**. Any of these selections directs you to entity overview pages.
- On the selected entity overview page, click anywhere on a row corresponding to the entity you want to load on the graph.
- An overlay slides in from the side of the screen to display the entity detail pane.
- On the entity detail pane, go to the **Neighborhood** tab.
- On the **Neighborhood** tab, click the small graph image, if available, to load the corresponding content onto the larger graph canvas.

×
Domain Traffic Blocking COA

🔗
Ingested: Today at 7:20 AM
Incoming feed: TAXII Stand Samples
● TLP Red

OVERVIEW
OBSERVABLES
NEIGHBORHOOD
JSON
VERSIONS
HISTORY

Graph

Directly related entities

Relationship type	Related entity
suggested_coas ←	🔗 "US-China" Phishing Indicator

EDIT RELATIONSHIPS

Actions ▾

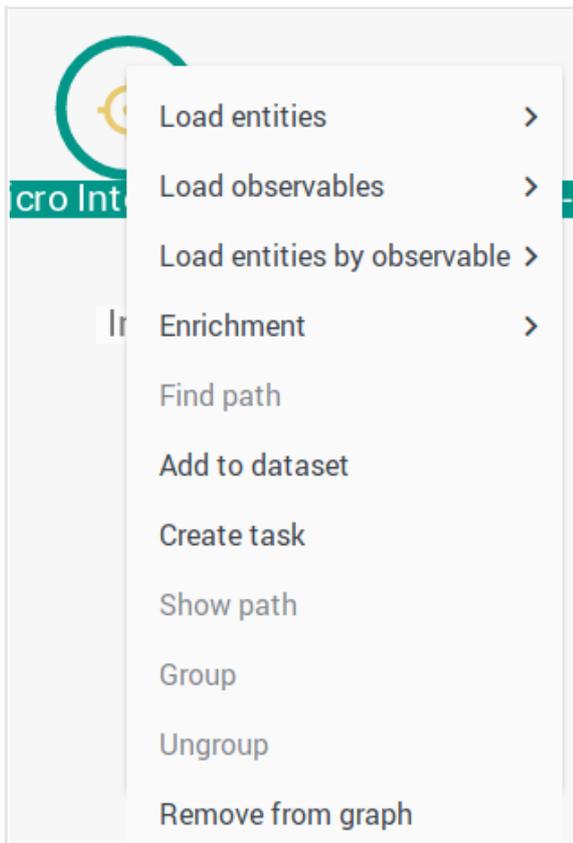
Analyze entities on the graph

The graph is a powerful toolbox to analyze cyber threat intelligence.

The right-click context menu is your Swiss Army knife to load entities, discover relationships, enrich entities, group and ungroup them.

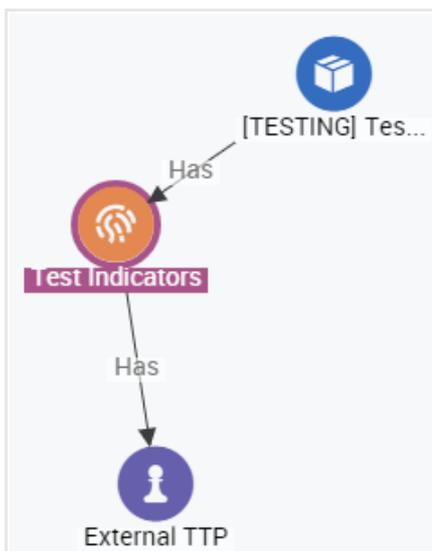
Grayed-out options in the menu are disabled for the selected item.

You can double-click an entity on the graph to view more details about it. An overlay slides in from the side of the screen with detailed entity information.

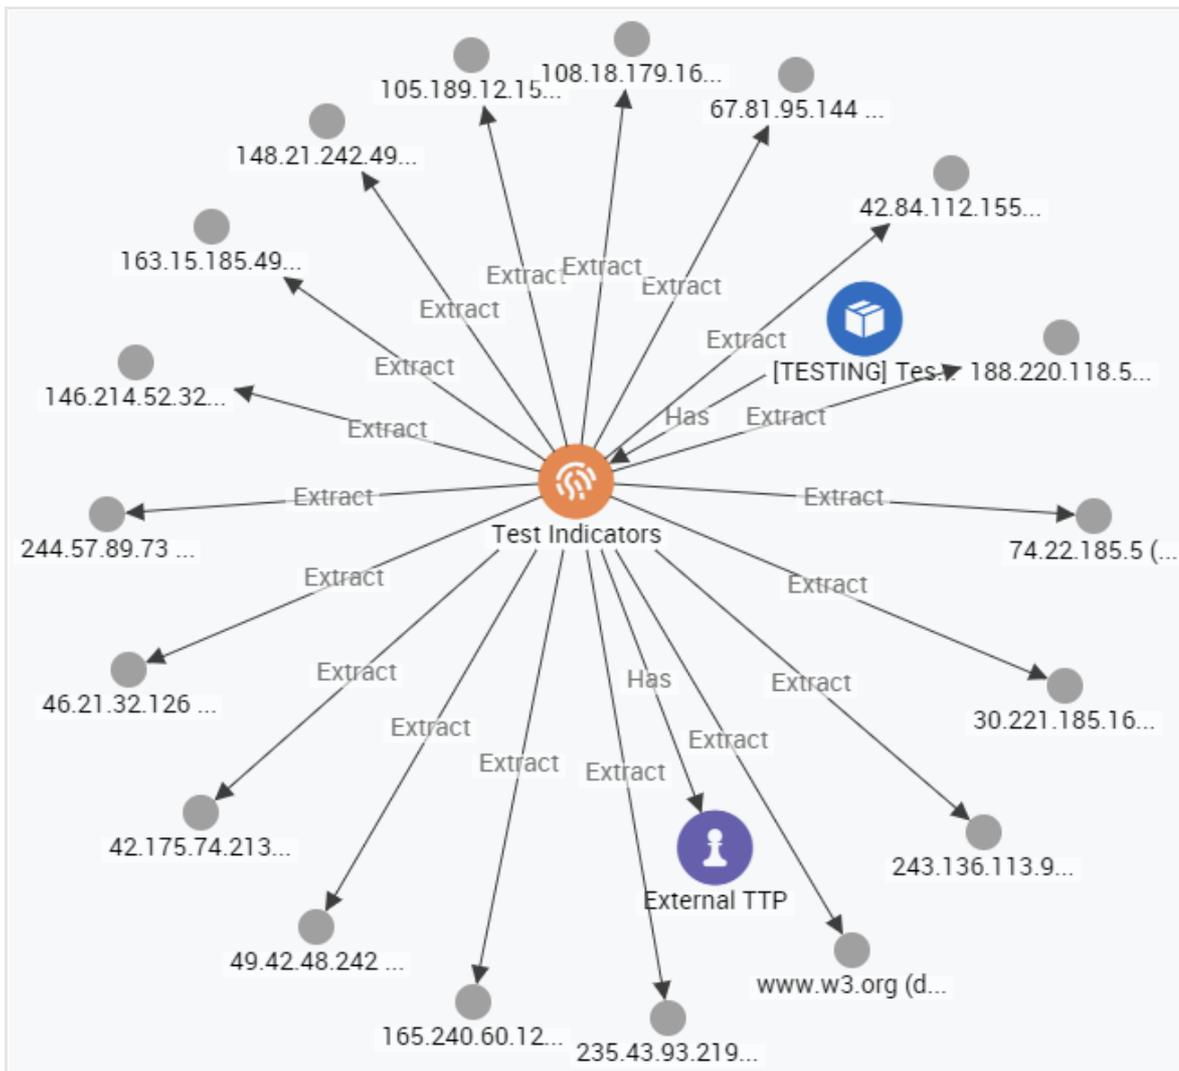


Look for and find relationships among entities and observables loaded on the graph

- **Load entities:** loads entities that are *directly related* to the selected one; for example, because they share common information such as one or more IP addresses, or target email addresses, and so on. From the **Load entities** submenu you can choose whether you want to load all related entities, or only specific related entity types.



- **Load observables:** loads any observables related to the selected entity; for example, an IP address, or an email address related to the entity. From the **Load observables** submenu you can choose whether you want to load all related observables, or only specific related observable types.



- Load entities by observable** : loads entities that are *indirectly related* to the selected one through one intermediate observable node.
 From the **Load entities by observable** submenu you can choose whether you want to load indirectly related entities based on all the available observables, or only on specific observable types.

i The maximum amount of relationships a graph query returns is capped at max. 500. When a query response returns more than 500 results, a dialog informs you about the limit, and it allows you to either cancel the query, or continue running it with a new cap at 5000 results in total.
 If you query multiple nodes, any nodes returning more than 500 relationships are automatically filtered out from the results.

There are too many results



The action returned more than 500 relations. If you load a large quantity of relations onto the graph, they are likely to clutter the view, while not adding value to the analysis.

If you selected more than one node, and if at least one of the selected nodes returns more than 500 results, you can run the query again and set the limit to a total of 5000 results. If a node returns less than 500 results, the items can be loaded onto the graph up to a maximum of 5000. Nodes returning more than 500 results are ignored. This may take a while to complete, as in: go grab a coffee in the meantime.

What would you like to do?

- Cancel the query and go back to the graph.
- Run the query again and limit it to 5000 results.

CANCEL THE QUERY

RUN THE QUERY

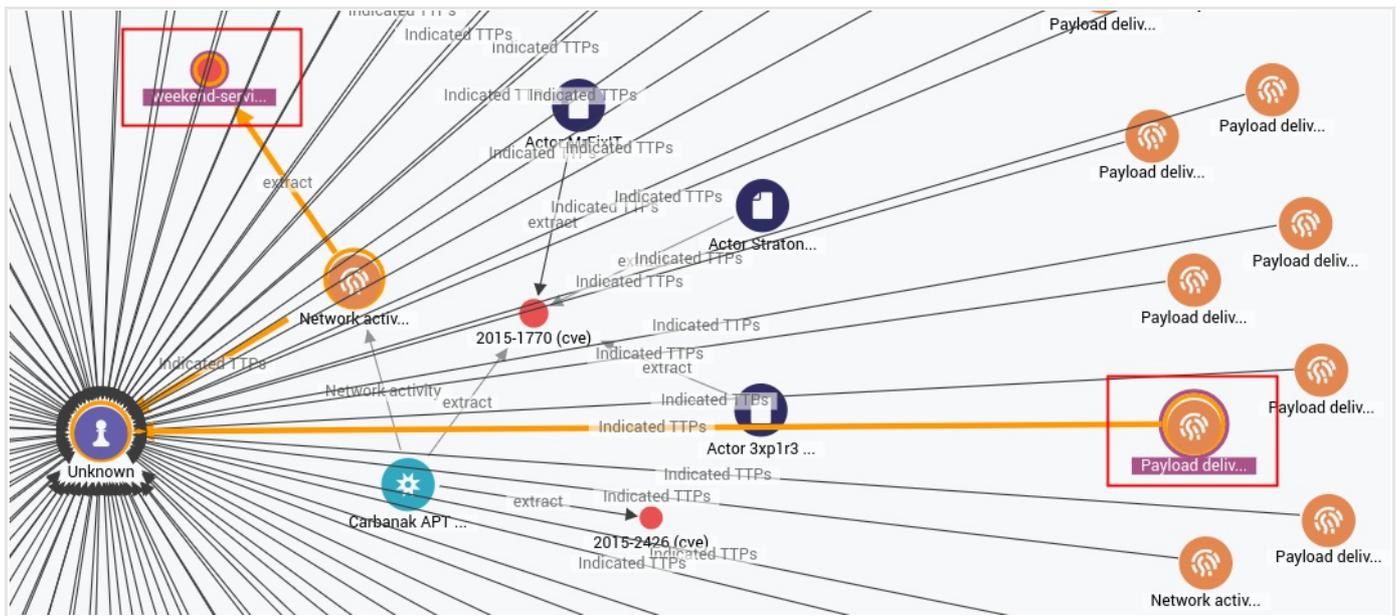
Enrich entities and observables on the fly

- **Enrichment:** manually triggers enricher tasks on the selected entities and observables. Any new enrichment observables are automatically loaded on the graph.
The **Enrichment** submenu enables you to apply granular control to the scope of the enrichment. Click one of the available options to:
 - Enrich all the selected entities and observables with all applicable enrichers
 - Enrich only the selected entities with all applicable enrichers
 - Enrich only the selected observables with all applicable enrichers
 - Apply a specific enricher from the list to all the selected entities and observables.

Explore connections between entities and observables

To see how entities, observables and enrichment observables are connected, and to inspect relationships between distant items, do the following:

- **CTRL** + click two nodes on the graph to select them.
- Right-click either selected node, and from the context menu select **Find path** to query the graph database about the existence of a path between the nodes, or **Show path** to highlight an existing path on the graph.
- If a path does exist, the selected nodes and all the intermediate ones are highlighted on the graph to show the path that links them.
 - **Find path:** queries the graph server to ask if there is a connection between the two selected nodes on the graph. If a connection does exist, the command loads any intermediate nodes, and then it highlights the connecting path. It differs from **Show path** because it first checks the existence of the path in the graph database.
 - **Show path:** highlights the shortest relationship path linking two nodes loaded on the graph. It differs from **Find path** because it does not check the existence of a path; it simply highlights the shortest path, if it exists on the graph.



Include selected entities in a workflow

- Add to dataset:** enables you to add the selected entity or entities to an existing dataset, which you can choose from a drop-down list. You can optionally assign the selection to an existing workspace. This option applies to entities only, it does not apply to observables.
- Create task:** allows you to create an actionable task related to the selected entities, which you can assign to a user, and to one or more stakeholders.

Select, group, ungroup, and remove entities and observables

- Group:** groups together entities and observables that you want to handle together.
 - Select the entities and/or the observables by dragging the mouse on the graph to highlight the area where the nodes are positioned.
 - Right-click any node in the selection, and then click **Group**.
- Ungroup:** undoes entity/observable grouping by restoring them as separate nodes on the graph.
- Remove from graph:** removes the selected entities and/or observables from the graph. It works on single, as well as multiple selections.

Add entities to an open graph

When you are working on an open graph, you can add and remove entities on the fly without leaving the graph canvas. Click the **Add from search**, **Add from discovery**, or **Add from workspaces** buttons on the graph top navigation bar to load more entities on the graph:



When you add entities from the search and the discovery services, as well as from one or more workspaces, you can apply filters to target specific entity types, data sources, and time ranges:

- Entity type:** select one or more checkboxes to select and load on the graph only the specified entity types.

- **Source:** select one or more checkboxes to select and load on the graph only entities belonging to the specified data sources.
- **Date:** select a time interval to select and load on the graph only entities ingested between the specified start and end dates.
- **Dataset:** select one or more checkboxes to select and load on the graph only entities belonging to the specified datasets.
The **Dataset** filter is not available when the results do not include any entities belonging to at least one dataset.

Add entities from search

You can add entities from the default search page. You can also run platform-wide search queries to look for specific entities to load, as well as filter search results to further narrow down your scope.

To open the search pane from the graph, do the following:

- Click **Q Add from search**.
- Enter search terms, **Elasticsearch queries** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>), or use the drop-down search filters to search for the entities you want to load on the graph.
The search or the filters return all the matches meeting your search/filtering criteria.
- Select the checkboxes corresponding to the entities you want to add to the graph.
- Click **Add to graph**.

The screenshot shows a search interface for a graph. At the top, there's a search bar with a magnifying glass icon and the text 'Filter...'. Below the search bar, there are three filter dropdowns: 'Entity types', 'Source', and 'Date'. To the right of these filters, it says '57064 results'. The main area contains a table of search results. Each row has a checkbox on the left, followed by a fingerprint icon and the title of the entity. The table has columns for 'Title', 'Source', and 'Ingestion time'. At the bottom left, it shows '1 - 6 of 57 064' and navigation arrows. At the bottom right, there is a teal button labeled 'ADD TO GRAPH'.

<input type="checkbox"/>	Title	Source	Ingestion time
<input type="checkbox"/>	untitled	Testing Group	24.07.2017 11:09
<input type="checkbox"/>	untitled	Testing Group	24.07.2017 11:09
<input type="checkbox"/>	untitled	Testing Group	24.07.2017 11:09
<input type="checkbox"/>	untitled	Testing Group	24.07.2017 18:01
<input type="checkbox"/>	!! E un fapt bine stabilit că cititorul va fi sustr...	Testing Group	02.08.2017 12:46
<input type="checkbox"/>	sid:2405021 "ET CNC Shadowserver Report...	INCTENRICH_System_Default	02.08.2017 14:58

Add entities from the discovery service

To limit your scope to discovered entities only, do the following:

- Click  **Add from discovery**.
- Enter search terms, **Elasticsearch queries** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl.html>), or use the drop-down search filters to search for the entities you want to load on the graph. The search or the filters return all the matches meeting your search/filtering criteria.
- Select the checkboxes corresponding to the entities you want to add to the graph.
- Click **Add to graph**.

Jnnamed graph*











x Default Workspace
x ▼

Filters: Entity types ▼ Source ▼ Date ▼ Datasets ▼ 101 results

<input type="checkbox"/>	Title	Source	Ingestion time	
<input checked="" type="checkbox"/>	 This domain tor.globenet.org has been identified as a TOR ...	tor	22.08.2017 17:15	
<input type="checkbox"/>	 This domain torex11.duffsdevice.com has been identified a...	tor	22.08.2017 17:14	
<input type="checkbox"/>	 This domain cassel.dogeneral.net has been identified as a ...	tor	22.08.2017 17:14	
<input checked="" type="checkbox"/>	 This domain sing-tor.cryptoligarch.com has been identified ...	tor	22.08.2017 17:18	
<input checked="" type="checkbox"/>	 This domain 55-21-50-84.dyn.estpak.ee has been identified ...	tor	22.08.2017 17:14	
<input type="checkbox"/>	 This ipAddress 202.85.233.34 has been identified as a TOR ...	tor	22.08.2017 17:18	

1 - 6 of 101





ADD TO GRAPH

Add entities from workspaces

To load on the graph entities belonging to one or more specific workspaces, do the following:

- Click  **Add from workspaces**.
- From the drop-down menu select one or more workspaces you want to load entities from.



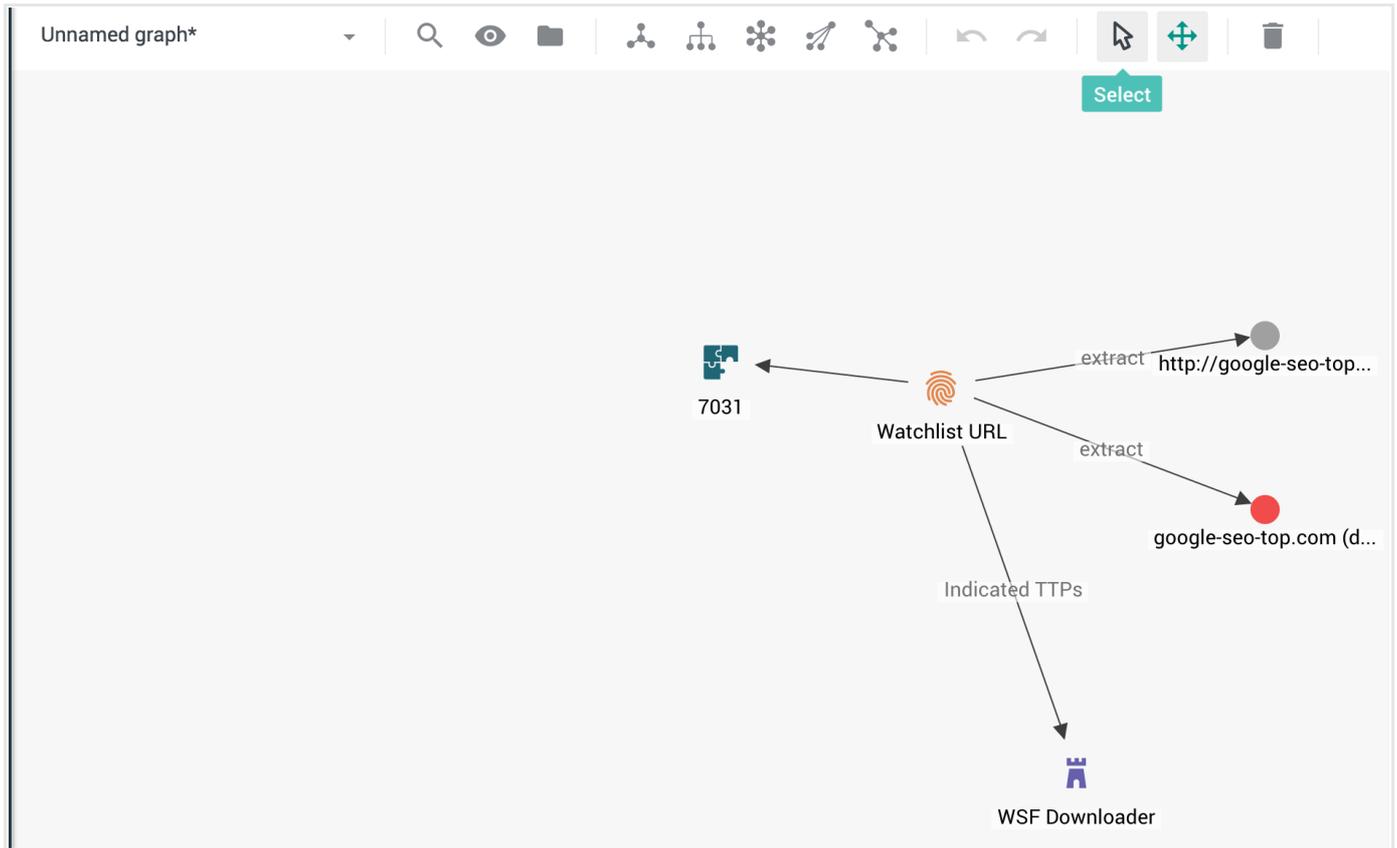
Layout type	Description
Standard	In the standard layout, links on the graph are a consistent length. Nodes and edges overlap as little as possible, and they are evenly distributed on the graph surface. It offers a consistent and clean view. It is a good starting point to begin analyzing any kind of data and any dataset size, especially when you are looking for patterns and symmetries.
Hierarchy	It is a tree structure with nodes. It displays child nodes horizontally below the corresponding parents. Connections flow top-down through the chart from the original subject. It is an efficient layout to visualize workflows and processes, impact analysis, and hierarchical relationships.
Radial	The radial layout arranges nodes in concentric circles around the original subject in a radial tree. Each set of nodes becomes a new orbit extending outwards from the original parent. This layout works best with networks with a large volume of child nodes to each parent.
Structural	It is similar to the standard layout. However, in the structural layout nodes with similar attributes are grouped together in fans. This visualization provides a clear overview of the clusters within a network, without focusing on a specific one.
Tweak	The tweak layout shows how networks evolve. The layout automatically adapts as links are created and destroyed, so that you can see where and how the changes occur. It is ideal for visualizing the behavior of dynamic and changing graphs.

Move around on the graph

You can move around on the graph canvas to zero in on a specific detail or to get an overall view of the entities and their relationships. You can select multiple entities and observables; for example, to group them together; deselect them, load new entities on the graph, and remove them.

You can toggle between cursor behaviors to move and select objects on the graph canvas:

- Click **Select** to select and deselect entities and observables on the graph. You can group, ungroup, and remove the selected entities, as well as further examine them using the context menu options.
- Click **Pan** to move up and down, as well as left and right on the graph canvas. This is helpful when working on a complex graph with a large amount of nodes.
- Use the mouse wheel to zoom in and out, for example to focus on a specific entity, and then to go back to the overall graph view.
- Click **Clear canvas** to remove everything from the graph and to start from a completely clean sheet. Confirm the action on the confirmation dialog to reset the graph to an empty canvas.



Undo and redo

On the graph you can recover from small slips and *oops!*:

- To annul the effect of an action, on the top navigation bar click **Undo**.
You can undo the 5 previous actions.
- To cancel the effect of an **Undo** action, on the top navigation bar click **Redo**.
This reapplies the effects of the previously undone action.
You can redo the 5 previous actions.



Filter entities with the histogram

When you analyze entities and observables on the graph canvas to explore relationships and to, almost literally, join the dots you may want to apply quick filters to the elements on the graph without having to move them around or temporarily remove them.

The histogram helps you filter and visually isolate specific subsets of the elements on the graph, based on shared/common properties and attributes.

To open the histogram pane, click the **Histogram** icon on the top navigation bar:



You can select one or more options by clicking the corresponding checkbox:

- Select a checkbox to display nodes with the corresponding property or attribute.
- Deselect a checkbox to hide nodes with the corresponding property or attribute.
- By default, all checkboxes are selected, that is, nothing is filtered out, and all the nodes and the relationships loaded on the graph are visible.

The histogram pane makes available many ready-to-use filters. You can stack and combine filters as you need.

- **Show singletons**: select this checkbox to view singleton nodes. They are isolated nodes with no relationships to any other nodes.
- **Entity type**: select one or more options in this category to view specific entity types.
 - **Multi-type-group**: select this checkbox to view grouped entities containing mixed entity types.
- **Observable type**: select one or more options in this category to view specific observable types.
- **Source**: select one or more options in this category to view entities and observables ingested from specific data sources, that is, incoming feeds and enrichers.
 - **Missing source**: select this checkbox to view entities and observables that are not associated with any data source.
- **TLP**: select one or more options in this category to view entities flagged with the specified TLP color codes. For example, you can use this filter to include in the resulting graph view only entities flagged as reserved, or that require immediate action.
 - **Missing TLP**: select this checkbox to view entities with no TLP flag.
- **Source reliability**: select one or more options in this category to view entities and observables flagged with the specified source reliability value. For example, you can use this filter to include in the resulting graph view only entities and observables originating from trustworthy data sources.
 - **Missing source reliability**: select this checkbox to view entities and observables that are not associated with any data source.
- **Confidence**: select one or more options in this category to view entities and observables flagged with the specified level of confidence; it flags the **estimated level of confidence** (https://docs.oasis-open.org/cti/stix/v1.2.1/csprd01/part14-vocabularies/stix-v1.2.1-csprd01-part14-vocabularies.html#_toc440440605) to assess the accuracy and trustworthiness of the entity information.
 - **Missing confidence**: select this checkbox to view entities whose confidence level is not set.
- **Observable classification**: select one or more options in this category to view observables flagged with the specified level of maliciousness. For example, you can use this filter to include in the resulting graph view only observable flagged as **Bad**.
 - **Missing observable classification**: select this checkbox to view entities and observables whose maliciousness confidence level is not set.
 - **Bad**: select this checkbox to view observables whose maliciousness confidence level is set to **Malicious (High confidence)**, **Malicious (Medium confidence)**, or **Malicious (Low confidence)**.
 - **Good**: select this checkbox to view observables marked as **Safe**.

- **Tags:** select one or more options in this category to view entities flagged with the specified tags. For example, you can use this filter to include in the resulting graph view only entities with specific Admiralty codes or kill chain values.
 - **Without tags:** select this checkbox to view untagged entities.

Filter entities with the timebar

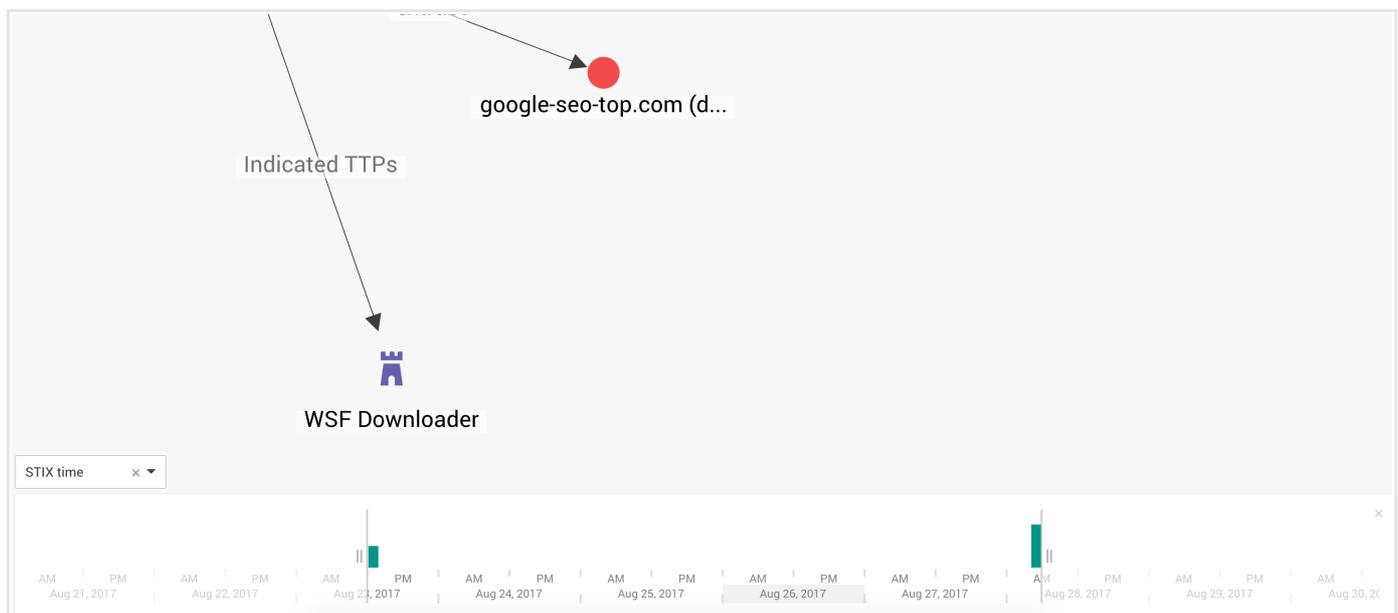
Besides filtering entities and observables based on specific properties and attributes, you can also filter by time range.

The graph timebar enables you to filter nodes on the graph based on a specified time interval; for example, to examine threat scenario evolution, or to focus on a specific phase in the development of the scenario under investigation.

To open the timebar, click the **Timebar** icon on the top navigation bar:



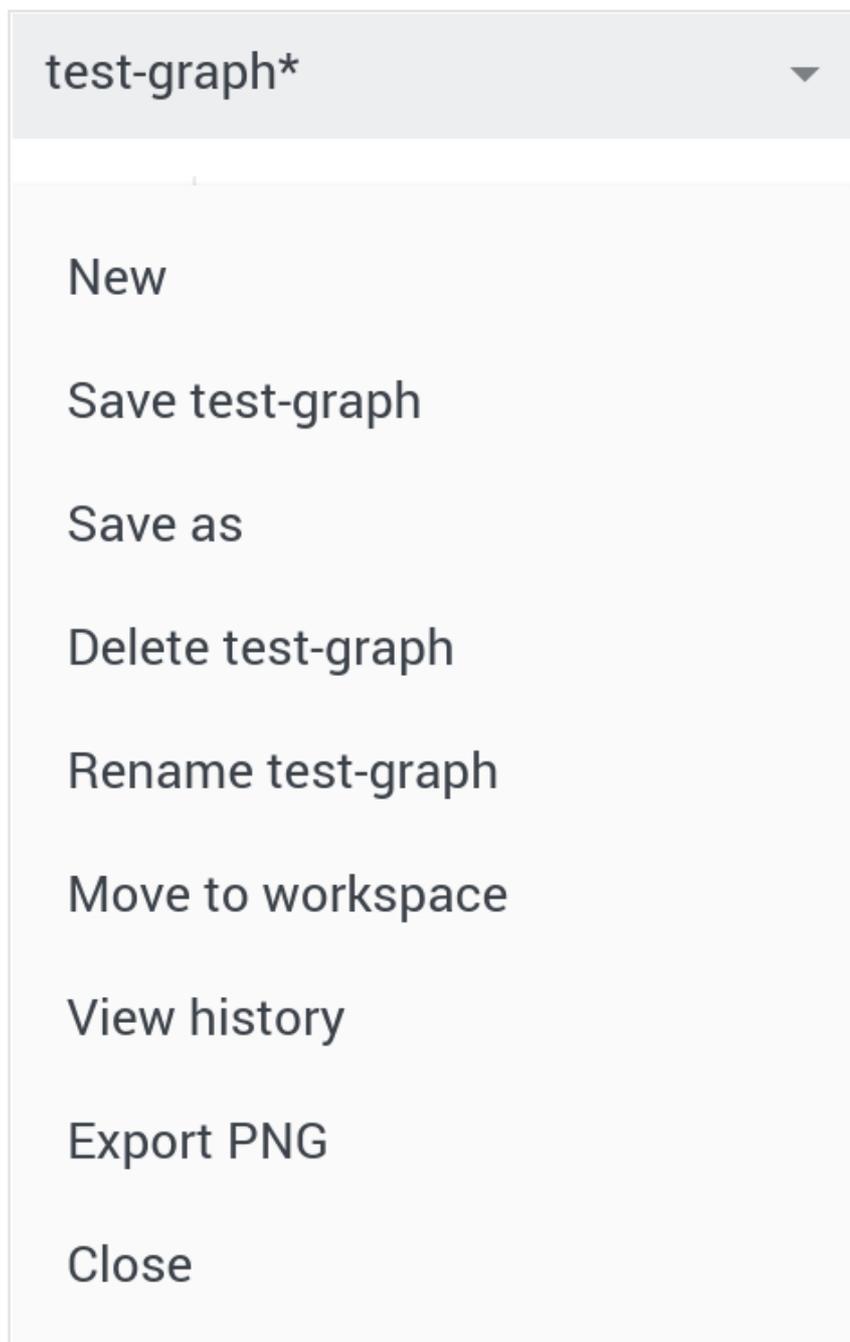
- From the drop-down menu select **STIX time** or **Ingestion time** to filter entities and observables on the graph based on either their STIX timestamp value (`data.timestamp` in the JSON representation of an entity), or their ingestion time into the platform (`created_at` in the JSON representation of an entity).
- Drag the sliders to set a time range to filter the nodes on the graph by.
- Click and drag the timeline on the timebar to the right or to the left to move forward or backward on the timeline, respectively.
- Use the mouse wheel to zoom in to focus on a shorter time range with a more granular level of detail; or to zoom out to get a more general overall view of the threats represented on the graph and their relationships.
- The bar chart on the timebar shows how the network grows over time. Double-click a column to zero in on it to view on the graph only entities and observables whose timestamp or ingestion date falls immediately around the time the selected column represents.



Save and export the graph

The graph drop-down menu enables you to save, rename, export, close, delete, and move the graph to a workspace. To display the drop-down graph menu, click the graph name or the downward-pointing ▼ arrow, and then select the option corresponding to the action you want to carry out.

An asterisk appended to the graph name indicates that the currently loaded data on the graph is not saved yet.



Besides saving, renaming, closing, and deleting the current graph, you can also do the following:

- **Move to workspace:** assigns the current graph to an existing workspace. On the dialog, from the drop-down menu select the destination workspace you want to assign the graph to. This enables you to add the graph to a workflow, so that you can access it also from the corresponding workspace. Moreover, based on the containing workspace access permissions, you can make the graph public — that is, any platform user can see it — or private — that is, only the workspace collaborators can see it.

- **View history:** displays the graph history pane with an overview in reverse chronological order of the actions performed on the graph and the corresponding users.
- **Export PNG:** exports the current graph as a *.png* file.
On the **Export graph as image** dialog you can set width and height in pixels of the exported image, as well as the graph area you want to export as image.
 - Click **Generate** to create the view to export.
 - Click **Download generated image** to export the view as a *.png* file.

Group entities in datasets

Datasets are generic containers that help you manage unordered data collections that do not need to be structured like data feeds.

About datasets

A dataset is an arbitrary and unordered data collection: you can edit and delete its content at any time.

A dataset is a generic container: you can create datasets to group entities for reference, for further analysis, to temporarily drop them and pick them up at a later time, and so on.

Why should you bother using datasets?

- They help you organize your intelligence. You can create datasets to group information based on any criteria that matter to you.

For example, you can create datasets to group entities based on:

- Entity type
- A specific threat scenario you are analyzing
- An incident
- A threat actor
- A targeted victim, and so on.

Subdividing a heterogeneous cyber threat intelligence corpus into smaller, more consistent, and more manageable chunks brings structure and clarity. This helps you identify relevant information more efficiently.

Static and dynamic datasets

A *static dataset* is a snapshot of the data selection it contains. You can modify it by adding and removing entities, but the dataset itself does not automatically evolve over time.

You can filter and search for entities by selecting one or more static datasets in the **Dataset** filter option available on most platform views.

A *dynamic dataset* is similar to a time-lapse video: a search query defines the scope, that is, the specific data you want to hold in the dataset. Every time you refresh (↻) the dataset view, the search runs, and any new matching data is added to the dataset.

Dynamic datasets help analyze the same data selection as it changes over time.

You cannot filter or search for entities by selecting dynamic datasets. Dynamic datasets are not included in the **Dataset** filter option available on most platform views.

Create a dataset

To create a dataset, do the following:

- On the top navigation bar click **Intelligence > All intelligence > Browse > Datasets** .
- On the dataset overview page, click the **+** icon.
- On the **Create dataset** page, enter a name for the dataset under **Dataset name**.
- Click **Save** to store your changes, or **Cancel** to discard them.

By default, new datasets are static.

- To create a dynamic dataset, click the **Dynamic** checkbox, and then specify a valid query string under **Search query**.
 - You can define the search query using the **Elasticsearch query syntax** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html>).
 - To point to a specific field in the entity JSON structure, use JSON path. This defines the target location for the search query.
 - The JSON path format is a string where dots (.) define JSON parent-child relationships.
 - Do not include square brackets ([]) in the path input: they are stripped during execution. It is not possible to use square brackets to point to specific array members.
 - In the specified location, you can look for literal values or for regex patterns.

Examples:

```
// Searches indicators for any of the following observables: IP addresses, or domain names, or URIs, or MD5 hashes

(extracts.kind:ipv4 or extracts.kind:domain or extracts.kind:uri or extracts.kind:hash-md5 ) AND types:("indicator")

// Searches for any observables containing the 'malware.win32.sample' value

extracts.value:malware.win32.sample

// Searches for any entities tagged exactly with 'Money Mule'

tags:"Money Mule"

// Searches for any entities whose original data source is 'Intel471'

meta.source_name:Intel471
```

Dataset name	Dynamic
0101	No
11617	No
8 Aug 2017	Yes
archi-reports	No
[DAILY DIGEST] Automated	Yes
[DAILY DIGEST] Manual	Yes
digest_report	No
fdsfds	No
hello	No
hh	No

Save options

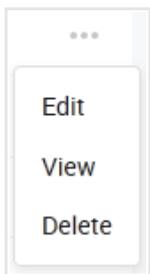
Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new**: saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate**: saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit a dataset

To edit an existing dataset, do the following:

- On the top navigation bar click **Intelligence > All intelligence > Browse > Datasets** .
- On the dataset overview page, click the click the  icon on the row corresponding to the dataset you want to modify.



- From the drop-down menu select **View** to inspect the content of the dataset. It shows an overview where you can review the dataset entities and edit them, if necessary.
- From the drop-down menu select **Edit** to rename an existing dataset or to enable/disable the static/dynamic property by selecting/deselecting the **Dynamic** checkbox.
- Click **Save** to store your changes, or **Cancel** to discard them.

Delete a dataset

To delete an existing dataset, do the following:

- On the top navigation bar click **Intelligence > All intelligence > Browse > Datasets** .
- On the dataset overview page, click the click the  icon on the row corresponding to the dataset you want to delete.
- From the drop-down menu select **Delete**.
- On the confirmation dialog, click **Delete** to confirm the action.

The dataset is deleted.

Deleting a dataset does not affect the entities that belong to it.

Add an entity to a dataset

You can manually add an entity to a dataset from almost anywhere in the platform.

To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse**, **Production**, **Discovery** or **Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.

- On the active view, browse to the entity you want to add to a dataset.
- Click the **i** icon corresponding to the entity you want to add to a dataset.
- From the drop-down menu select **Add to dataset**.

Alternatively:

- On the active view, browse to the entity you want to add to a dataset.
- Click anywhere on the row corresponding to the entity you want to add to a dataset. The entity detail pane slides in from the side of the screen.
- On the bottom half of the detail pane, click **Actions**.
- From the pop-up menu select **Add to dataset**.



You can add to datasets only published entities.
You cannot add draft entities to a dataset.

- On the **Add entity to a dataset** dialog, click the **Dataset** drop-down menu to select an existing dataset to add the entity to.
- To add the entity to a new dataset, which you can create on the fly, select **Create new dataset**.
- In this case, the **New dataset name** input field is displayed: enter a name for the new dataset.
- Optionally, you can assign the entity to a workspace by clicking the **Workspace** drop-down menu, and then by selecting an existing destination workspace or by creating a new one on the fly.
- Click **Add to dataset** to add the entity to the selected dataset and to close the dialog.

Add multiple entities to a dataset

- On the active view, select the checkboxes corresponding to the entities you want to add in bulk to a dataset.
- Click the **Add to** menu on the top-right-corner of the table view, above the table header row.
- From the drop-down menu select **Dataset**.



You can add to datasets only published entities.
You cannot add draft entities to a dataset.

- On the **Add entities to a dataset** dialog, click the **Dataset** drop-down menu to select an existing dataset to add the entities to.
- To add the entities to a new dataset, which you can create on the fly, select **Create new dataset**.
- In this case, the **New dataset name** input field is displayed: enter a name for the new dataset.
- Optionally, you can assign the entities to a workspace by clicking the **Workspace** drop-down menu, and then by selecting an existing destination workspace or by creating a new one on the fly.
- Click **Add to dataset** to add the entities to the selected dataset and to close the dialog.

View entity details

The entity detail pane enables you to drill down into an entity to closely examine it. The Overview tab shows the core details of the entity.

About the entity detail pane

The entity detail pane is a structured container holding a detailed and exhaustive information overview of an entity. You can use the entity detail pane as a reference resource you can look up when you want to zero in on a specific entity to review its structure, any observables it contains or it is related to, any relationships with external entities and observables, where it comes from, and if you are leveraging its intelligence value or not.

Through the entity detail pane you can also edit the entity — for example, to update information — load it on the graph, as well as include it in a workflow by adding it to a dataset and/or to a workspace, or by creating a follow-up task.

Access the entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.
- On the active view, browse to the entity you want to inspect, and click it. The entity detail pane slides in from the side of the screen.

You can edit the entity on the fly or by selecting **Actions > Edit**.

Examine the entity overview

The default view on the entity detail pane is the **Overview** tab. It is divided in stacked areas that structure the available information for the entity:

- **TLP:** the TLP color code the entity is flagged with.
 - Click the **TLP** button to override the current value with a new one.
- **Title:** the name of the entity, as shown also on the detail pane header section.
- **Confidence:** it flags the **estimated level of confidence** (https://docs.oasis-open.org/cti/stix/v1.2.1/csprd01/part14-vocabularies/stix-v1.2.1-csprd01-part14-vocabularies.html#_toc440440605) to assess the accuracy and trustworthiness of the entity information.
- **Analysis:** it is a free-text input field to include non-structured information such as additional context, references, links, and so on.

- **Tags:** select one or more tags to flag the entity with.
Tags help you structure and categorize entities based on criteria like confidence and attack stage. Tags improve findability, and they represent quick reference pointers to place entities in a broader cyber threat context. You can select existing taxonomy tags from the drop-down list, as well as create tags on the fly by typing them in the input field.
You can manage tags and their parent-child relationships under **Taxonomy**.
 - Click a tag to display an overview listing all entities sharing the same tag.
 - To remove a tag from the input field, click the corresponding **✕** icon.
 - To completely clear the **Tags** field, click the **✕** icon on the right-hand side of the field.
- **Estimated time**
 - **Start time:** sets the estimated inception time of the threat activity, based on observation, reports and other intelligence.
If no start date is indicated, you can click the edit button for this field, select a start date, and save it.
 - **End time:** if the threat is no longer active, this field sets the estimated end time of the threat activity, based on observation, reports and other intelligence.
If no start date is indicated, you can click the edit button for this field, select a start date, and save it.
 - **Observed:** defines the point in time when the entity was first observed/detected.
If no start date is indicated, you can click the edit button for this field, select a start date, and save it.
 - **Half life:** *Half life* is a numeric value representing the amount of time required to decrease the initial threat intelligence value of a malicious entity by 50%.
In other words, it indicates how long it takes for a threat to cut its malicious potential by half.
This value affects relevancy.
 - **Half life relevancy:** *Relevancy* is a numerical value based on the current time and the estimated start time of the threat. You can use it to sort and filter entities. *0%* = low relevancy — *100%* = high relevancy. Its value is 100% when the current time (*now*) is included between the threat start and end times. Otherwise, its value is 0. If the estimated end time is not available, relevancy is calculated using the estimated start time and the half-life value.
This field or value is non-editable.
- **Source**
 - **Name:** the data source of the entity. It can refer to a single source, for example a specific incoming feed, or to more sources grouped together.
You can group sources by intel type, for example IP addresses and domains, locations like countries and cities, forums, and so on; or by source type, for example incoming feeds vs. enrichers.
You can configure group sources under **⚙ > User management > Groups > <group_name> > Overview > Allowed sources**.
 - **Type:** defines the source type, for example a feed or a group.
 - **Reliability:** a reliability flag serves as an indication to assess the level of accuracy and trustworthiness of the source the entity originates from.

■ Exposure

- **Exposed:** Exposed entities are ingested and processed. However, their intelligence value is not leveraged to produce follow-up actions.

For example, triggering a detection event in a malware detection application downstream in the system; or a prevention event such as creating a firewall rule; or a community event such as sending a notification message to inform other parties about the possible threat the entity represents.

The entities hold intelligence value that is not consumed.

- **Detection:** If the dot is gray, no follow-up action has been undertaken to respond to the possible threat described in the entity.

If the dot is green, the entity information is used to carry out a follow-up action. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.

- **Prevention:** If the dot is gray, no follow-up action has been undertaken to respond to the possible threat described in the entity.

If the dot is green, the entity information is used to carry out a follow-up action. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.

- **Community:** If the dot is gray, no follow-up action has been undertaken to respond to the possible threat described in the entity.

If the dot is green, the entity information is used to carry out a follow-up action. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.

- **Sighting:** if an entity has been sighted, it means that it has been detected in the system and the system is compromised.

- **Outgoing feeds:** if one or more outgoing feeds are configured for the platform, and if the selected entity is included in at least one of them, you can see here how you are relaying entity information.

- **Datasets:** if the entity belongs to one or more datasets, they are listed here.

- **Name:** the name of the dataset.

Click it to go to the dataset overview page, where you can view and interact with the dataset entities and observables.

- **Entities:** the total amount of entities in the dataset.

- **Workspaces:** if the entity belongs to one or more workspaces, they are listed here.

- **Name:** the name of the workspace.

Click it to go to the workspace overview page, where you can view and interact with the workspace contents.

- **Last changed:** indicates the last time a user modified the workspace.

- **Collaborator:** if you are a collaborator of a workspace in the list, the corresponding flag is **Yes**.

- **Tasks:** actionable user tasks associated with the entity are listed here.
You can create tasks and assign them to yourself or to other users to request follow-ups; for example, further investigation or a call to action.
 - **Name:** the name identifying the task.
 - **Status:** the workflow stage the task is in: **Open**, **In progress**, **Done**, or **Canceled**.
 - **Assigned to:** the designated platform user who should carry out the task.
 - **Due date:** the deadline for the task to be completed.
- **Direct link to entity:** click the direct link to the entity to copy it to the clipboard and share it with other team members or threat analysts.

Manage the entity

Click the **Actions** pop-up menu on the bottom half of the entity detail pane tab and select the desired option to manage the entity and act on it. You can:

- Edit it;
- Delete it;
- Add it to a dataset;
- Load it on the graph for analysis;
- Create a follow-up task for the entity;
- Export it as JSON or STIX;
- Download it in its original data format; for example, the original STIX package containing the entity.

View entity observables

The Observables tab provides an overview of any observables belonging to the entity. You can analyze them, set their maliciousness confidence level, as well as create new observables or indicators and sightings from existing observables.

About the entity detail pane

The entity detail pane is a structured container holding a detailed and exhaustive information overview of an entity. You can use the entity detail pane as a reference resource you can look up when you want to zero in on a specific entity to review its structure, any observables it contains or it is related to, any relationships with external entities and observables, where it comes from, and if you are leveraging its intelligence value or not.

Through the entity detail pane you can also edit the entity — for example, to update information — load it on the graph, as well as include it in a workflow by adding it to a dataset and/or to a workspace, or by creating a follow-up task.

Access the entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.
- On the active view, browse to the entity you want to inspect and click it. The entity detail pane slides in from the side of the screen.

Examine the entity observables

On the entity detail pane click the **Observables** tab to display an overview listing any observables related to the entity.

You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.

- **Kind:** specifies the observable data type.
- **Value:** shows the observable value.

- **Lv:** level — either **1** or **2** — indicates the origin of the observable data, and if its intelligence value should be used for prevention/detection (level **1**) or for analysis (level **2**).
 - **Level 1:** the observable value originates from an embedded CybOX object. It is closer to the entity it refers to, and possibly more relevant.
You can decide to leverage the observable for prevention and/or detection purposes; for example, by instrumenting a firewall downstream in the system to block it.
 - **Level 2:** the observable value originates from a STIX field. It is more generic, and possibly less relevant.
You can decide to further analyze the observable to assess its maliciousness and relevance; for example, by creating a task and by assigning it to an analyst.
- **Conn:** indicates the number of connections/links the observable has with other entities in the platform.
- **Origins:** specifies the data source(s) the observable value was retrieved from; for example, an incoming feed, an enricher, the parent entity the observable is embedded in (only if the corresponding level is **1**), or a user who manually created it.
 - Hover the mouse pointer over an origin to display a pop-up with extra information such as the name(s) of the original data source(s), and links to download the original blob(s) containing the observable data.
- **Created:** the creation date of the observable
- The colored dot indicates if the observable is safe or malicious, as well as the maliciousness confidence level:
 - *1 gray dot:* there is not enough information or evidence to assess whether the observable is safe or malicious.
 - *1 green dot:* the observable is safe.
 - *1 red dot:* the observable *might* be malicious (low confidence).
 - *2 red dots:* the observable *may/can* be malicious (medium confidence).
 - *3 red dots:* the observable *is* malicious (high confidence).
- To refresh the view, if necessary, click the refresh icon: .

Apply bulk actions

On the **Observables** tab you can apply actions to multiple observables at the same time:

- On the active view, select the checkboxes corresponding to the observables you want to process in bulk.
- Click the **Actions** menu on the top-left corner of the table view, above the table header row.
- From the drop-down menu select the action you want to run on the selected observables.
- If a confirmation dialog is displayed, confirm the action as appropriate to run it.

The screenshot shows a context menu for an indicator. The menu items are: Promote, Demote, Ignore observable, Create indicator, Create sighting, Add to graph, and Set maliciousness >. Below the menu is a table with the following structure:

Actions	Filters: Kind	Lv
<input type="checkbox"/> kind	value	
<input type="checkbox"/> domain	wintorius.gecko.rs	
<input checked="" type="checkbox"/> ipv4	82.117.205.78	
<input checked="" type="checkbox"/> address	8.8.8.8	
<input checked="" type="checkbox"/> uri	http://boo.com/lalala	

Promote and demote observables

You can change the level of importance of an observable related to one or more entities to match the findings of an analysis, so that the observable correctly reflects the information value it represents.

Observable promoting and demoting affects the type of relationship an observable has with an entity — either direct or indirect — and the observable level — either 1 or 2.

In turn, levels provide an indication of how you can leverage the observable intelligence value by either using it for prevention/detection (level 1), or analysis (level 2).

Action	Current lv.	New lv.	Current rel.	New rel.	Results
Promote	2	1	indirect	direct	+ relevance; use for prevention/detection
Demote	1	2	direct	indirect	- relevance; use for analysis

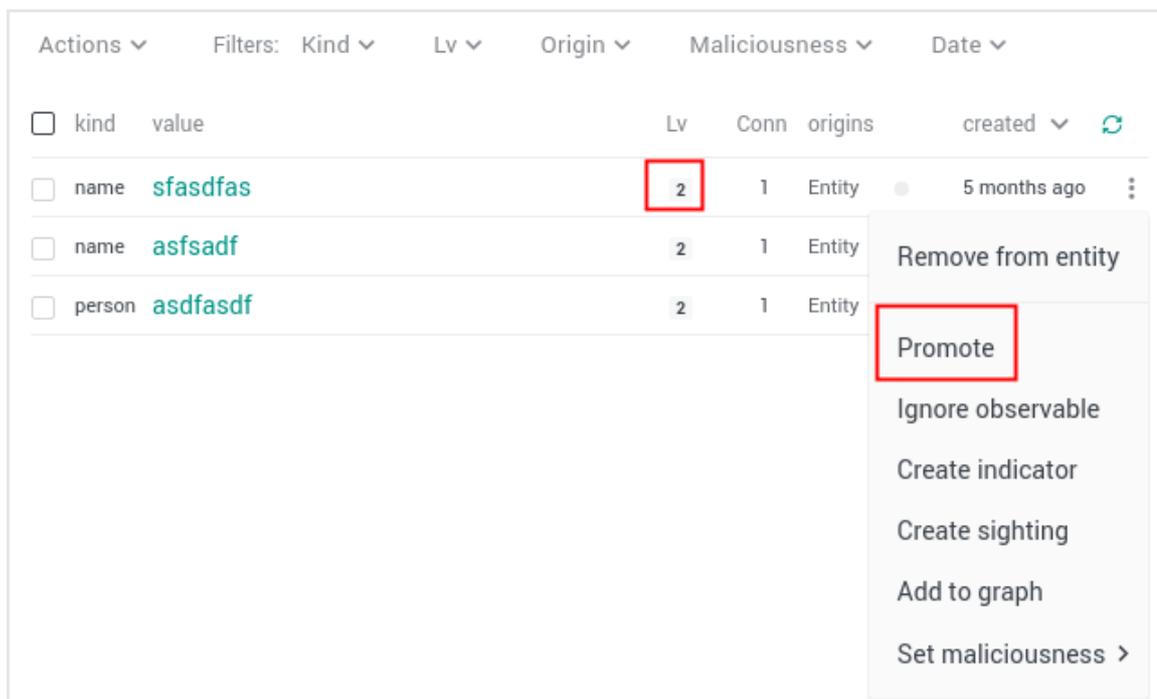
Promote an observable

To promote an observable from level 2 to level 1, do the following:

- On the row corresponding to the observable you want to promote, click the  icon.
- From the drop-down menu select **Promote**.
 - The platform creates a new version of the entity to reflect the entity-observable relationship change.
 - On the entity saving dialog click **Continue**.
 - The observable **Lv** value changes from **2** to **1**.

Alternatively:

- Select the checkbox corresponding to the observable you want to promote.
- From the **Actions** drop-down menu on the **Observables** tab click **Promote**.
 - The platform creates a new version of the entity to reflect the entity-observable relationship change.
 - On the entity saving dialog click **Continue**.
 - The observable **Lv** value changes from **2** to **1**.



The screenshot shows a table with columns: Actions, Filters: Kind, Lv, Origin, Maliciousness, Date. The table has three rows of data. The first row has a checkbox, 'name', 'sfasdfas', '2', '1', 'Entity', '5 months ago', and a three-dot menu icon. The context menu is open over the three-dot icon, showing options: Remove from entity, Promote, Ignore observable, Create indicator, Create sighting, Add to graph, and Set maliciousness >. The 'Promote' option is highlighted with a red box. The '2' in the 'Lv' column of the first row is also highlighted with a red box.

Actions	Filters: Kind	Lv	Origin	Maliciousness	Date
<input type="checkbox"/>	kind value	Lv	Conn	origins	created
<input type="checkbox"/>	name sfasdfas	2	1	Entity	5 months ago
<input type="checkbox"/>	name asfsadf	2	1	Entity	
<input type="checkbox"/>	person asdfasdf	2	1	Entity	

Promote an observable from level 2 to level 1

Saving Entity
✕

Create entity	Done.	✓
Create relations	Done.	✓

CONTINUE

The platform creates a new version of the entity and it rewires the entity-observable relationships

Actions		Filters: Kind		Lv	Origin	Maliciousness	Date
<input type="checkbox"/>	kind	value		Lv	Conn	origins	created
<input type="checkbox"/>	name	sfasdfas		1	1	User Created	5 months ago
<input type="checkbox"/>	name	asfsadf		2	1	Entity	5 months ago
<input type="checkbox"/>	person	asdfsadf		1	1	Direct	6 months ago

The new, promoted observable level is 1

Demote an observable

To demote an observable from level 1 to level 2, do the following:

- On the row corresponding to the observable you want to demote, click the icon.
- From the drop-down menu select **Demote**.
 - The platform creates a new version of the entity to reflect the entity-observable relationship change.
 - On the entity saving dialog click **Continue**.
 - The observable **Lv** value changes from **1** to **2**.

Alternatively:

- Select the checkbox corresponding to the observable you want to demote.
- From the **Actions** drop-down menu on the **Observables** tab click **Demote**.
 - The platform creates a new version of the entity to reflect the entity-observable relationship change.
 - On the entity saving dialog click **Continue**.
 - The observable **Lv** value changes from **2** to **1**.



You can also select multiple observables with mixed level values, and then promote or demote them by applying the desired bulk action.

Create an indicator from an observable

During an analysis, you may find out that an observable gains weight in the form of relevant contextual information that expands its intelligence value beyond the sheer statement of a fact such as an IP address, a hash value, or a threat actor's name. Therefore, you may want to consolidate, organize, and integrate this information in a consistent way; for example, by creating an indicator.

You can create an indicator from an observable in one of the following ways:

On the observable overview page

- On the top navigation bar click **Browse > Observables**.
- On the row corresponding to the observable you want to transform into a new indicator, click the  icon, and then select **Create indicator**.
The entity editor opens and you can proceed to enter the relevant details to create the indicator.

On the observable detail pane

- On the top navigation bar click **Browse > Observables**.
- Click anywhere on the row corresponding to the observable you want to transform into a new indicator.
The observable detail pane slides in from the side of the screen.
- On the bottom half of the detail pane, click **Actions**.
- From the pop-up menu select **Create indicator**.
The entity editor opens and you can proceed to enter the relevant details to create the indicator.

On the Observable tab on the entity detail pane

- Go to the entity detail pane of the entity related to the observable you want to transform into a new indicator.
- On the entity detail pane, go to the **Observables** tab.
- On the row corresponding to the observable you want to transform into a new indicator, click the  icon.
- From the drop-down menu select **Create indicator**.
The entity editor opens and you can proceed to enter the relevant details to create the indicator.

Alternatively:

- Select the checkbox corresponding to the observable you want to transform into a new indicator.
- Click the **Actions** drop-down menu on the **Observables** tab, and then select **Create indicator**.
The entity editor opens and you can proceed to enter the relevant details to create the indicator.

Create a sighting from an observable

When your organization records a discrete instance of an observable, for example, in a log file, the observable is sighted and the organization environment is compromised.

To represent this scenario in the platform, you can create a sighting from the sighted observable.

You can create a sighting from an observable in one of the following ways:

On the observable overview page

- On the top navigation bar click **Browse > Observables**.

- On the row corresponding to the observable you want to transform into a new sighting, click the  icon, and then select **Create sighting**.
The entity editor opens and you can proceed to enter the relevant details to create the sighting.

On the observable detail pane

- On the top navigation bar click **Browse > Observables**.
- Click anywhere on the row corresponding to the observable you want to transform into a new sighting.
The observable detail pane slides in from the side of the screen.
- On the bottom half of the detail pane, click **Actions**.
- From the pop-up menu select **Create sighting**.
The entity editor opens and you can proceed to enter the relevant details to create the sighting.

On the Observable tab on the entity detail pane

- Go to the entity detail pane of the entity related to the observable you want to transform into a new sighting.
- On the entity detail pane, go to the **Observables** tab.
- On the row corresponding to the observable you want to transform into a new sighting, click the  icon.
- From the drop-down menu select **Create sighting**.
The entity editor opens and you can proceed to enter the relevant details to create the sighting.

Alternatively:

- Select the checkbox corresponding to the observable you want to transform into a new sighting.
- Click the **Actions** drop-down menu on the **Observables** tab, and then select **Create sighting**.
The entity editor opens and you can proceed to enter the relevant details to create the sighting.

Add observables to the graph

On an observable view and on the entity detail pane **Observables** tab you can load one or more observables on the graph to analyze them and to examine their relationships.

To load an observable on the graph, do the following:

- On the row corresponding to the observable you want to load on the graph, click the  icon.
- From the drop-down menu select **Add to graph**.
- You can then proceed to open the graph, where you can start analyzing the observable.

Alternatively:

- Select the checkbox corresponding to the observable you want to load on the graph.
- From the **Actions** drop-down menu on the **Observables** tab click **Add to graph**.
- You can then proceed to open the graph, where you can start analyzing the observable.



You can also select multiple observables, and then load them all on the graph by clicking **Add to > Graph** on the horizontal bar above the column header row.

Set observable maliciousness

Gauging maliciousness helps you assess how dangerous an observable threat potential can be. In the platform you can set a confidence level to estimate the likelihood of an observable being malicious or not. The maliciousness values you can set help you answer the following question:

“Based on the factual evidence and the intelligence gathered so far, how likely is it for the observable to be malicious?”

Maliciousness confidence level	Represented as	Meaning
Unknown	● (gray)	It is not possible to assess whether the observable is malicious or not.
Safe	● (green)	The observable is not malicious.
Malicious (Low confidence)	● (red)	The observable might be malicious, but I am not sure.
Malicious (Medium confidence)	● ● (red)	I am confident to a point that the observable may be malicious.
Malicious (High confidence)	● ● ● (red)	I am confident that the observable is malicious.

Setting a maliciousness confidence level makes it easier to triage and to prioritize threat severity. You can set the maliciousness confidence level of an observable in one of the following ways:

On the observable overview page

- On the top navigation bar click **Browse > Observables**.
- On the row corresponding to the observable whose maliciousness confidence level you want to set, click the  icon, and then select **Set maliciousness**.
- From the sub-menu, click the maliciousness confidence level you want to assign to the observable.

On the observable detail pane

- Open the detail pane of the observable whose maliciousness confidence level you want to set.
- On the top half of the **Overview** tab under **Maliciousness** click **Edit**, and then select a maliciousness confidence level for the observable.
- Alternatively, on the bottom half of **Overview** tab, click **Actions > Set maliciousness**.
- From the sub-menu, click the maliciousness confidence level you want to assign to the observable.

On the Observable tab on the entity detail pane

- Go to the entity detail pane of the entity related to the observable whose maliciousness confidence level you want to set.
- On the entity detail pane, go to the **Observables** tab.
- On the row corresponding to the observable whose maliciousness confidence level you want to set, click the  icon.
- From the drop-down menu select **Set maliciousness**.
- From the sub-menu, click the maliciousness confidence level you want to assign to the observable.

Alternatively:

- Go to the entity detail pane of the entity related to the observable whose maliciousness confidence level you want to set.
- On the entity detail pane, go to the **Observables** tab.
- Select the checkbox corresponding to the observable whose maliciousness confidence level you want to set.
- From the **Actions** drop-down menu on the **Observables** tab click **Set maliciousness**.
- From the sub-menu, click the maliciousness confidence level you want to assign to the observable.



You can select multiple observables, and then you can assign the same maliciousness level to them clicking **Actions > Set maliciousness**.

Manually enrich observables

You can manually enrich observables by:

- Run all applicable enrichers for the entity to enrich all the observables it holds by selecting **Actions > Enrich > Enrich with all**:

Malicious files detected

Ingested: 06.10.2017 9:20 Incoming feed: TAXII Stand Samples TLP Not Set

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

Type	Value	Relation	Sighted	Conn.	First seen	Maliciousness
<input type="checkbox"/>	hash-sha256: e3b0c44298fc1c149afb4c899...	Related +1		2	06.10.2017 9:20	
<input type="checkbox"/>	hash-sha256: d7a8fbb307d7809469ca9abcb...	Related +1		1	06.10.2017 9:20	
<input type="checkbox"/>	file: readme.doc.exe	Related +1		1	06.10.2017 9:20	

- Edit
- Delete
- Add to dataset
- Add to graph
- Create task
- Export
- Download original
- Enrich
 - Enrich with all (5)
 - Censys Enricher
 - CrowdStrike Enricher
 - FireEye
 - Flashpoint AggregINT Enricher
 - Flashpoint Thresher Enricher

- Select and run a specific enricher on all the observables listed on the tab by clicking all observable checkboxes, and then **Enrich > <enricher_name>** on the horizontal bar above the column header row:

Malicious files detected

Ingested: 06.10.2017 9:20 Incoming feed: TAXII Stand Samples TLP Not Set

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

3 selected Deselect all

Type	Value	Relation	Sighted	Conn.	Business
<input checked="" type="checkbox"/>	hash-sha256: e3b0c44298fc1c149afb4c899...	Related +1		2	
<input checked="" type="checkbox"/>	hash-sha256: d7a8fbb307d7809469ca9abcb...	Related +1		1	
<input checked="" type="checkbox"/>	file: readme.doc.exe	Related +1		1	

Enrich dropdown menu options:

- Enrich with all (5)
- Censys Enricher
- CrowdStrike Enricher**
- FireEye
- Flashpoint AggregINT Enricher
- Flashpoint Thresher Enricher

- Select some observables by clicking the corresponding checkboxes, and then run all applicable enrichers by clicking **Enrich > Enrich with all** on the horizontal bar above the column header row:

Malicious files detected

Ingested: 06.10.2017 9:20 Incoming feed: TAXII Stand Samples TLP Not Set

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

2 selected Deselect all

Type	Value	Relation	Sighted	Conn.	Business
<input type="checkbox"/>	hash-sha256: e3b0c44298fc1c149afb4c899...	Related +1		2	
<input type="checkbox"/>	hash-sha256: d7a8fbb307d7809469ca9abcb...	Related +1		1	
<input checked="" type="checkbox"/>	file: readme.doc.exe	Related +1		1	

Enrich dropdown menu options:

- Enrich with all (5)**
- Censys Enricher
- CrowdStrike Enricher
- FireEye
- Flashpoint AggregINT Enricher
- Flashpoint Thresher Enricher

Manually add observables

You can add observables on the fly to associate them with the corresponding entity, that is, either the current entity displayed on the entity detail panel, or an entity you are creating or editing.

To manually add an observable, do the following:

- On the left-hand navigation sidebar click **+** > **Observable**
or:
- On the top navigation bar click **Intelligence > All intelligence > Production > Observables**
or:
- Click an entity on the **Browse, Production, Discovery, or Exposure** views.
- On the entity detail pane, click the **Observables** tab.
- On the active view, click **+** (**Create observable**) to begin creating an observable.

The observable editor opens, and you can start describing the new observable in the **Add observables** view. You can add as many observables to an entity as you need.

- **Value(s)**: enter the value of the observable. The value and its format should match the specified observable type (kind).
Insert one value entry per line.
If you enter multiple values on one line, use a comma (,) as a separator.
Example: *75.23.125.231, ipwnu.biz, Kansas City, 1.37bn@rivercitymedia.com, Alvin Slocombe*
- **Level**: from the drop-down menu select **1** or **2** to indicate if the observable is retrieved from an embedded CybOX object — level 1, closer to the entity it refers to, and possibly more relevant/urgent — or if it is retrieved from a STIX field — level 2, more generic, and possibly less relevant/less urgent.
- **Maliciousness**: from the drop-down menu select a maliciousness confidence level to assess the likelihood of the potential threat.

This option corresponds to the value you set under **Data configuration > Rules > Observable > + (Create rule) > Action > Mark as malicious > Confidence**.

When you flag an observable with a maliciousness confidence level, it cannot transition to *safe* or *ignore* anymore. It can only become more malicious, that is, it can only transition to a higher, never to a lower, maliciousness confidence level. Once an observable is flagged as harmful, it cannot go back to being safe or irrelevant anymore.

Link name labels vary based on the relationship the observable has with the specific entity type it belongs to.

If the entity type is *course of action*:

- **Parameter**: it is the only link name option available for entities.
It enables defining specific technical parameters, settings, and configurations related to the using the CybOX Language.

You can set parameters for a course of action to define automated courses of action designed to to carry out follow-up actions. It can be a detection follow-up; for example, it can trigger adjusting the settings of a malware detection application accordingly. It can be a prevention follow-up; for example, it can instrument a third-party system to block a range of malicious IP addresses or domain names. Or it can produce a community follow-up; for example, creating and publishing a report to notify other parties about the possible threat the entity represents.

If the entity type is *exploit target*:

- **Affected**: describes an affected, impacted resource.
- **Configuration**: enter the **Common Configuration Enumeration (CCE)**
(<https://nvd.nist.gov/config/cce/index>) code defining a specific security system configuration issue, as well as the related configuration guidance statement containing preferred or required settings or policies for the system configuration it refers to.
Example: *CCE-5770-3*

- **Vulnerability:** enter the **Common Vulnerabilities and exposures (CVE)** (<https://cve.mitre.org/cve/identifiers/>) **identifier** (https://en.wikipedia.org/wiki/common_vulnerabilities_and_exposures#cve_identifiers) to reference the security threat.
 Example: **CVE-2017-6394 on CVE** (<https://cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2017-6394>) or **CVE-2017-6394 on NVD** (<https://web.nvd.nist.gov/view/vuln/detail?vulnid=cve-2017-6394>).
- **Weakness:** enter the **Common Weakness Enumeration (CWE)** (<https://cwe.mitre.org/>) **identifier** (https://en.wikipedia.org/wiki/common_weakness_enumeration) to reference the software security weakness.
 Example: **CWE-319** (<http://cwe.mitre.org/data/definitions/319.html>), **CWE-642** (<http://cwe.mitre.org/data/definitions/642.html>).

If the entity type is *incident*:

- **Affected asset:** defines an affected, impacted resource or **asset type** (<https://stixproject.github.io/data-model/1.2/stixvocabs/assettypevocab-1.0/>).
- **Related:** holds one or more observables that are related to this one.

If the entity type is *indicator*:

- **Observable:** the observable related to the entity is an embedded CybOX observable object. It has been detected *outside* the organization.
- **Sighted:** the observable related to the entity is an embedded CybOX observable object. At least one specific occurrence of the observable related to the entity has been detected, that is, sighted, *inside* the organization.
- **Test mechanism:** a **test mechanism** (<https://stixproject.github.io/data-model/1.2/indicator/testmechanismtype/>) enables the platform to share entity information with external tools and systems. In particular, it is useful to send information to an **IDS/HIDS/NIDS** (https://en.wikipedia.org/wiki/intrusion_detection_system) to test it against a tool-specific rule.

For example, an observable with a **Test mechanism** link name can trigger follow-up actions in external systems:

- **Rule: generic test mechanism** (<https://stixproject.github.io/data-model/1.2/generictm/generictestmechanismtype/>) to interact with a generic system supporting plain text format as an input.
- **Snort: Snort test mechanism** (<https://stixproject.github.io/data-model/1.2/snorttm/snorttestmechanismtype/>).
 You can include the observable in an outgoing feed to a Snort instance. The Snort rules in the indicator are used to look for **matching patterns** (<https://stixproject.github.io/documentation/idioms/snort-test-mechanism/>) in the Snort logs. You can configure Snort so that matching hits trigger a follow-up action. For example, creating a sighting or adding a malicious entry to a blacklist.
- **YARA: YARA test mechanism** (<https://stixproject.github.io/data-model/1.2/yaratm/yaratestmechanismtype/>).
 You can include the observable in an outgoing feed to a YARA instance. YARA uses the rules in the indicator to look for **matching patterns** (<https://stixproject.github.io/documentation/idioms/yara-test-mechanism/>) in the target files or locations you specify in YARA.
 You can feed indicators from the platform to YARA to look for, identify, and classify malware samples.

If the entity type is *TTP*:

- **Malicious infrastructure:** describes a component of the infrastructure — gear, equipment, tools, software and hardware, services — used to carry out the malicious activities described in the TTP.
- **Targeted victim:** describes a component of the targeted victim's assets and resources.

If the entity type is *campaign*, *report*, *sighting* or *threat actor*:

- N/A. Campaign-related observables do not have link types.

- Click **Save** to store your changes, or **Cancel** to discard them.



You can use the specified observable values to set up automation processes, so that the (potential) threat the entity represents can trigger an action in a security system or another device in the toolchain.

For example, if the observable **Type** is *Email*, the **Link name** is *Parameter*, and the **Value(s)** are *scam@honestpaul-superdeals.com*, *info@honestpaul-superdeals.com*, and *support@honestpaul-superdeals.com*, create a rule in the email server to block all incoming messages from the *honestpaul-superdeals.com* email domain.

Manage the entity

Click the **Actions** pop-up menu on the bottom half of the entity detail pane tab and select the desired option to manage the entity and act on it. You can:

- Edit it;
- Delete it;
- Add it to a dataset;
- Load it on the graph for analysis;
- Create a follow-up task for the entity;
- Export it as JSON or STIX;
- Download it in its original data format; for example, the original STIX package containing the entity.

Inspect the neighborhood

The Neighborhood tab in the entity detail pane includes a small graph canvas showing the immediate context of the entity, as well as related observables, datasets, workspaces, and tasks.

About the entity detail pane

The entity detail pane is a structured container holding a detailed and exhaustive information overview of an entity. You can use the entity detail pane as a reference resource you can look up when you want to zero in on a specific entity to review its structure, any observables it contains or it is related to, any relationships with external entities and observables, where it comes from, and if you are leveraging its intelligence value or not.

Through the entity detail pane you can also edit the entity — for example, to update information — load it on the graph, as well as include it in a workflow by adding it to a dataset and/or to a workspace, or by creating a follow-up task.

Access the entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.
- On the active view, browse to the entity you want to inspect, and click it. The entity detail pane slides in from the side of the screen.

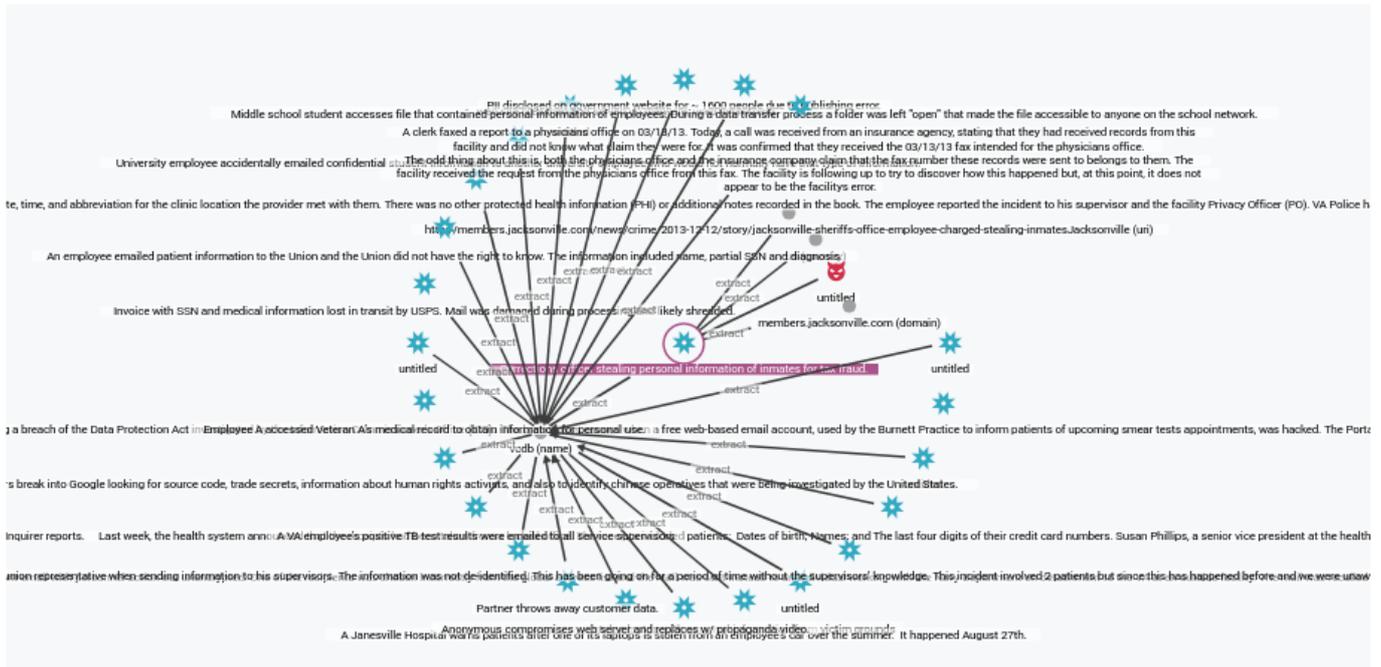
Explore the entity neighborhood

During an analysis you may want to quickly inspect an entity to check relationships with other entities and observables. Normally, you would load the selected entity onto the graph, open the graph, and proceed with the inspection.

Without leaving the entity detail pane, the **Neighborhood** tab offers a faster alternative: click it to see a small graph displaying close-range relationships the entity has with nearby entities and observables.

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

Graph



! Too many items to show, showing only most relevant 30 items. ✕

The embedded graph is a snapshot of the graph canvas view. The embedded snapshot is refreshed when accessing the **Neighborhood** tab, but it is not updated in real time. When the entity relationship landscape changes, for example, after adding or removing relationships, the embedded graph is not in sync anymore. In this case, a notification message is displayed to inform the user:

i DATA PROCESSING IN PROGRESS — It may take some time before the latest entity data is available in the graph.

OVERVIEW OBSERVABLES **NEIGHBORHOOD** JSON VERSIONS HISTORY

! DATA PROCESSING IN PROGRESS ✕

It may take some time before the latest entity data is available in the graph.

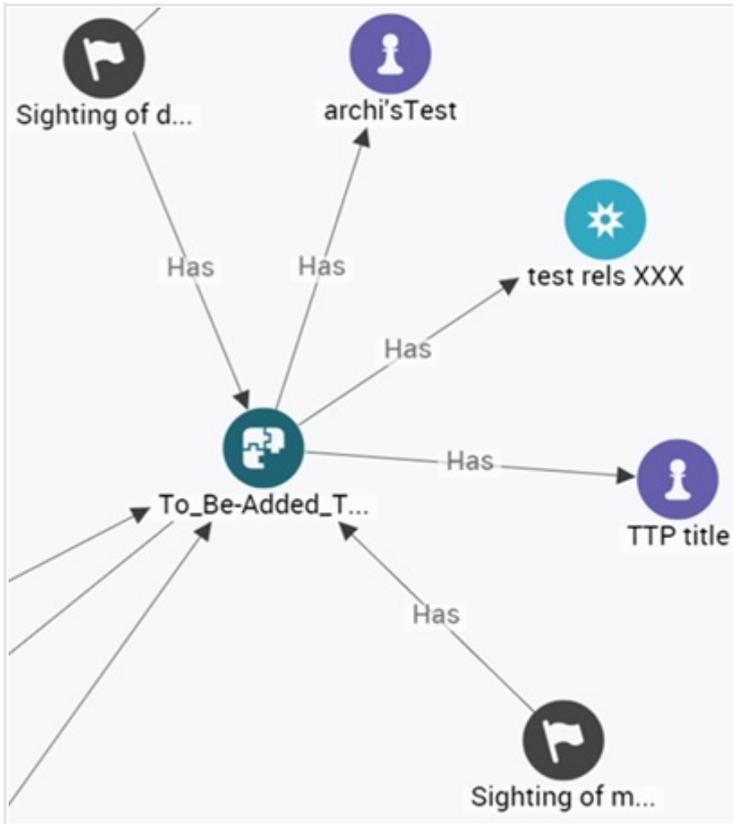
Graph

The graph displays a central node representing an incident associated with the Red October campaign, which is connected to numerous other nodes. These nodes include various URLs, domains, and technical details, all linked by relationships. The graph is rendered with a light blue background and a grid pattern.

After the platform graph completes indexing, the embedded graph is back in sync. The time this task requires varies, depending on the size of the graph queue.

View related entities and observables

On the graph view you can inspect any relationships the entity may have with other entities and observables in the platform. Relationships can be *direct* — the entities and/or observables are immediately related to each other — as well as *indirect* — the entities and/or observables are related through a shared entity or a shared observable.



Entities with direct relationships



Entities with indirect relationships

To visually examine the entity more closely, click the small graph to launch the larger and feature-richer graph.

Directly related entities

This section displays entities that are directly related to the active entity. You can see the entity the current entity is related to, the relationship type, and the relationship direction, that is, if it outgoing (from the current entity to the related one) or incoming (from the related entity to the current one).

Click an entity name to display the corresponding detail pane in full page format. To edit entity relationships, click **Edit relationships**.

! Too many items to show, showing only most relevant 30 items.

Directly related entities

Relationship type	Related entity
indicated_ttps	→ 🚩 Fake TTP d0412409924b #72

EDIT RELATIONSHIPS

Entities related through observables

This section displays entities that are indirectly related to the active entity, that is, the relationship exists through an intermediate entity or observable.

Each entry reports entity name, entity TLP color code, if available, and entity ingestion time.

Click an entity name to display the corresponding detail pane in full page format.

Entities related through observables		
Type	TLP	Ingested
🚩 6796		24.08.2017 20:48
🚩 6795		24.08.2017 20:04
🚩 6793		24.08.2017 20:03
🚩 6788		24.08.2017 18:39
🚩 6794		24.08.2017 18:57
🚩 6791		24.08.2017 18:53
🚩 6792		24.08.2017 18:53
🚩 6790		24.08.2017 18:08
🚩 6789		24.08.2017 17:43
🚩 6787		24.08.2017 17:39

1 - 10 of 200
|< < > >|

Set relationships

On the **Neighborhood** tab you can update entity information by adding and removing entity relationships.

To do so, do the following:

- Under **Directly related entities** click **Edit relationships**.
- From the drop-down menu select the option corresponding to the relationship you want to create.
- On the **Search an entity** dialog, click the checkbox(es) to select one or more entities to relate them to the current one.



You can refine the displayed results by specifying a search string in the filter input field. Alternatively, click one of the available filter options to select and filter by specific:

- **Entity types**
- **Source**
- **Date**
- **Datasets**

- Click **Select**.
- From the **Source** drop-down menu, select a data source for the entity or entities you are relating to the current one. You can select only one data source at a time, regardless the number of entities you choose on the **Search an entity** dialog.
- Click **Save** to store your changes, or **Cancel** to discard them.
- To *remove* a relationship or a relationship type, click the **✕** icon on the row displaying the relationship or next to the relationship type you want to remove. The row and the corresponding relationship or the relationship type are removed. You cannot undo this action.

Set campaign relationships

Select this option...	... to create this relationship for the campaign
Associated Neighborhoods	Outgoing relationship — Relates the campaign to the selected campaign(s) on the Search an entity dialog.
Attributions	Outgoing relationship — Relates the campaign to the selected threat-actor(s) on the Search an entity dialog.
Related incidents	Outgoing relationship — Relates the campaign to the selected incident(s) on the Search an entity dialog.
Related TTPs	Outgoing relationship — Relates the campaign to the selected TTP(s) on the Search an entity dialog.

Select this option...	... to create this relationship for the campaign
Indicator → Related campaigns	Incoming relationship — Relates the selected indicator(s) on the Search an entity dialog to the campaign.
Report → Campaigns	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the campaign.
Threat actor → Associated campaigns	Incoming relationship — Relates the selected threat-actor(s) on the Search an entity dialog to the campaign.
Sighting → Campaign	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the campaign.

Set course of action relationships

Select this option...	... to create this relationship for the course of action
Related exploit targets	Outgoing relationship — Relates the course of action to the selected exploit target(s) on the Search an entity dialog.
Related incidents	Outgoing relationship — Relates the course of action to the selected incident(s) on the Search an entity dialog.
Related courses of action	Outgoing relationship — Relates the course of action to the selected course(s) of action on the Search an entity dialog.
Exploit target → Potential courses of action	Incoming relationship — Relates the selected exploit target(s) on the Search an entity dialog to the course of action.
Indicator → Suggested courses of action	Incoming relationship — Relates the selected indicator(s) on the Search an entity dialog to the course of action. Recommends carrying out a course of action to respond to an indicator.
Incident → Courses of action requested	Incoming relationship — Relates the selected indicator(s) on the Search an entity dialog to the course of action. Requests to carry out a course of action to respond to an incident.
Incident → Courses of action taken	Incoming relationship — Relates the selected indicator(s) on the Search an entity dialog to the course of action. Reports the course of action carried out as a response to an incident.
Report → Courses of action	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the course of action.
Sighting → Course of action	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the course of action.

Set exploit target relationships

Select this option...	... to create this relationship for the exploit target
Potential courses of action	Outgoing relationship — Relates the exploit target to the selected potential course(s) of action on the Search an entity dialog
Related exploit targets	Outgoing relationship — Relates the exploit target to the selected exploit target(s) on the Search an entity dialog
Course of action → Related exploit targets	Incoming relationship — Relates the selected course(s) of action on the Search an entity dialog to the exploit target.
Report → Exploit targets	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the exploit target.
TTP → Exploit targets	Incoming relationship — Relates the selected TTP(s) on the Search an entity dialog to the exploit target.
Sighting → Exploit target	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the exploit target.

Set incident relationships

Select this option...	... to create this relationship for the incident
Related indicators	Outgoing relationship — Relates the incident to the selected indicator(s) on the Search an entity dialog.
Leveraged TTPs	Outgoing relationship — Relates the incident to the selected TTP(s) on the Search an entity dialog.
Attributed threat actors	Outgoing relationship — Relates the incident to the selected threat-actor(s) on the Search an entity dialog.
Related incidents	Outgoing relationship — Relates the incident to the selected incident(s) on the Search an entity dialog.
Courses of action requested	Outgoing relationship — Relates the incident to the selected course(s) of action on the Search an entity dialog to respond to the incident.
Courses of action taken	Outgoing relationship — Relates the incident to the selected course(s) of action on the Search an entity dialog that are carried out as a response to the incident.
Campaign → Related incidents	Incoming relationship — Relates the selected campaign(s) on the Search an entity dialog to the incident.
Course of action → Related incidents	Incoming relationship — Relates the selected course(s) of action on the Search an entity dialog to the incident.
Report → Incidents	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the incident.
Sighting → Incident	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the incident.

Edit indicator relationships

Select this option...	... to create this relationship for the indicator
Indicated TTPs	Outgoing relationship — Relates the indicator to the selected TTPs(s) on the Search an entity dialog.
Suggested courses of action	Outgoing relationship — Relates the indicator to the selected course(s) of action on the Search an entity dialog. Recommends carrying out a course of action to respond to the indicator.
Related Neighborhoods	Outgoing relationship — Relates the indicator to the selected indicator(s) on the Search an entity dialog.
Related campaigns	Outgoing relationship — Relates the indicator to the selected campaign(s) on the Search an entity dialog.
Incident → Related indicators	Incoming relationship — Relates the selected incident(s) on the Search an entity dialog to the indicator.
Report → Indicators	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the indicator.
Sighting → Indicator	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the indicator.

Set report relationships

Select this option...	... to create this relationship for the report
Indicators	Outgoing relationship — Relates the report to the indicator(s) on the Search an entity dialog.
TTPs	Outgoing relationship — Relates the report to the selected TTP(s) on the Search an entity dialog. Recommends carrying out a course of action to respond to the report.
Exploit targets	Outgoing relationship — Relates the report to the selected exploit target(s) on the Search an entity dialog.
Incidents	Outgoing relationship — Relates the report to the selected incident(s) on the Search an entity dialog.
Courses of action	Outgoing relationship — Relates the report to the selected course(s) of action on the Search an entity dialog.
Campaigns	Outgoing relationship — Relates the report to the selected campaign(s) on the Search an entity dialog.
Threat actors	Outgoing relationship — Relates the report to the selected threat actor(s) on the Search an entity dialog.
Sighting → Neighborhood	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the report.

Set sighting relationships

Select this option...	... to create this relationship for the sighting
Campaign	Outgoing relationship — Relates the sighting to the selected campaign(s) on the Search an entity dialog.
Course of action	Outgoing relationship — Relates the sighting to the selected course(s) of action on the Search an entity dialog.
Exploit target	Outgoing relationship — Relates the sighting to the selected exploit target(s) on the Search an entity dialog.
Indicator	Outgoing relationship — Relates the sighting to the selected indicator(s) on the Search an entity dialog.
Incident	Outgoing relationship — Relates the sighting to the selected incident(s) on the Search an entity dialog.
Report	Outgoing relationship — Relates the sighting to the selected report(s) on the Search an entity dialog.
Threat actor	Outgoing relationship — Relates the sighting to the threat actor(s) on the Search an entity dialog.
TTP	Outgoing relationship — Relates the sighting to the selected TTP(s) on the Search an entity dialog.

Set threat actor relationships

Select this option...	... to create this relationship for the threat actor
Observed TTPs	Outgoing relationship — Relates the threat actor to the selected TTP(s) on the Search an entity dialog.
Associated campaigns	Outgoing relationship — Relates the threat actor to the selected campaign(s) on the Search an entity dialog.
Associated actors	Outgoing relationship — Relates the threat actor to the selected threat actor(s) on the Search an entity dialog.
Campaign → Attributions	Incoming relationship — Relates the selected campaign(s) on the Search an entity dialog to the threat actor.
Incident → Attributed threat actors	Incoming relationship — Relates the selected incident(s) on the Search an entity dialog to the threat actor.
Report → Threat actors	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the threat actor.

Select this option...	... to create this relationship for the threat actor
Sighting → Threat actor	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the threat actor.

Set TTP relationships

Select this option...	... to create this relationship for the TTP
Exploit targets	Outgoing relationship — Relates the TTP to the selected exploit target(s) on the Search an entity dialog.
Related TTPs	Outgoing relationship — Relates the TTP to the selected TTP(s) on the Search an entity dialog.
Campaign → Related TTPs	Incoming relationship — Relates the selected campaign(s) on the Search an entity dialog to the TTP.
Indicator → Indicated TTPs	Incoming relationship — Relates the selected indicator(s) on the Search an entity dialog to the TTP.
Incident → Leveraged TTPs	Incoming relationship — Relates the selected incident(s) on the Search an entity dialog to the TTP.
Report → TTPs	Incoming relationship — Relates the selected report(s) on the Search an entity dialog to the TTP.
Threat actor → Observed TTPs	Incoming relationship — Relates the selected threat actor(s) on the Search an entity dialog to the TTP.
Sighting → TTP	Incoming relationship — Relates the selected sighting(s) on the Search an entity dialog to the TTP.

View related datasets

Related datasets

Entities can belong to one, more, or no datasets. If the entity is included in one or more datasets, they are listed here. Each entry reports the total amount of entities the corresponding dataset contains.

When a dataset is related to an entity, it shares data with it. Datasets and entities can be related in the following ways:

- The entity is included in the dataset.
- The entity and the dataset share common observables.
- The dataset contains an entity that bears a direct or indirect relationship with the active entity displayed on the entity detail pane.

Click a dataset name to display the corresponding detail pane in full page format where you can modify and edit it, if necessary.

View related workspaces

Related workspaces

Entities can belong to one, more, or no workspaces. If the entity belongs to one or more workspaces, they are listed here.

Each entry reports the most recent workspace modification date/time, and whether or not you are a collaborator of the workspace.

Click a workspace name to display the corresponding detail pane in full page format where you can modify and edit it, if necessary.

View related tasks

Related tasks

Any actionable user tasks associated with the entity are listed here.

You can create tasks and assign them to yourself or to other users to request follow-ups; for example, further investigation or a call to action.

This overview lists any actions that have been requested, are in progress, or have been carried out as a response or a follow-up action to the entity information. It shows what is being done to leverage the entity intelligence value.

Each entry reports task name, task progress status, task assignee, and task deadline.

Click a task name to display the corresponding detail pane where you can modify and edit it, if necessary.

Manage the entity

Click the **Actions** pop-up menu on the bottom half of the entity detail pane tab and select the desired option to manage the entity and act on it. You can:

- Edit it;
- Delete it;
- Add it to a dataset;
- Load it on the graph for analysis;
- Create a follow-up task for the entity;
- Export it as JSON or STIX;
- Download it in its original data format; for example, the original STIX package containing the entity.

View the JSON structure

The JSON tab displays the entity JSON structure.

About the entity detail pane

The entity detail pane is a structured container holding a detailed and exhaustive information overview of an entity. You can use the entity detail pane as a reference resource you can look up when you want to zero in on a specific entity to review its structure, any observables it contains or it is related to, any relationships with external entities and observables, where it comes from, and if you are leveraging its intelligence value or not.

Through the entity detail pane you can also edit the entity — for example, to update information — load it on the graph, as well as include it in a workflow by adding it to a dataset and/or to a workspace, or by creating a follow-up task.

Access the entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.
- On the active view, browse to the entity you want to inspect, and click it. The entity detail pane slides in from the side of the screen.

Examine the entity JSON structure

This tab displays a reader-friendly representation the of entity as a JSON object. You can expand and compress the nodes to show or hide the corresponding data.

This view can be handy as a reference to quickly look up the value of a specific field without or before exporting the entity as JSON or STIX.

Manage the entity

Click the **Actions** pop-up menu on the bottom half of the entity detail pane tab and select the desired option to manage the entity and act on it. You can:

- Edit it;
- Delete it;

- Add it to a dataset;
- Load it on the graph for analysis;
- Create a follow-up task for the entity;
- Export it as JSON or STIX;
- Download it in its original data format; for example, the original STIX package containing the entity.

View entity versions

The Versions tab displays an overview in reverse chronological order of any alternative versions of the entity in the platform.

About the entity detail pane

The entity detail pane is a structured container holding a detailed and exhaustive information overview of an entity. You can use the entity detail pane as a reference resource you can look up when you want to zero in on a specific entity to review its structure, any observables it contains or it is related to, any relationships with external entities and observables, where it comes from, and if you are leveraging its intelligence value or not.

Through the entity detail pane you can also edit the entity — for example, to update information — load it on the graph, as well as include it in a workflow by adding it to a dataset and/or to a workspace, or by creating a follow-up task.

Access the entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse**, **Production**, **Discovery** or **Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.
- On the active view, browse to the entity you want to inspect, and click it. The entity detail pane slides in from the side of the screen.

View alternative entity versions

When an entity is updated, modified or edited, the platform creates a new version of the entity to reflect the changes. Over time, multiple versions of an entity can exist in the platform.

Click the **Versions** tab to display an overview in reverse chronological order of any existing versions of the entity, along with the corresponding data source — an incoming feed, an enricher or a group — a TLP color code, if available, and the creation date and time.

Click an entity version name to display the corresponding detail pane in full page format where you can modify and edit it, if necessary.

The most recent version of the entity is the active one. Older versions exist to document the history of the entity and to act as reference. Therefore, most actions you can normally carry out on an entity are disabled and grayed out on the **Actions** menu.

The platform notifies you when an entity version is obsolete:

OVERVIEW	OBSERVABLES	NEIGHBORHOOD	JSON	VERSIONS	HISTORY
Title		Source	TLP	Creation time	
	Test-TTP-new  outdated	Testing Group	 Amber	Yesterday at 2:09 PM	
	Test-TTP-new  outdated	Testing Group	 Amber	Yesterday at 2:07 PM	
	Test-TTP-new  outdated	Testing Group		Yesterday at 1:41 PM	
	Test-TTP  outdated	Testing Group		Yesterday at 1:40 PM	
	Test-TTP  outdated	Testing Group		Yesterday at 1:28 PM	

If you click an outdated entity version name, the corresponding entity detail pane displays in full page format. On the **Actions** menu most actions are grayed out (disabled). A warning message notifies you about the version being outdated:



Warning: OUTDATED — There are newer versions for this entity. Please view the versions tab to see the latest version. Most actions will be disabled.



OUTDATED

There are newer versions for this entity. Please view the **versions** tab to see the latest version. Most actions will be disabled.

Manage the entity

Click the **Actions** pop-up menu on the bottom half of the entity detail pane tab and select the desired option to manage the entity and act on it. You can:

- Edit it;
- Delete it;
- Add it to a dataset ;
- Load it on the graph for analysis;
- Create a follow-up task for the entity;
- Export it as JSON or STIX;
- Download it in its original data format; for example, the original STIX package containing the entity.

View entity history

The History tab displays an overview in reverse chronological order of the actions applied to the entity.

About the entity detail pane

The entity detail pane is a structured container holding a detailed and exhaustive information overview of an entity. You can use the entity detail pane as a reference resource you can look up when you want to zero in on a specific entity to review its structure, any observables it contains or it is related to, any relationships with external entities and observables, where it comes from, and if you are leveraging its intelligence value or not.

Through the entity detail pane you can also edit the entity — for example, to update information — load it on the graph, as well as include it in a workflow by adding it to a dataset and/or to a workspace, or by creating a follow-up task.

Access the entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.
- On the active view, browse to the entity you want to inspect, and click it. The entity detail pane slides in from the side of the screen.

View the entity history

Click the **History** tab to display an overview in reverse chronological order of the actions performed on the entity since its creation. Subsequent creation actions record the creation of new versions of the entity after a change was applied.

This reference view enables you to inspect *what happened* to the entity (the action), *who did it* (the user), and *when it happened* (the date and time).

Manage the entity

Click the **Actions** pop-up menu on the bottom half of the entity detail pane tab and select the desired option to manage the entity and act on it. You can:

- Edit it;

- Delete it;
- Add it to a dataset ;
- Load it on the graph for analysis;
- Create a follow-up task for the entity;
- Export it as JSON or STIX;
- Download it in its original data format; for example, the original STIX package containing the entity.

Filter menus

Filter platform entities with quick filters and drop-down context filters to isolate and drill down into specific bits of information.

Use the quick filters

Filters () make it easy to highlight and isolate specific clusters of information you want to zero in on during an analysis. They help you identify entities sharing properties and attributes, which in turn can disclose connections and links among entities.

By default, quick filters are switched off. To toggle quick filter visibility click .

Most platform views include a range of quick filters with selectable checkboxes or input fields to quickly filter content on the current view, and to show or hide specific subsets, based on the selected shared properties and attributes.

The number and type of available quick filters may vary: quick filters are enabled and they become available only when there are selectable values, and therefore filtering options, for the data points the filters process.

If the current view does not allow filtering based on, for example, dataset or TLP because no entities on the view belong to a dataset or have a TLP color code, the corresponding filters are disabled and they are not displayed. This behavior applies to all filters.

Click a filter group name to expand the corresponding child elements with the selectable options.

Some filters with many options — for example, **Dataset** — feature a search input field with autocompletion. If a filter includes the search field, you can start typing the names of the desired values to look them up without browsing through the option list. You can stack and combine filters as you need.



About the Dataset filter

- You can filter and search for entities by selecting one or more static datasets in the **Dataset** filter option available on most platform views.
- You **cannot** filter or search for entities by selecting dynamic datasets. Dynamic datasets are not included in the **Dataset** filter option available on most platform views.
- The **Dataset** filter is not available when the results do not include any entities belonging to at least one dataset.

Classification

Filters objects on the current view by maliciousness classification.

Select one or more checkboxes to include in the resulting filtered view observables flagged as malicious, safe, or unknown.

The filter is available in the following platform areas:

- **Browse > Observables**

Connections

Filters objects on the current view by number of connections with other entities in the platform.

Enter a minimum and a maximum number of connections to include in the resulting filtered view observables whose numbers of connections/links with other entities match the specified range.

The filter is available in the following platform areas:

- **Browse > Observables**

Date

Filters objects on the current view by date.

Select a start and an end date to include in the resulting filtered view entities ingested within the specified time range.

The filter is available in the following platform areas:

- **Browse > Entities**
- **Browse > Published**
- **Browse > Draft**
- **Discovery**

Dataset

Filters objects on the current view by dataset.

Select one or more checkboxes to include in the resulting filtered view entities belonging to the specified datasets.

The filter is available in the following platform areas:

- **Browse > Entities**
- **Browse > Published**
- **Browse > Draft**
- **Discovery**
- **Exposure**

Discovery rules

Filters objects on the current view by one or more specific discovery rules.

Select one or more checkboxes to include in the resulting filtered view entities whose properties and attributes match the selection criteria of the specified discovery rules.

The filter is available in the following platform areas:

- **Discovery**

Entity

Filters objects on the current view by entity type.

Select one or more checkboxes to include in the resulting filtered view the specified entity types.

The filter is available in the following platform areas:

- **Browse > Entities**
- **Browse > Published**
- **Browse > Draft**
- **Discovery**
- **Exposure**
- **Datasets**
- Incoming feed detail pane, **Content** tab

Kind

Filters objects on the current view by observable data type.

Select one or more checkboxes to include in the resulting filtered view observables whose data types match the specified values.

The filter is available in the following platform areas:

- **Browse > Observables**

Reliability

Filters objects on the current view by data source reliability.

Select one or more checkboxes to include in the resulting filtered view entities ingested from data sources whose reliability level matches the specified value(s).

The filter is available in the following platform areas:

- **Browse > Entities**
- **Browse > Published**
- **Discovery**
- **Exposure**

Source

Filters objects on the current view by data source (incoming feeds, enrichers, user groups).

Select one or more checkboxes to include in the resulting filtered view entities ingested from the specified data sources.

The filter is available in the following platform areas:

- **Browse > Entities**
- **Browse > Observables**
- **Browse > Published**
- **Browse > Draft**
- **Discovery**

Timestamp

Filters objects on the current view by observable timestamp.

Select a start and an end date to include in the resulting filtered view observables whose timestamps fall inside the specified time range.

The filter is available in the following platform areas:

- **Browse > Observables**

TLP

Filters objects on the current view by

Select one or more checkboxes to include in the resulting filtered view entities flagged with the specified TLP color codes.

The filter is available in the following platform areas:

- **Browse > Entities**
- **Browse > Published**
- **Browse > Draft**
- **Discovery**

Use the context filters

Besides quick filters, you can also access contextual drop-down filter menus. The drop-down filter menus are available on selected platform views, and on most detail panes for entities, observables, datasets, and feeds.

Filtering options for these menus may vary, depending on where in the platform they are available.

Classification

Filters objects on the current view by maliciousness classification.

Select one or more checkboxes to include in the resulting filtered view observable rules that apply an action to flag matching observables as malicious, safe, or unknown.

The filter is available in the following platform areas:

- **Rules > Observable**

Date

Filters objects on the current view by date.

Select a start and an end date to include in the resulting filtered view observables created within the specified time range.

The filter is available in the following platform areas:

- Entity detail pane > **Observables** tab
- Observable detail pane > **Observables** tab
- Incoming feed detail pane > **Content** tab

Entity types

Filters objects on the current view by entity type.

Select one or more checkboxes to include in the resulting filtered view the specified incoming feed entity types.

The filter is available in the following platform areas:

- Incoming feed detail pane > **Content** tab

Kind

Filters objects on the current view by observable data type.

Select one or more checkboxes to include in the resulting filtered view observables whose data types match the specified values.

The filter is available in the following platform areas:

- Entity detail pane > **Observables** tab
- Observable detail pane > **Observables** tab

Lv

Filters objects on the current view by observable level, either **1** or **2**.

Select a checkbox to include in the resulting filtered view either embedded CyBOX observables (level 1, possibly more relevant) or STIX-field level observables (level 2, possibly less relevant).

The filter is available in the following platform areas:

- Entity detail pane > **Observables** tab

Maliciousness

Filters objects on the current view by maliciousness classification.

Select one or more checkboxes to include in the resulting filtered view observables flagged as malicious, safe, or

unknown.

The filter is available in the following platform areas:

- Entity detail pane > **Observables** tab
- Observable detail pane > **Observables** tab

My tasks

Filters objects on the current view by task user role.

Select one or more checkboxes to include in the resulting filtered view tasks the current user created, or tasks the current user is assigned to.

The filter is available in the following platform areas:

- **Workspaces > <workspace_name> > Tasks**

Origin

Filters objects on the current view by observable data source.

Select one or more checkboxes to include in the resulting filtered view observables whose data sources match the specified values.

The platform creates observables as a result of the ingestion process, after running an enricher, or when a user manually adds an observable to an entity.

The filter is available in the following platform areas:

- Entity detail pane > **Observables** tab

Show

Filters objects on the current view by the specified properties or attributes.

Select one or more checkboxes to include in the resulting filtered view items whose properties or attributes match the filter selection.

The filter is available in the following platform areas:

- **Workspaces**, where the available filtering options are **Case**, **Generic**, **Team**, and **Topic**.
- **Rules > Entity**, where the available filtering options are **Disabled** and **Enabled**.
- **Rules > Observable**, where the available filtering options are **Disabled** and **Enabled**.

Source

Works like the **Source** quick filter.

The filter is available in the following platform areas:

- **Rules > Entity**
- **Rules > Observable**

Status

Filters objects on the current view by task status.

Select one or more checkboxes to include in the resulting filtered view tasks whose workflow status matches the specified values.

The filter is available in the following platform areas:

- **Workspaces > <workspace_name> > Tasks**

Filter by source reliability

Filter entities based on the level of reliability of the data source to retrieve and focus on relevant intelligence and minimize data noise.

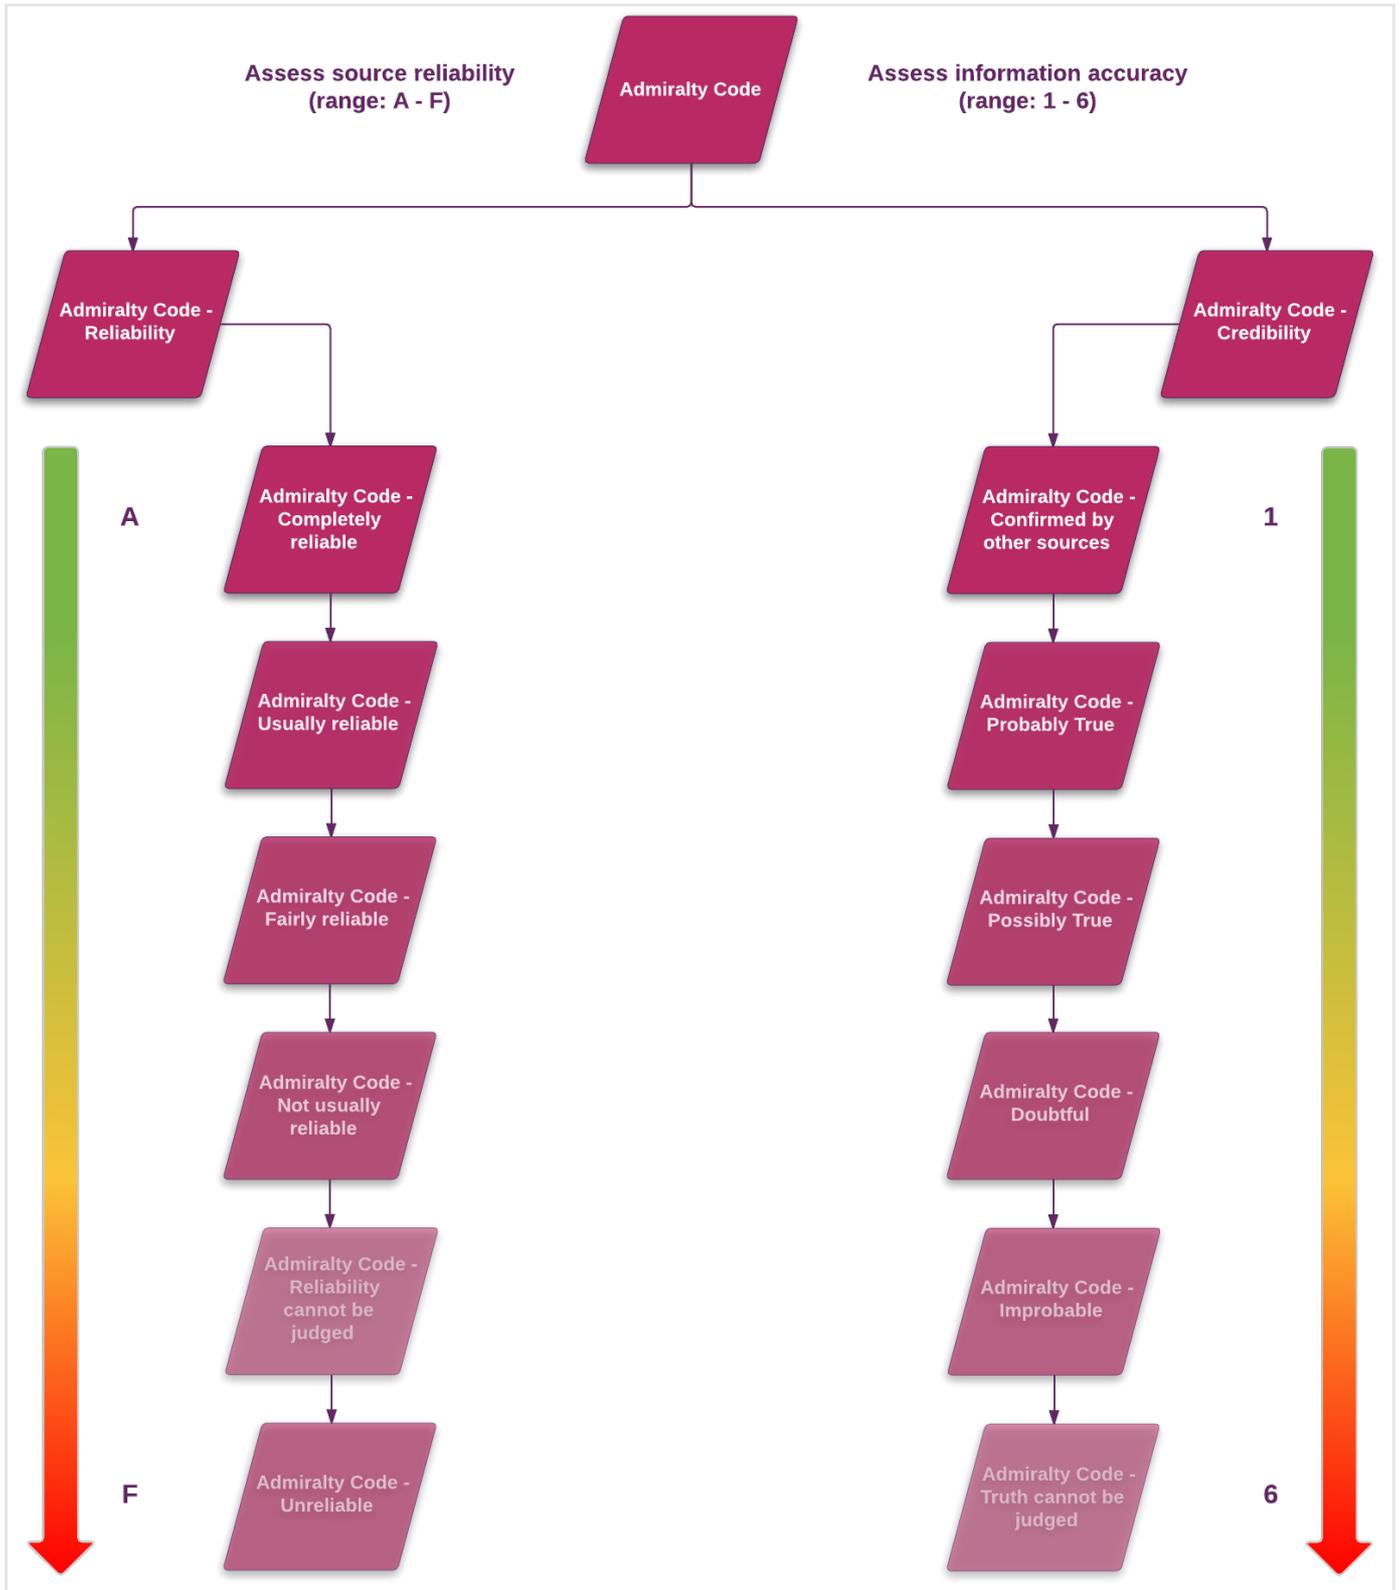
About source reliability

You can set a source reliability value for ingested and created entities:

- When you create a new entity, you can include a reliability flag in the entity `meta.source_reliability` metadata field.
- When you configure an incoming feed, you can set a source reliability value that is applied to all entities ingested through that feed.

It serves as an indication to help assess the level of accuracy and trustworthiness of the data source the entity originates from.

Values in this menu have the same meaning as the first character in the **two-character Admiralty System code** (https://en.wikipedia.org/wiki/admiralty_code).



Filter entities by source reliability

You can search for entities flagged with a specific source reliability value or that fall within a reliability range:

- On the top navigation bar click the **Q** search icon.

- In the search input field prepend `meta.source_reliability:` to the specified source reliability value(s) you want to use as search criteria.

Example:

```
/* Searches for all entities whose  
   source reliability value is 'A' */  
meta.source_reliability:A  
  
/* Searches for all entities whose  
   source reliability value is either 'A', or 'B', or 'C' */  
meta.source_reliability:(A B C)
```

Filter by tag and taxonomy

Tags and taxonomies organize and structure data, so that it is easier to search for and to find entities in the platform.

Search for entities by tag

You can search for entities sharing the same tag:

- On the top navigation bar click the **Q** search icon.
- In the search input field prepend `tags:` to the specified search term(s) to look up entities tagged with custom (`meta.tags`) and with taxonomy (`meta.taxonomy`) tags.
- In the search input field prepend `meta.tags:` to the specified search term(s) to look up entities tagged with custom tags — taxonomy tags are excluded.
- In the search input field prepend `meta.taxonomy:` to the specified search term(s) to look up entities tagged with taxonomy tags — custom tags are excluded.

Alternatively:

- Open an entity detail pane.
- Under **Tags**, click a tag to automatically run a search based on it, and to display the corresponding results.

Filter entities by tag on the graph

Select and deselect tags in the **Tags** category on the histogram to show and hide entities on the graph based on the corresponding tags.

To open the histogram pane, click the **Histogram** icon on the top navigation bar:



- Select a checkbox to display nodes with the corresponding property or attribute.
- Deselect a checkbox to hide nodes with the corresponding property or attribute.
- By default, all checkboxes are selected, that is, nothing is filtered out, and all the nodes and the relationships loaded on the graph are visible.

To filter entity visibility on the graph by tag, browse to the **Tags** category, which groups all the tags belonging to the entities loaded on the current graph.

- **Tags:** select one or more options in this category to view entities flagged with the specified tags. For example, you can use this filter to include in the resulting graph view only entities with specific Admiralty codes or kill chain values.
 - **Without tags:** select this checkbox to view untagged entities.

Filter by TLP color

Filter entities by Traffic Light Protocol (TLP) color code to provide a quick reference to assess how sensitive a piece of information is, how urgent, and the appropriate audience recipients can share it with.

About TLP

Traffic Light Protocol (<https://www.us-cert.gov/tlp>) color code.

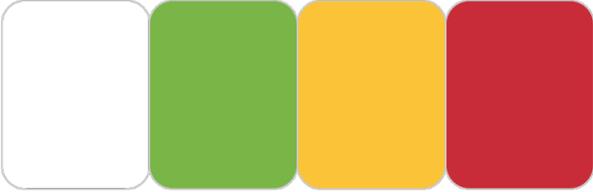
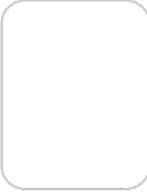
TLP is used to flag information to provide handling and sharing guidelines. It indicates if the information:

- Is sensitive/reserved, or if you can share it with other parties.
- Holds high risk, if it is useful to promote awareness of the content it describes, or if it holds no foreseeable risk of misuse.
- Requires immediate action (deter/prevail), or if it can be part of a longer term strategy (prevent).

You can override the original or the current TLP color code of an entity, an incoming feed, or an outgoing feed.

When working as a filter, TLP colors select a decreasing range: if you set a TLP color as a filter the enricher, the feed, or the returned filtered results include all the entities flagged with the selected TLP color code, as well as all the entities whose TLP color indicates that they are progressively lower risk, less sensitive, and suitable for disclosure to broader audiences.

For example, if you select green the filtered results include entities with a TLP color set to green, as well as entities with a TLP color set to white, and entities with no TLP color code flag.

TLP color	Description	Filter results
 <p>None</p>	<p>TLP set to None:</p> <ul style="list-style-type: none"> • There are no guidelines about sharing and disclosing the information. • When you filter by None, the filter returns all entities. In other words, None does not filter out any entities. 	
 <p>White</p>	<p>TLP set to White:</p> <ul style="list-style-type: none"> • You can publicly share and disclose the information. • The risk level it carries is minimal. • When you filter by White, the filter returns all entities whose TLP value is set to White. 	
 <p>Green</p>	<p>TLP set to Green:</p> <ul style="list-style-type: none"> • You can share the information within the organization and with peer/partner organizations. • The risk level it carries is low/awareness. • When you filter by Green, the filter returns all entities whose TLP value is set to Green or White. 	
 <p>Amber</p>	<p>TLP set to Amber:</p> <ul style="list-style-type: none"> • You can share and disclose the information within a team and/or with selected organizations that need to act on it. • The risk level it carries is moderate, and action is required. • When you filter by Amber, the filter returns all entities whose TLP value is set to Amber, Green or White. 	
 <p>Red</p>	<p>TLP set to Red:</p> <ul style="list-style-type: none"> • This is reserved/sensitive information that you should not share outside the recipients' circle. • The risk level it carries is serious, and it requires immediate action. • When you filter by Red, the filter returns all entities whose TLP value is set to Red, Amber, Green or White. 	

Filter entities by TLP color code

You can search for entities flagged with specific TLP color codes:

- On the top navigation bar click the **Q** search icon.
- In the search input field prepend `meta.tlp_color:` to the specified TLP color code you want to use as search criterion.
- Enter color names in all uppercase.

Example:

```
/* Searches for entities whose TLP color code is
  'RED', 'AMBER', 'GREEN', or 'WHITE'. */
meta.tlp_color:RED

/* Searches for entities whose TLP color code is
  'AMBER', 'GREEN', or 'WHITE'. */
meta.tlp_color:AMBER

/* Searches for entities whose TLP color code is
  'GREEN' or 'WHITE'. */
meta.tlp_color:GREEN

/* Searches for entities whose TLP color code is
  'WHITE'. */
meta.tlp_color:WHITE
```

Tag and classify entities

Tag entities to organize them and make them easier to find. Structure tags into taxonomies to create a set of controlled categories to classify data and to improve information retrieval in the platform.

About taxonomies

Taxonomies are structured categories. Taxonomies make it easier for you to organize and maintain content, and they help other users find what they are looking for. They provide a hierarchical framework to structure tags and to describe parent-child relationships between tagged topics. Tag relationships provide a reference grid that makes content easier to navigate and to retrieve.

The main benefits of implementing a taxonomy are:

- Label information in a structured way to make it easier to navigate and to retrieve.
- Provide a reference framework to control entity tagging in the platform, so that tags remain meaningful and consistent.
- Deliver more accurate search results.

Platform taxonomies enable you to define specific categories to organize tagged entities. Besides the predefined ones, you can create as many taxonomies as you need to make it easier for users to discover meaningful information in the platform data corpus.

Predefined taxonomies

EclecticIQ Platform ships with the following predefined taxonomy sets:

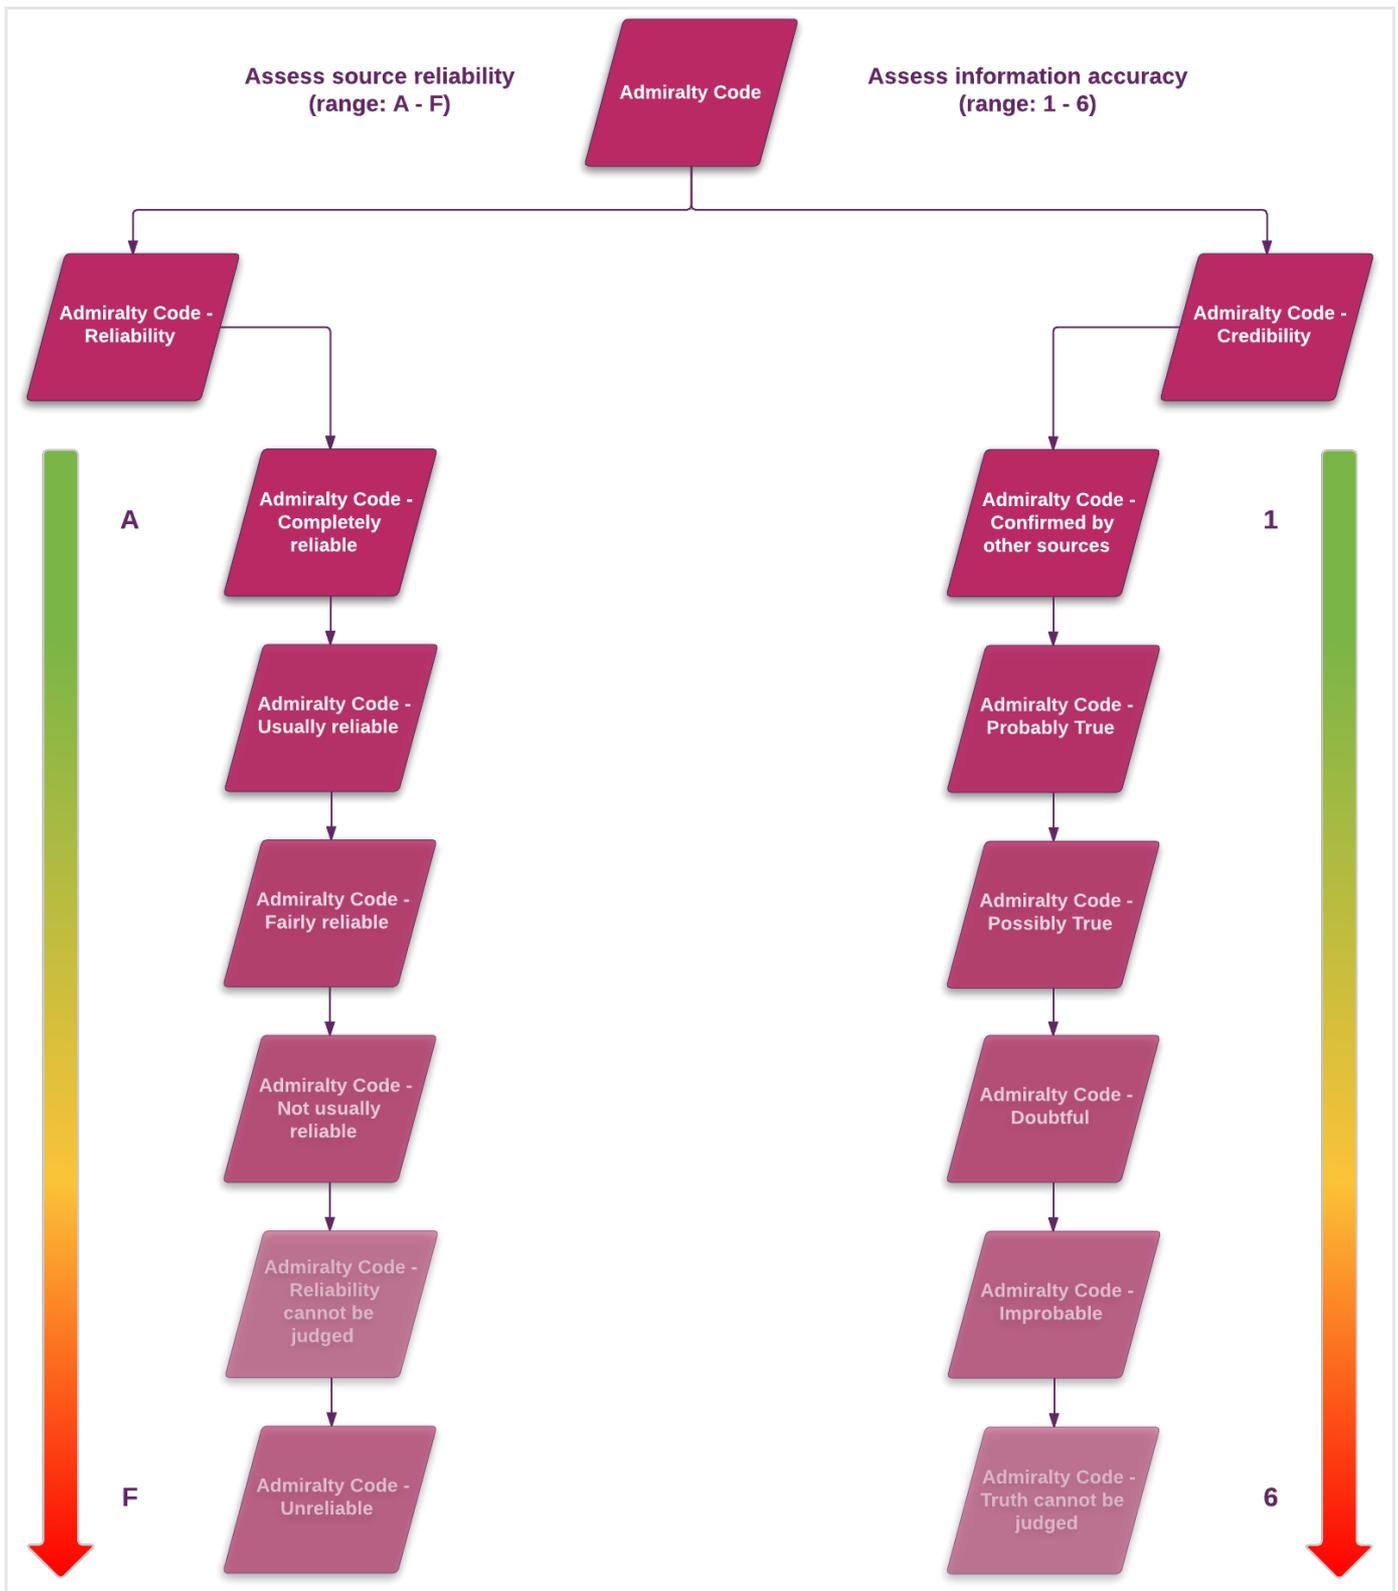
- **Admiralty code**: based on the **two-character Admiralty System code** (https://en.wikipedia.org/wiki/admiralty_code), it helps assess and categorize the reliability of a data source, and the accuracy of the information obtained through a data source.
- **Kill chain phases**: describes the different stages of an attack or an intrusion. By doing so, it helps identify the point(s) in the **kill chain** (<http://www.net-security.org/article.php?id=2220&p=1>) where it is possible to intervene with a mitigation action.

The Admiralty code

Use the **Admiralty code** (https://en.wikipedia.org/wiki/admiralty_code) taxonomy to label entities with tags that define the level of reliability of the data source and the level of accuracy of the entity information. The Admiralty code taxonomy makes it easier to filter entities and information based on criteria such as relevance and credibility. It provides intuitive guidance to retrieve reliable and accurate information more easily, while leaving out unwanted data noise.

Data source reliability	Data accuracy
Completely reliable	Confirmed by other sources
Usually reliable	Probably True

Data source reliability	Data accuracy
Fairly reliable	Possibly True
Not usually reliable	Doubtful
Reliability cannot be judged	Improbable
Unreliable	Truth cannot be judged



The kill chain

In the context of cyber threat defense, a **kill chain** (https://en.wikipedia.org/wiki/kill_chain) aims at encouraging proactive defense, and at implementing adequate courses of action as early as possible in the chain.

The kill chain provides a structured model to:

- Break down the actions of an adversary. This helps understand the TTPs the adversary is implementing.
- In case of an ongoing attack or intrusion, identify the current stage of the intrusion and quantify damage.
- Inspect the kill chain to identify the root cause of the attack or the intrusion.

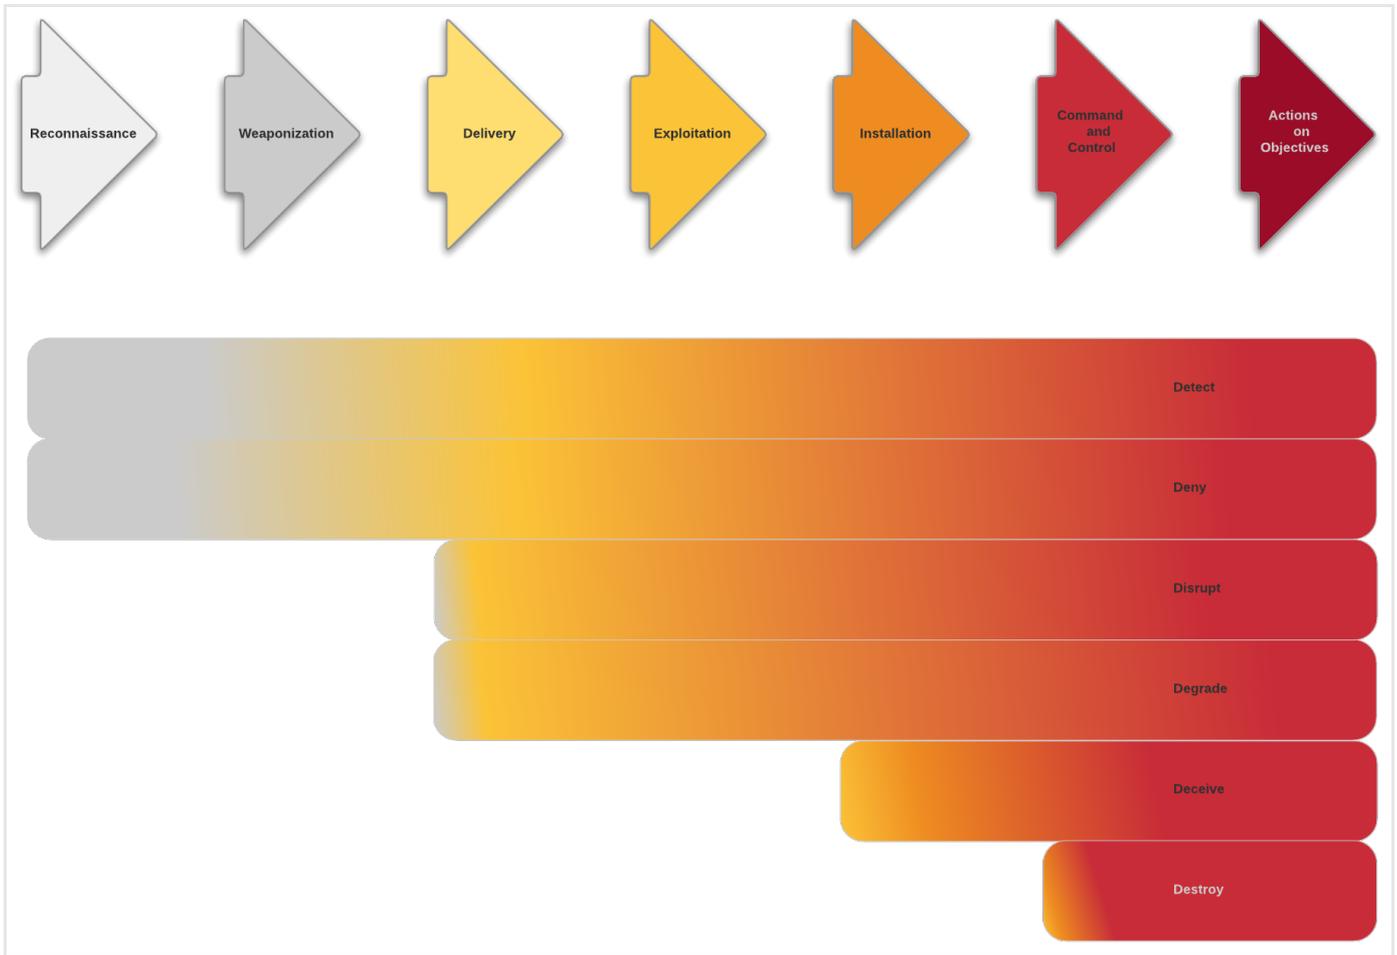
- Plan a defensive course of action to neutralize the adversary.

Kill chain phase	Description
Reconnaissance	Research, identification and selection of targets, often represented as crawling Internet websites such as conference proceedings and mailing lists for email addresses, social relationships, or information on specific technologies.
Weaponization	Coupling a remote access trojan with an exploit into a deliverable payload, typically by means of an automated tool (weaponizer). Increasingly, client application data files such as Adobe Portable Document Format (PDF) or Microsoft Office documents serve as the weaponized deliverable.
Delivery	Transmission of the weapon to the targeted environment. The three most prevalent delivery vectors for weaponized payloads by APT actors, as observed by the Lockheed Martin Computer Incident Response Team (LM-CIRT) for the years 2004-2010, are email attachments, websites, and USB removable media.
Exploitation	After the weapon is delivered to victim host, exploitation triggers intruders' code. Most often, exploitation targets an application or operating system vulnerability, but it could also more simply exploit the users themselves or leverage an operating system feature that auto-executes code.
Installation	Installation of a remote access trojan or backdoor on the victim system allows the adversary to maintain persistence inside the environment.
Command and Control (C2)	Typically, compromised hosts must beacon outbound to an Internet controller server to establish a C2 channel. APT malware especially requires manual interaction rather than conduct activity automatically. Once the C2 channel establishes, intruders have "hands on the keyboard" access inside the target environment.
Actions on Objectives	Only now, after progressing through the first six phases, can intruders take actions to achieve their original objectives. Typically, this objective is data exfiltration which involves collecting, encrypting and extracting information from the victim environment; violations of data integrity or availability are potential objectives as well. Alternatively, the intruders may only desire access to the initial victim box for use as a hop point to compromise additional systems and move laterally inside the network.

(Source: **Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains** (<http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/lm-white-paper-intel-driven-defense.pdf>), by Eric M. Hutchins, Michael J. Cloppert, Rohan M. Amin, Ph.D. Paper presented at the 6th Annual International Conference on Information Warfare and Security, Washington, DC, 2011.

Course of action	Description
Detect	Example: use analytics, auditing, logging tools, and intrusion detection systems (IDS) to detect the intrusion.
Deny	Example: use patching, firewall rules, access control lists (ACL), and intrusion prevention systems (IPS) to deny exploitation.
Disrupt	Example: use data execution prevention (DEP) and intrusion prevention systems to block or otherwise disturb exploitation.
Degrade	Example: use queuing or a tarpit to hinder or otherwise reduce exploitation.
Deceive	Example: use DNS redirection or a honeypot to divert exploitation to a decoy.

Course of action	Description
Destroy	Take control of the attacker’s system to neutralize it.



Create a taxonomy entry

✔ Input fields marked with an asterisk are required.

To create a new taxonomy entry to categorize entity tags, do the following:

- On the top navigation bar click **Data configuration > Taxonomies**.

Alternatively:

- On the top navigation bar click **Data configuration > Taxonomies > Taxonomies > +**.

- On the **Data configuration > Taxonomies > + > Create taxonomy** page, fill out the input fields to define the new taxonomy entry:
 - **Name:** enter a name for the taxonomy entry. The name you specify here corresponds to the tag name you can assign to entities.
 - **Description:** enter a short explanation of what the entry represents or refers to.
 - **Parent:** you can structure taxonomy entries hierarchically by flagging them as either *parent* top-level entries, or subordinate *child* entries.
 - To create a parent entry, leave the field empty.
 - To create a child entry, from the drop-down menu select the parent entry you want to relate the child to. A child taxonomy entry can be the parent of another child entry nested one level beneath.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new:** saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate:** saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit a taxonomy entry

You can edit only user-created, custom taxonomy entries. You cannot edit the predefined Admiralty code and Kill chain taxonomies.

To edit an existing taxonomy entry, do the following:

- On the top navigation bar click **Data configuration > Taxonomies**.
 The **Data configuration > Taxonomies** page displays an overview of the existing entries.
 You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.
- On the overview table, click the  icon.
- From the drop-down menu select **Edit**.

Entry name ^	Description	Parent	Last changed
Kill chain phase - Weaponization		Kill Chain Phases	13.07.2017 10:38
Kill Chain Phases	Kill Chain Phases as defined by Lockheed Martin		13.07.2017 10:38
Kill chain phase - Reconnaissance		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Installation		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Exploitation		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Delivery		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Command and Control		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Actions on Objectives		Kill Chain Phases	13.07.2017 10:38
foo	foo		Last Tuesday at 15:06
foo	foo	Kill Chain Phases	12.08.2017 10:01
bar	bar	foo	Last Tuesday at
Admiralty Code - Usually reliable		Admiralty Code - Reliability	13.07.2017 10:38

- On the **Data configuration > Taxonomies > ⋮ > Edit taxonomy** page, edit the name, the description, or the parent-child hierarchy relationship as needed.
- Click **Save** to store your changes, or **Cancel** to discard them.

Delete a taxonomy entry

You can delete only user-created, custom taxonomy entries. You cannot delete the predefined Admiralty code and Kill chain taxonomies.

To delete an existing taxonomy entry, do the following:

- On the top navigation bar click **Data configuration > Taxonomies**.
The **Data configuration > Taxonomies** page displays an overview of the existing entries.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.
- On the overview table, click the ⋮ icon.
- From the drop-down menu select **Delete**.

Entry name ▲	Description	Parent	Last changed
Kill chain phase - Weaponization		Kill Chain Phases	13.07.2017 10:38
Kill Chain Phases	Kill Chain Phases as defined by Lockheed Martin		13.07.2017 10:38
Kill chain phase - Reconnaissance		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Installation		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Exploitation		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Delivery		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Command and Control		Kill Chain Phases	13.07.2017 10:38
Kill chain phase - Actions on Objectives		Kill Chain Phases	13.07.2017 10:38
foo	foo		Last Tuesday at 15:06 ⋮
foo	foo	Kill Chain Phases	12.08.2017 10:01 Edit
bar	bar	foo	Last Tuesday at Delete
Admiralty Code - I actually reliable		Admiralty Code - Reliability	13.07.2017 10:38

- On the confirmation pop-up dialog, click **Delete** to confirm the action.
- The taxonomy entry is deleted.

If you delete a taxonomy entry that is a parent to one or more children entries, the children related to the removed parent remain available in the taxonomy. However, they lose the parent-child relationship, and they become top-level taxonomy entries.

About tags

Tags provide a quick and intuitive way to sort heaps of data into manageable mounds. They help organize a data corpus by dividing it into subsets containing entities that have some details in common. They provide basic information about an entity, so that analysts can get an idea without having to read all the available details about the entity. They make content index- and search-friendlier. They make it easier for users to find what they are looking for.

You can add to an entity as many tags as you need. However, when it comes to tagging less is more: efficient tagging is consistent, relevant, meaningful, and economical.

Manually tag entities

You can manually tag an entity from almost anywhere in the platform.

To do so, go to an entity overview page. For example, by selecting **Intelligence > All intelligence > Browse , Production, Discovery or Exposure** on the top navigation bar, or by clicking **Intelligence > All Intelligence > Browse > Datasets > <dataset_name>** or **Data configuration > Incoming feeds > <incoming_feed_name> > Entities** tab on the feed detail pane.

- On the active view, browse to the entity you want to tag.
- Click the  icon corresponding to the entity you want to tag.
- From the drop-down menu select **Edit**.
- On the entity edit page, go to **Tags** under the **Meta** section: from the drop-down menu select one of the available tags to add it to the entity.
- If the tag you want to add to the entity does not exist, you can create it on the fly: type the label designating the new tag in the input field, and then press **ENTER** to confirm the new tag. The tag is automatically created and it is assigned to the entity.
- To remove a selection from the input field, click the  icon corresponding to the item(s) you want to remove.

Alternatively:

- On the active view, browse to the entity you want to tag.
- Click anywhere on the row corresponding to the entity you want to tag. The entity detail pane slides in from the side of the screen.
- Under **Tags** from the drop-down menu select one of the available tags to add it to the entity.
- If the tag you want to add to the entity does not exist, you can create it on the fly: type the label designating the new tag in the input field, and then press **ENTER** to confirm the new tag. The tag is automatically created and it is assigned to the entity.
- To remove a selection from the input field, click the  icon corresponding to the item(s) you want to remove.

Auto-tag entities

It is possible to configure custom entity rules to automatically assign one or more predefined tags to specific ingested entities:

- In the rule editor, from the **+ Actions** drop-down menu select **Add tags**.
- From the **Tags** drop-down menu select one or more tags, or create them on the fly.
- Under **Criteria selection**, define the conditions to implement in the rule to specify which entities should be tagged automatically.

When the rule is active, it looks for any published entities matching the rule criteria, and it assigns them the specified tags.

Create tags

To create and to manage tags, use [Taxonomy](#).

Delete tags

To manage and to delete tags, use [Taxonomy](#).

Search for entities by tag

You can search for entities sharing the same tag:

- On the top navigation bar click the **Q** search icon.
- In the search input field prepend `tags:` to the specified search term(s) to look up entities tagged with custom (`meta.tags`) and with taxonomy (`meta.taxonomy`) tags.
- In the search input field prepend `meta.tags:` to the specified search term(s) to look up entities tagged with custom tags — taxonomy tags are excluded.
- In the search input field prepend `meta.taxonomy:` to the specified search term(s) to look up entities tagged with taxonomy tags — custom tags are excluded.

Alternatively:

- Open an entity detail pane.
- Under **Tags**, click a tag to automatically run a search based on it, and to display the corresponding results.

Filter entities with the histogram

The histogram enables you to apply one or more quick filters to the entities and observables on the graph to isolate specific subsets, based on shared properties and attributes.

When you analyze entities and observables on the graph canvas to explore relationships and to, almost literally, join the dots you may want to apply quick filters to the elements on the graph without having to move them around or temporarily remove them.

The histogram helps you filter and visually isolate specific subsets of the elements on the graph, based on shared/common properties and attributes.

To open the histogram pane, click the **Histogram** icon on the top navigation bar:



You can select one or more options by clicking the corresponding checkbox:

- Select a checkbox to display nodes with the corresponding property or attribute.
- Deselect a checkbox to hide nodes with the corresponding property or attribute.
- By default, all checkboxes are selected, that is, nothing is filtered out, and all the nodes and the relationships loaded on the graph are visible.

The histogram pane makes available many ready-to-use filters. You can stack and combine filters as you need.

- **Show singletons**: select this checkbox to view singleton nodes. They are isolated nodes with no relationships to any other nodes.
- **Entity type**: select one or more options in this category to view specific entity types.
 - **Multi-type-group**: select this checkbox to view grouped entities containing mixed entity types.
- **Observable type**: select one or more options in this category to view specific observable types.
- **Source**: select one or more options in this category to view entities and observables ingested from specific data sources, that is, incoming feeds and enrichers.
 - **Missing source**: select this checkbox to view entities and observables that are not associated with any data source.
- **TLP**: select one or more options in this category to view entities flagged with the specified TLP color codes. For example, you can use this filter to include in the resulting graph view only entities flagged as reserved, or that require immediate action.
 - **Missing TLP**: select this checkbox to view entities with no TLP flag.
- **Source reliability**: select one or more options in this category to view entities and observables flagged with the specified source reliability value. For example, you can use this filter to include in the resulting graph view only entities and observables originating from trustworthy data sources.
 - **Missing source reliability**: select this checkbox to view entities and observables that are not associated with any data source.
- **Confidence**: select one or more options in this category to view entities and observables flagged with the specified level of confidence; it flags the **estimated level of confidence** (https://docs.oasis-open.org/cti/stix/v1.2.1/csprd01/part14-vocabularies/stix-v1.2.1-csprd01-part14-vocabularies.html#_toc440440605) to assess the accuracy and trustworthiness of the entity information.
 - **Missing confidence**: select this checkbox to view entities whose confidence level is not set.

- **Observable classification:** select one or more options in this category to view observables flagged with the specified level of maliciousness.
For example, you can use this filter to include in the resulting graph view only observable flagged as **Bad**.
- **Missing observable classification:** select this checkbox to view entities and observables whose maliciousness confidence level is not set.
- **Bad:** select this checkbox to view observables whose maliciousness confidence level is set to **Malicious (High confidence)**, **Malicious (Medium confidence)**, or **Malicious (Low confidence)**.
- **Good:** select this checkbox to view observables marked as **Safe**.
- **Tags:** select one or more options in this category to view entities flagged with the specified tags.
For example, you can use this filter to include in the resulting graph view only entities with specific Admiralty codes or kill chain values.
- **Without tags:** select this checkbox to view untagged entities.

Filter entities with the timebar

Open the timebar on the graph to filter entities and observables based on a specified time range, and to observe how a network evolves over time.

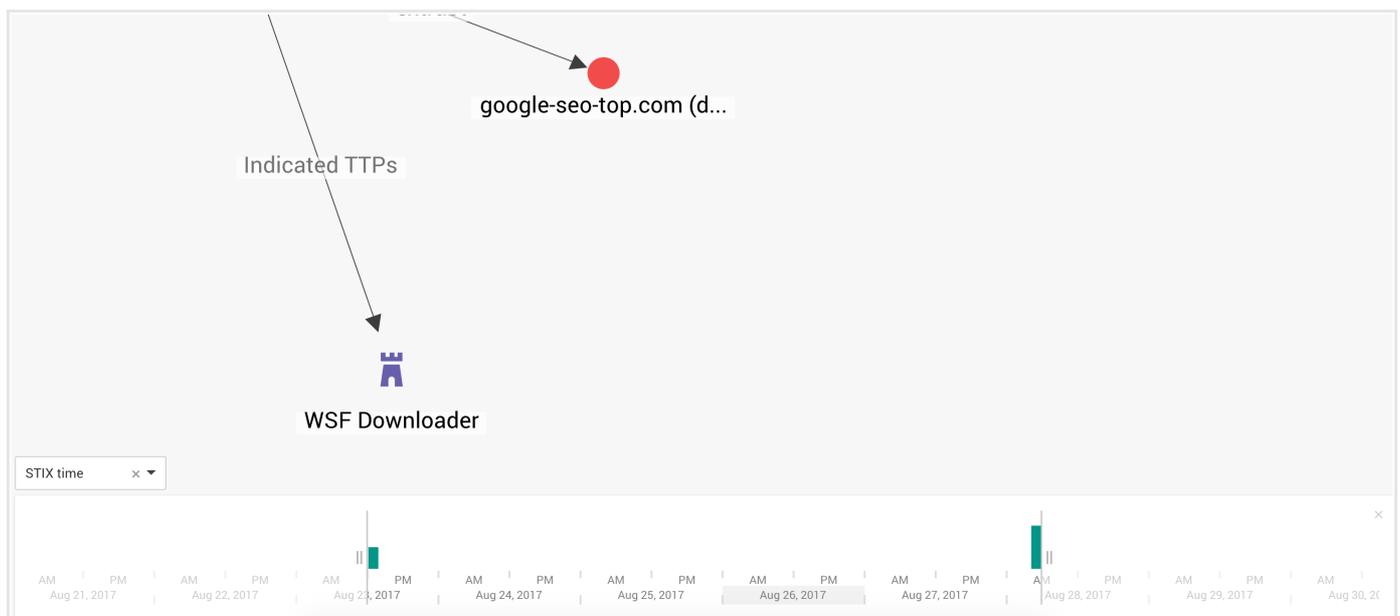
Besides filtering entities and observables based on specific properties and attributes, you can also filter by time range.

The graph timebar enables you to filter nodes on the graph based on a specified time interval; for example, to examine threat scenario evolution, or to focus on a specific phase in the development of the scenario under investigation.

To open the timebar, click the **Timebar** icon on the top navigation bar:



- From the drop-down menu select **STIX time** or **Ingestion time** to filter entities and observables on the graph based on either their STIX timestamp value (`data.timestamp` in the JSON representation of an entity), or their ingestion time into the platform (`created_at` in the JSON representation of an entity).
- Drag the sliders to set a time range to filter the nodes on the graph by.
- Click and drag the timeline on the timebar to the right or to the left to move forward or backward on the timeline, respectively.
- Use the mouse wheel to zoom in to focus on a shorter time range with a more granular level of detail; or to zoom out to get a more general overall view of the threats represented on the graph and their relationships.
- The bar chart on the timebar shows how the network grows over time. Double-click a column to zero in on it to view on the graph only entities and observables whose timestamp or ingestion date falls immediately around the time the selected column represents.



Entity rules

Entity rules enable you to automatically assign taxonomy tags, add entities to a dataset at the end of the ingestion process, or merge almost identical versions of an entity.

When you ingest large quantities of data, you are likely to introduce noise that can clutter your database. Noisy data can make analysis and research more time-consuming and labor-intensive. Wading through a large data soup that includes meaningful information, as well as unnecessary data that does not yield any relevant intelligence value can slow down analysts' decision-making process, and it can make it more error-prone. This has an impact, among others, on prevention and response timeliness.

Entity and observable rules are highly customizable to give you granular control over your data.

For example, you can create rules to target specific entities or observables from predefined data sources, and then automatically add them to a detection or prevention system, or mark them for exclusion to reduce data noise.

Entity rules allow you to automatically assign free tags or taxonomy tags to, merge, and add ingested entities to a dataset. Besides adding semantic relevance, you can use tags within a workflow to group entities sharing similar characteristics. Datasets act as containers providing a focused insight into specific entity subsets. Entity merging helps you reduce unnecessary noise in the database.

Entity rules

Entity rules help you automate tagging entities, adding entities to one or more datasets, and merging almost identical entities, based on a predefined set of criteria.

Add an entity rule

The screenshot shows a web interface for creating an entity rule. At the top, there is a navigation bar with 'Data Configuration' selected, and sub-menus for 'Incoming feeds', 'Outgoing feeds', 'Taxonomies', 'Enrichers', and 'Rules'. Below this, there are tabs for 'OBSERVABLE', 'ENTITY', 'ENRICHMENT', and 'DISCOVERY', with 'ENTITY' being the active tab. The main content area is titled 'Create entity rule' and contains the following fields and sections:

- Rule name ***: A text input field with an asterisk indicating it is required.
- Enabled**: A checkbox to toggle the rule's status.
- Criteria selection**: A section with the instruction 'Entities should match ALL of the following conditions:'. It contains one dropdown menu with a plus sign and the text 'Content'.
- Actions**: A section with one dropdown menu with a plus sign and the text 'Actions'.
- Buttons**: 'CANCEL' and 'SAVE' buttons with a dropdown arrow.

✓ Input fields marked with an asterisk are required.

To create a new entity rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Entity** .

- On the **Rules > Entity > Create** page, define the new rule criteria to automatically tag entities, add entities to datasets, or merge entities:
 - **Rule name:** enter a name to identify the rule. It should be descriptive and easy to remember.
 - Select the **Enabled** checkbox to enable the rule immediately after creating it.
 - **Actions:** from the drop-down menu select at least one of the following options:
 - **Add tags:** *all* entities matching *all* the conditions defined under **Criteria selection** are tagged with *all* selected tags.
 - **Tags:** from the drop-down menu select one or more tags to assign to the entities matching the rule criteria. You can select predefined taxonomy tags that follow the Admiralty code system or the Kill chain model, any existing free tags, as well as start typing to create a new tag on the fly. To remove a selected item from the input field, click the **✕** icon on the item(s) you want to deselect. This option is not available if you do not select **Add tags**.
 - **Add to dataset:** *all* entities matching *all* the conditions defined under **Criteria selection** are added to *all* selected datasets.
 - **Datasets:** from the drop-down menu select one or more datasets to add the entities matching the rule criteria to. To remove a selected item from the input field, click the **✕** icon on the item(s) you want to deselect. This option is not available if you do not select **Add to dataset**.
 - Select the **Enabled** checkbox to enable the rule immediately after creating it.

A valid rule needs to include a name, an action, and at least one condition, which you can select and configure under **Criteria selection**.

Click **+ Condition** to define one or more conditions:

- **Entity types:** from the drop-down menu select one or more entity types to apply the rule to. The rule applies the same **Actions** to all selected entity types, that is, it handles all selected entities in the same way.

To remove a selected item from the input field, click the **✕** icon on the item(s) you want to deselect:

Data Configuration | Incoming feeds | Outgoing feeds | Taxonomies | Enrichers | **Rules**

Rule name *

Enabled

Criteria selection

Entities should match ALL of the following conditions:

- > Entity types - ×
- > Content criteria ×
- ▼ Source - ×
 - Source *
- + Content ▼

Actions

- + Actions ▼

PREVIEW RULE

- **Content criteria:** key/value pairs define the content criteria the rule should apply.

The input format for the *key* field is a *JSON* path. It points to an entity field/entity location in the entity structure.

The input format for the *value* field is a *regex*. It specifies the content pattern.

By default, **Content criteria** JSON path expressions are relative to the `data` field, which is the root of the JSON path expression.

The `data` root is implied. To point to the title or to the description fields of an entity, you only need to specify `title` or `description`, instead of `data.title` or `data.description`.

- **Content > Path:** based on the specified JSON path, the rule searches for a corresponding match in the JSON data structure representing entities in the platform.

The JSON path root is the `data` field.

The JSON path is a string that points to a location, that is, a field inside a JSON object. It tells the rule *where* in the entity structure it should go look for the corresponding data value.

Think of it as a friend's address you scribble on the back of a postcard before dropping it into the mailbox.

The JSON path format is a string where dots (.) define JSON parent-child relationships.

Do not include square brackets ([]) in the path input: they are stripped during execution. It is not possible to use square brackets to point to specific array members.

Wildcards are currently not supported.

Examples:

- Input string pattern example: `related_extracts.value`
- The path matching the specified pattern points to any `value` key in the following array:

```
{
  "data": {
    "related_extracts": [
      {
        "kind": "domain",
        "value": "robohelptestng.biz"
      },
      {
        "kind": "ipv4",
        "value": "195.22.28.199"
      },
      {
        "kind": "ipv4",
        "value": "188.200.164.50"
      }
    ]
  }
}
```



To examine the JSON data structure of an entity:

- Go the entity detail pane, and then click the **JSON** tab.

Alternatively:

- On the selected entity detail pane, click **Actions > Export > JSON** to save the entity in JSON format.

- **Content > Value**: define a regex to specify the data pattern the rule should apply to search for the desired content. The regex tells the rule *what* to look for at the location the JSON path points to. Think of it as the front of the postcard you're sending to a friend, the side with the picture of a very stereotypical landscape that can match a number of actual places.

This field supports only **Elasticsearch regular expression syntax**

(<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-regexp-query.html#regexp-syntax>).

The main peculiarities of the Elasticsearch query regex syntax are:

- Anchors (^ and \$) are implied at the beginning and at the end of the regex. You do not need to include them in the regex you input.
- If you insert explicit anchor characters in the **Value** field, they are interpreted as literal values.
- You need to escape special characters (. ? + * | { } [] () " \). To escape a special character, prepend a backslash \ to it. Example: \{ \}

✓ At this moment, Elasticsearch regular expression syntax **optional operators**

(https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-regexp-query.html#_optional_operators) are not supported.

- Click **+ Add** or **+ More** to add new rows as needed, where you can enter additional criteria.
- **Source**: from the drop-down menu select an incoming feed or an enricher to use as a data source for the rule.
- **TLPs**: the TLP color code you want to use to filter data. **TLP** (<https://www.us-cert.gov/tlp>) provides an intuitive reference to assess how sensitive information is, focusing in particular on how serious it is, and whom it should or should not be shared with.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new**: saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate**: saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit rules

To edit an existing observable or entity rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Entity** or **⚙ > Rules > Observable**.
- On the **Rules** page, go to **Rules > Entity** or to **Rules > Observable**, and then click the row corresponding to the rule you want to modify.
- On the entry detail pane, click **Actions > Edit** to go to the form where you can modify the selected rule.
- Enter your changes as needed.
- Click **Save** to store your changes, or **Cancel** to discard them.

Alternatively:

- On the rule overview page, click the  icon corresponding to the rule you want to edit.
- From the drop-down menu select **Edit** to go to the form where you can modify the selected rule.
- Enter your changes as needed.
- Click **Save** to store your changes, or **Cancel** to discard them.

Delete rules

To delete an existing observable or entity rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Entity** or  **> Rules > Observable**.
- On the **Rules** page, go to **Rules > Entity** or to **Rules > Observable**, and then click the row corresponding to the rule you want to delete.
- On the entry detail pane, click **Actions > Delete**.
- On the pop-up confirmation dialog, confirm your choice.
- The rule is removed from the list.

Alternatively:

- On the rule overview page, click the  icon corresponding to the rule you want to delete.
- From the drop-down menu select **Delete**.
- On the pop-up confirmation dialog, confirm your choice.
- The rule is removed from the list.

To disable an active rule, follow the same procedure but instead of selecting **Delete**, from the context menu select **Disable**.

Filter rules

The **Rules** page shows overviews of the existing entity and observable rules.

You can narrow down the displayed results by clicking one or more quick filters above the table view to select and filter by specific:

- **Source**: select the incoming feed(s) and enrichers used as data sources for the rules.
- **Show**: select if you want to display only **Enabled** or **Disabled** rules.
- **Classification**: select if you want to display only **Malicious** rules, only **Safe** rules, only **Ignore** rules, or any combination of these options.
This option is available only for observable rules.

Example

For example, let's assume we want to apply an observable rule that zeroes in on ipv4 IP address observables. We want the rule to target IP address observables only when they are included in a sighting.

The criteria we set for the rule are:

- **Entity types:** *Sightings*
- **Observable types:** *Ipv4*
- **Paths:** *related_extracts.value*
- **Value matches:** *(.+)*.abc.com*
- **Source:** in this example, we want the rule to be enabled for all incoming feeds. Therefore, we do not set this condition.

Create observable rule

Rule name *

 Enabled

Criteria selection

Observables should match ALL of the following conditions:

+ Condition ▾

- Entity types
- Observable types
- Paths
- Source
- Value matches
- Derivation
- Link types filter

The path matching the specified pattern points to the ipv4 values in the second and third members of the following array:

```
{
  "data": {
    "related_extracts": [
      {
        "kind": "domain",
        "value": "robohelptestesting.biz"
      },
      {
        "kind": "ipv4",
        "value": "195.22.28.199"
      },
      {
        "kind": "ipv4",
        "value": "188.200.164.50"
      }
    ]
  }
}
```

The array contains the observables related to the sighting, which can have a JSON data structure like this:

```
{
  "alternative_versions": [],
  "attachments": [],
  "created_at": "2016-06-03T10:20:21.515918+00:00",
  "created_by": null,

  "data": {
    "confidence": {
      "type": "confidence",
      "value": "High"
    },

    "description": "Sinowal trojan identified to inform robohelptestesting.biz|195.22.28.199 from 188.200.164.50",
    "impact": "High",

    "raw_events": "{\"trojanfamily\": \"Sinowal\", \"_geo_env_server_addr\": {\"postal_code\": \"1300-125\", \"latitude\": 38.7167, \"region_code\": \"14\", \"longitude\": -9.1333, \"path\": \"env.server_addr\", \"asn_name\": \"ClaraNET LTD\", \"asn\": 8426, \"region\": \"Lisboa\", \"country_code\": \"PT\", \"netmask\": 24, \"city\": \"Lisbon\", \"country_name\": \"Portugal\", \"ip\": \"195.22.28.199\"}, \"_geo_env_remote_addr\": {\"postal_code\": \"3430\", \"latitude\": 52.0148, \"region_code\": \"09\", \"longitude\": 5.1004, \"path\": \"env.remote_addr\", \"asn_name\": \"KPN B.V.\", \"asn\": 1136, \"region\": \"Utrecht\", \"country_code\": \"NL\", \"netmask\": 24, \"city\": \"Nieuwegein\", \"country_name\": \"Netherlands\", \"ip\": \"188.200.164.50\"}, \"env\": {\"server_name\": \"robohelptestesting.biz\", \"remote_port\": \"3805\", \"remote_addr\": \"188.200.164.50\", \"request_method\": \"POST\", \"server_addr\": \"195.22.28.199\", \"path_info\": \"/search2\", \"server_port\": \"80\"}, \"args\": \"fr=altavista&itag=ody&q=ca8584331d1264912bd2e298c38eb88b%2Cdd5701fc75f672e%2C6AS2Me0aD0dEag3aS0h kgs=1&kls=0\", \"_ts\": 1464949055, \"_origin\": \"banktrojan\", \"sd\": 1}\"",

    "related_extracts": [{
      "kind": "domain",
      "value": "robohelptestesting.biz"
    }],

    {
```

```
    "kind": "ipv4",
    "value": "195.22.28.199"
  },

  {
    "kind": "ipv4",
    "value": "188.200.164.50"
  }
],

"title": "Sighting robohelptesting.biz",
"type": "eclecticiq-sighting"
},

"destinations": [],

"exposure": {
  "affected": true,
  "affected_override": null,
  "community_feed": false,
  "detect_feed": false,
  "detect_ok": false,
  "detect_override": null,
  "exposed": true,
  "prevent_feed": false,
  "prevent_ok": false,
  "prevent_override": null,
  "sighted": true
},

"group_id": "1632265a-ac31-49a6-9dd2-3127dcc3a39e",
"id": "00000b8e-8b59-49b3-b04e-d3ddf540a516",
"incoming_stix_relations": [],
"intel_sets": [],
"last_updated_at": "2016-06-03T10:20:21.515918+00:00",

"meta": {
  "blob": 3586667,
  "estimated_observed_time": "2016-06-03T10:17:35",
  "estimated_threat_start_time": "2016-06-03T10:17:35",
  "incoming_feed": 237,
  "ingest_time": "2016-06-03T10:20:21.590912+00:00",
  "source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
  "source_name": "Anubis",
  "source_type": "incoming_feed",
  "title": "Sighting robohelptesting.biz"
},

"outgoing_stix_relations": [{
  "alternative_versions": [],
  "attachments": [],
  "created_at": "2016-06-03T10:20:21.768998+00:00",
  "created_by": null,

  "data": {
    "key": "indicators",
    "source": "00000b8e-8b59-49b3-b04e-d3ddf540a516",
    "source_type": "eclecticiq-sighting",
    "target": "952c4de5-9abe-4904-9211-9c694d775046",
    "target_type": "indicator",
    "type": "relation"
  }
},

"destinations": [],

"exposure": {
  "affected": false,
```

```
"affected_override": null,
"community_feed": false,
"detect_feed": false,
"detect_ok": false,
"detect_override": null,
"exposed": true,
"prevent_feed": false,
"prevent_ok": false,
"prevent_override": null,
"sighted": false
},

"group_id": "1632265a-ac31-49a6-9dd2-3127dcc3a39e",
"id": "a0040965-b3d7-4c91-b247-8d9a5d3d614b",
"intel_sets": [],
"last_updated_at": "2016-06-03T10:20:21.768998+00:00",

"meta": {
  "blob": 3586667,
  "incoming_feed": 237,
  "source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
  "source_name": "Anubis",
  "source_type": "incoming_feed"
},

"relevancy": 1,
"source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
"workspaces": [],
"workspaces_public": []
}],

"relevancy": 1,
"source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
"type": "entities",
"workspaces": [],
"workspaces_public": []
}
```

Observable rules

Observable rules give you granular control over entity observables by automatically flagging them as malicious, safe, or irrelevant. This allows you to automate default reactions by triggering follow-up actions in other components of your prevention/detection toolchain.

When you ingest large quantities of data, you are likely to introduce noise that can clutter your database. Noisy data can make analysis and research more time-consuming and labor-intensive. Wading through a large data soup that includes meaningful information, as well as unnecessary data that does not yield any relevant intelligence value can slow down analysts' decision-making process, and it can make it more error-prone. This has an impact, among others, on prevention and response timeliness.

Entity and observable rules are highly customizable to give you granular control over your data.

For example, you can create rules to target specific entities or observables from predefined data sources, and then automatically add them to a detection or prevention system, or mark them for exclusion to reduce data noise.

Observable rules enable you to automate default reactions by triggering follow-up actions in other components of your prevention/detection toolchain, based on automatically assigned ignore, safe, or malicious flags.

Observable rules

Observable rules filter and flag entity observable as malicious, safe, or ignorable, based on predefined sets of criteria.

Add an observable rule

The screenshot shows a web interface for creating an observable rule. At the top, there is a navigation bar with tabs for 'OBSERVABLE', 'ENTITY', 'ENRICHMENT', and 'DISCOVERY'. The 'OBSERVABLE' tab is selected. Below the navigation bar, the page title is 'Create observable rule'. The form contains the following elements:

- A text input field for 'Rule name *'.
- A dropdown menu for 'Action *' with the placeholder text 'Please select one'.
- An 'Enabled' checkbox.
- A section titled 'Criteria selection' with the instruction 'Observables should match ALL of the following conditions:'.
- A button with a plus sign and the text 'Condition' next to a dropdown arrow.
- At the bottom right, there are two buttons: 'CANCEL' and 'SAVE'.

✓ Input fields marked with an asterisk are required.

To create a new observable rule, do the following:

- On the top navigation bar click **+ > Rules > Observable**.

- On the **Rules > Observable > Create** page, define the new rule criteria to flag entity observables:
 - **Rule name:** enter a name to identify the rule. It should be descriptive and easy to remember.
 - **Action:** from the drop-down menu select one of the following options:
 - **Ignore:** *all* entity observables matching *all* the conditions defined under **Criteria selection** are ignored. If any observables are found that can be ignored, you can delete them in bulk from the platform by selecting **Delete all matching observables**.
It is a good idea to review the specified observables before deleting them.
 - **Mark as safe:** *all* entity observables matching *all* the conditions defined under **Criteria selection** are flagged as safe, and therefore non-threatening.
 - **Mark as malicious** *all* entity observables matching *all* the conditions defined under **Criteria selection** are flagged as malicious. These are the ones you may want to drill down into; for example, by defining rules that trigger follow-up actions in external prevention/detection components, or by requesting further analysis on the potential threat. Setting a maliciousness confidence level makes it easier to triage and to prioritize threat severity.

When you flag an observable with a maliciousness confidence level, it cannot transition to *safe* or *ignore* anymore. It can only become more malicious, that is, it can only transition to a higher, never to a lower, maliciousness confidence level. Once an observable is flagged as harmful, it cannot go back to being safe or irrelevant anymore.

When you select **Mark as malicious**, you can fine-tune the option by making a further distinction based on **Confidence**:

- **Malicious - Low confidence:** based on the available intelligence, the threat represented by the entity observable(s) may or may not be malicious.
 - **Malicious - Medium confidence:** based on the available intelligence, the threat represented by the entity observable(s) is likely to be malicious.
 - **Malicious - High confidence:** based on the available intelligence, the threat represented by the entity observable(s) is malicious.
- Select the **Enabled** checkbox to enable the rule immediately after creating it.

A valid rule needs to include a name, an action, and at least one condition, which you can select and configure under **Criteria selection**.

Click **+ Condition** to define one or more conditions:

- **Entity types:** from the drop-down menu select one or more entity types to apply the rule to.

To remove a selected item from the input field, click the **✕** icon on the item(s) you want to deselect:

The screenshot shows the 'Rules' configuration interface. At the top, there are navigation tabs: 'Data Configuration', 'Incoming feeds', 'Outgoing feeds', 'Taxonomies', 'Enrichers', and 'Rules'. The 'Rules' tab is active. Below the tabs, there is a 'Rule name *' text input field. Underneath is an 'Enabled' checkbox. The 'Criteria selection' section is titled 'Entities should match ALL of the following conditions:'. It contains three criteria items: 'Entity types', 'Content criteria', and 'Source'. The 'Source' item is expanded, showing a 'Source *' dropdown menu with the text 'Please select one'. Below the criteria items is a '+ Content' dropdown menu. The 'Actions' section contains a '+ Actions' dropdown menu. At the bottom of the form is a green 'PREVIEW RULE' button. On the left side, there is a dark sidebar with icons for search, home, add, notifications, and settings.

- **Observable types:** from the drop-down menu select one or more observable types, that is, the data types describing the corresponding entity observable data. For example, you may want to include in the rule processing a specific city name, an actor, a range of IP addresses, or a telephone number.

- **Paths:** based on the specified JSON path, the rule searches for a corresponding match in the JSON data structure representing entities in the platform.
If you specify multiple values, enter one value per line.

The JSON path root is the `data` field.

The JSON path is a string that points to a location, that is, a field inside a JSON object. It tells the rule *where* in the entity structure it should go look for the corresponding data value.

Think of it as a friend's address you scribble on the back of a postcard before dropping it into the mailbox.

The JSON path format is a string where dots (.) define JSON parent-child relationships.

Do not include square brackets ([]) in the path input: they are stripped during execution. It is not possible to use square brackets to point to specific array members.

Wildcards are currently not supported.

Examples:

- Input string pattern example: `related_extracts.value`
- The path matching the specified pattern points to any `value` key in the following array:

```
{
  "data": {
    "related_extracts": [
      {
        "kind": "domain",
        "value": "robohelptestng.biz"
      },
      {
        "kind": "ipv4",
        "value": "195.22.28.199"
      },
      {
        "kind": "ipv4",
        "value": "188.200.164.50"
      }
    ]
  }
}
```



To examine the JSON data structure of an entity:

- Go the entity detail pane, and then click the **JSON** tab.

Alternatively:

- On the selected entity detail pane, click **Actions > Export > JSON** to save the entity in JSON format.

- **Value matches:** define a regex to specify the data pattern the rule should apply to search for the desired content. If you specify multiple values, enter one value per line.

The regex tells the rule *what* to look for at the location the JSON path points to.

Think of it as the front of the postcard you're sending to a friend, the side with the picture of a very stereotypical landscape that can match a number of actual places.

This field supports only **Elasticsearch regular expression syntax**

(<https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-regexp-query.html#regexp-syntax>).

The main peculiarities of the Elasticsearch query regex syntax are:

- Anchors (^ and \$) are implied at the beginning and at the end of the regex. You do not need to include them in the regex you input.
- If you insert explicit anchor characters in the **Value** field, they are interpreted as literal values.
- You need to escape special characters (. ? + * | { } [] () " \). To escape a special character, prepend a backslash \ to it. Example: \{ \}

✓ At this moment, Elasticsearch regular expression syntax **optional operators**

(https://www.elastic.co/guide/en/elasticsearch/reference/current/query-dsl-regexp-query.html#_optional_operators) are not supported.

- **Source:** from the drop-down menu select an incoming feed or an enricher to use as a data source for the rule.
- **Derivation:** from the drop-down menu select **Original** or **Derived**.
- **Levels:** from the drop-down menu select **1** or **2**.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save observable rules

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new:** saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate:** saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Derivation and levels

Derivation — **Original** vs **Derived** observables — and levels — level **1** and level **2** observables — work together to make it easier to act efficiently on observables and to use them to trigger follow-up actions in your prevention/detection toolchain.

The platform can flag observables to automate processes such as:

- Add potentially malicious threats to a prevention and/or a detection system;
- Exclude non-malicious observables that do not represent a potential threat for the organization.

Rules handle the flags, and they can initiate actions on observables; for example, routing them to a prevention and/or a detection system, or marking them as ignorable and filter them out to reduce unwanted data noise.

Observable

An observable records a distinct bit of information: it can be an item such as an IP address, a city name, or a hash; or an action such as the creation of a registry value, or a file deletion.

It is atomic: the information it contains is complete and meaningful, but it cannot be split into smaller components without losing meaning and value.

It is factual: it records a bare fact as is, with no additional context.

Observable inside structured text

Observables inside structured text are typically inside the fields of CybOX objects. Observables whose values are retrieved from CybOX fields usually have higher intelligence value, and they are more relevant because they are directly related to the parent STIX entity they refer to.

For example, you may decide to act on these observables by instrumenting your prevention/detection components to automatically block or blacklist them.

Example:

- A URI in a CybOX `URIObj` field: *http://www.evil.com/big/info.php*

Observable inside unstructured text

Observables inside unstructured text are typically not included in CybOX objects. They can be inside STIX fields such as headers, titles, descriptions, where they are included for reference. Observables whose values are retrieved from STIX fields usually have lower intelligence value, and they are less relevant because they are indirectly related to the source STIX entity they belong to.

For example, you may decide to ignore these observables, or to request further analysis to better assess their threat potential.

Example:

- A URI in a STIX `Reference` field: *http://www.evil.com/big/info.php*

Original observable

Original observables hold data that is extracted from the original source and stored as is.

Example:

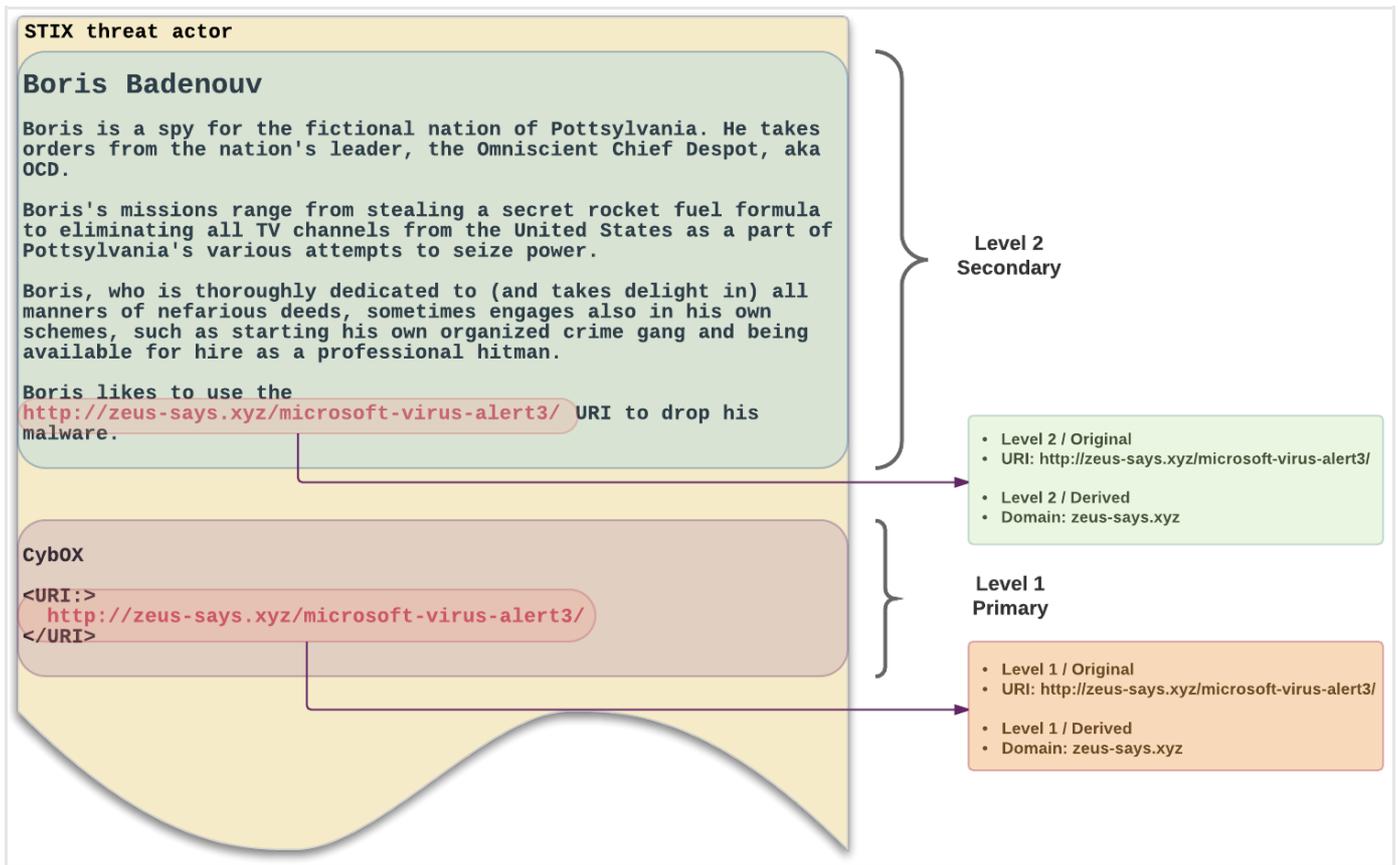
- The original data is a URI in a CybOX or a STIX field: *http://www.evil.com/big/info.php*
- The original observable stores the original data as is: *http://www.evil.com/big/info.php*

Derived observable

Derived observables hold a data subset of the original observable data. The original observable is analyzed to look for specific patterns. When matching data is found, it is extracted and saved to a distinct, derived observable that holds the analyzed data.

Example:

- The original observable is a URI: *http://www.evil.com/big/info.php*
- The extracted derived observable is a domain name: *evil.com*



Original + level 1

Derivation	Original
Level	1

- **Original / 1**: the extracted data is directly retrieved as is from a CybOX object embedded in a STIX indicator.
- **Original**: the value is extracted as is, that is, the observable holds the actual value found in the CybOX object. For example, a URI value extracted from:

```
<URIObj:Value condition="Equals">http://x4z9arb.cn/4712</URIObj:Value>
```

- **1**: the extracted data is inside a CybOX object. For example, a URI in a CybOX object embedded in a STIX indicator.

When the platform flags an observable as **Original / 1**, it handles it as follows:

- It assigns the observable an initially *low confidence maliciousness* level.
- It flags it as *level 1* extracted data to indicate that it originates from a CybOX object, it is directly related to its parent STIX entity, and it is probably relevant.
- It marks it as a potential threat that needs to be added to a detection and/or prevention system.

Original + level 2

Derivation	Original
Level	2

- **Original / 2**: the extracted data is directly retrieved as is from a STIX field, not from a CybOX object.
- **Original**: the value is extracted as is, that is, the observable holds the actual value found in the STIX field. For example, a URI value extracted from:

```
<stixCommon:Reference>https://technet.microsoft.com/library/security/2887505</stixCommon:Reference>
```

- **2**: the source of the extracted data is a value inside a STIX field, not a value inside a CybOX object. For example, a URI in a STIX field like a header, a title, or a reference.

When the platform flags an observable as **Original / 2**, it handles it as follows:

- It does not assign the observable any maliciousness level.
- It flags it as *level 2* extracted data to indicate that it does not originate from a CybOX object, but from a STIX field; it is indirectly related to its source, and possibly less relevant.
- It does not mark it for inclusion in a detection and/or prevention system.

Derived + level 1

Derivation	Derived
Level	1

- **Derived / 1**: the source of the extracted data is a value inside a CybOX field.
- **Derived**: the extracted data is the result of an analysis of the original value found inside a CybOX object. For example, a domain name extracted from a URI:

```
<!-- The original observable value, in this example a URI -->
<URIObj:Value condition="Equals">http://x4z9arb.cn/4712</URIObj:Value>

<!-- The derived observable obtained from the URI, that is, a domain -->
x4z9arb.cn
```

- **1**: the extracted data is inside a CybOX object. For example, a URI in a CybOX object embedded in a STIX indicator.

When the platform flags an observable as **Derived / 1**, it handles it as follows:

- It does not assign the observable any maliciousness level.
- It flags it as *level 1* extracted data to indicate that it originates from a CybOX object, it is directly related to its parent STIX entity, and it is probably relevant.
- It does not mark it for inclusion in a detection and/or prevention system.

Derived + level 2

Derivation	Derived
Level	2

- **Derived / 2**: the source of the extracted data is a value inside a STIX field, not a value inside a CybOX object.

- **Derived**: the extracted data is the result of an analysis of the original value found inside a STIX field. For example, a domain name extracted from a URI:

```
<!-- The original observable value, in this example a URI -->
<stixCommon:Reference>https://technet.microsoft.com/library/security/2887505</stixCommon:Reference>

<!-- The derived observable obtained from the URI, that is, a domain -->
technet.microsoft.com
```

- **2**: the source of the extracted data is a value inside a STIX field, not a value inside a CybOX object. For example, a URI in a STIX field like a header, a title, or a reference.

When the platform flags an observable as **Derived / 2**, it handles it as follows:

- It does not assign the observable any maliciousness level.
- It flags it as *level 2* extracted data to indicate that it does not originate from a CybOX object, but from a STIX field; it is indirectly related to its source, and possibly less relevant.
- It does not mark it for inclusion in a detection and/or prevention system.

Edit rules

To edit an existing observable or entity rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Entity** or **⚙ > Rules > Observable**.
- On the **Rules** page, go to **Rules > Entity** or to **Rules > Observable**, and then click the row corresponding to the rule you want to modify.
- On the entry detail pane, click **Actions > Edit** to go to the form where you can modify the selected rule.
- Enter your changes as needed.
- Click **Save** to store your changes, or **Cancel** to discard them.

Alternatively:

- On the rule overview page, click the **ⓘ** icon corresponding to the rule you want to edit.
- From the drop-down menu select **Edit** to go to the form where you can modify the selected rule.
- Enter your changes as needed.
- Click **Save** to store your changes, or **Cancel** to discard them.

Delete rules

To delete an existing observable or entity rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Entity** or **⚙ > Rules > Observable**.
- On the **Rules** page, go to **Rules > Entity** or to **Rules > Observable**, and then click the row corresponding to the rule you want to delete.

- On the entry detail pane, click **Actions > Delete**.
- On the pop-up confirmation dialog, confirm your choice.
- The rule is removed from the list.

Alternatively:

- On the rule overview page, click the **⋮** icon corresponding to the rule you want to delete.
- From the drop-down menu select **Delete**.
- On the pop-up confirmation dialog, confirm your choice.
- The rule is removed from the list.

To disable an active rule, follow the same procedure but instead of selecting **Delete**, from the context menu select **Disable**.

Filter rules

On the **Rules** page you can see table format overviews of the existing observable and entity rules.

You can narrow down the displayed results by clicking one or more quick filters above the table view to select and filter by specific:

- **Source**: select the incoming feed(s) and enrichers used as data sources for the rules.
- **Show**: select if you want to display only **Enabled** or **Disabled** rules.
- **Classification**: select if you want to display only **Malicious** rules, only **Safe** rules, only **Ignore** rules, or any combination of these options.
This option is available only for observable rules.

Example

For example, let's assume we want to apply an observable rule that zeroes in on ipv4 IP address observables. We want the rule to target IP address observables only when they are included in a sighting.

The criteria we set for the rule are:

- **Entity types**: *Sightings*
- **Observable types**: *Ipv4*
- **Paths**: *related_extracts.value*
- **Value matches**: *(.+)*abc.com*
- **Source**: in this example, we want the rule to be enabled for all incoming feeds. Therefore, we do not set this condition.

Create observable rule

Rule name *

Enabled

Criteria selection

Observables should match ALL of the following conditions:

+ Condition

- Entity types
- Observable types
- Paths
- Source
- Value matches
- Derivation
- Link types filter

The path matching the specified pattern points to the ipv4 values in the second and third members of the following array:

```
{
  "data": {
    "related_extracts": [
      {
        "kind": "domain",
        "value": "robohelptestng.biz"
      },
      {
        "kind": "ipv4",
        "value": "195.22.28.199"
      },
      {
        "kind": "ipv4",
        "value": "188.200.164.50"
      }
    ]
  }
}
```

The array contains the observables related to the sighting, which can have a JSON data structure like this:

```
{
  "alternative_versions": [],
  "attachments": [],
  "created_at": "2016-06-03T10:20:21.515918+00:00",
  "created_by": null,

  "data": {
    "confidence": {
      "type": "confidence",
      "value": "High"
    },

    "description": "Sinowal trojan identified to inform robohelptesting.biz|195.22.28.199 from
188.200.164.50",
    "impact": "High",

    "raw_events": "{\"trojanfamily\": \"Sinowal\", \"_geo_env_server_addr\": {\"postal_code\":
\\\"1300-125\\\", \"latitude\": 38.7167, \"region_code\": \"14\", \"longitude\": -9.1333, \"path\":
\\\"env.server_addr\\\", \"asn_name\": \"ClaraNET LTD\", \"asn\": 8426, \"region\": \"Lisboa\",
\\\"country_code\": \"PT\", \"netmask\": 24, \"city\": \"Lisbon\", \"country_name\": \"Portugal\",
\\\"ip\": \"195.22.28.199\"}, \"_geo_env_remote_addr\": {\"postal_code\": \"3430\", \"latitude\":
52.0148, \"region_code\": \"09\", \"longitude\": 5.1004, \"path\": \"env.remote_addr\",
\\\"asn_name\": \"KPN B.V.\", \"asn\": 1136, \"region\": \"Utrecht\", \"country_code\": \"NL\",
\\\"netmask\": 24, \"city\": \"Nieuwegein\", \"country_name\": \"Netherlands\", \"ip\":
\\\"188.200.164.50\"}, \"env\": {\"server_name\": \"robohelptesting.biz\", \"remote_port\":
\\\"3805\\\", \"remote_addr\": \"188.200.164.50\", \"request_method\": \"POST\", \"server_addr\":
\\\"195.22.28.199\\\", \"path_info\": \"/search2\", \"server_port\": \"80\"}, \"args\":
\\\"fr=altavista&itag=ody&q=ca8584331d1264912bd2e298c38eb88b%2Cdc5701fc75f672e%2C6AS2Me0aD0dEag3aS0h
kgs=1&kls=0\", \"_ts\": 1464949055, \"_origin\": \"banktrojan\", \"sd\": 1}\"",

    "related_extracts": [{
      "kind": "domain",
      "value": "robohelptesting.biz"
    },

    {
      "kind": "ipv4",
      "value": "195.22.28.199"
    },

    {
      "kind": "ipv4",
      "value": "188.200.164.50"
    }
  ]},

  "title": "Sighting robohelptesting.biz",
  "type": "eclecticiq-sighting"
},

"destinations": [],

"exposure": {
  "affected": true,
  "affected_override": null,
  "community_feed": false,
  "detect_feed": false,
  "detect_ok": false,
  "detect_override": null,
  "exposed": true,
  "prevent_feed": false,
  "prevent_ok": false,
  "prevent_override": null,
  "sighted": true
}
```

```
    signed: true
  },
  "group_id": "1632265a-ac31-49a6-9dd2-3127dcc3a39e",
  "id": "00000b8e-8b59-49b3-b04e-d3ddf540a516",
  "incoming_stix_relations": [],
  "intel_sets": [],
  "last_updated_at": "2016-06-03T10:20:21.515918+00:00",
  "meta": {
    "blob": 3586667,
    "estimated_observed_time": "2016-06-03T10:17:35",
    "estimated_threat_start_time": "2016-06-03T10:17:35",
    "incoming_feed": 237,
    "ingest_time": "2016-06-03T10:20:21.590912+00:00",
    "source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
    "source_name": "Anubis",
    "source_type": "incoming_feed",
    "title": "Sighting robohelptestesting.biz"
  },
  "outgoing_stix_relations": [{
    "alternative_versions": [],
    "attachments": [],
    "created_at": "2016-06-03T10:20:21.768998+00:00",
    "created_by": null,
    "data": {
      "key": "indicators",
      "source": "00000b8e-8b59-49b3-b04e-d3ddf540a516",
      "source_type": "eclecticiq-sighting",
      "target": "952c4de5-9abe-4904-9211-9c694d775046",
      "target_type": "indicator",
      "type": "relation"
    },
    "destinations": [],
    "exposure": {
      "affected": false,
      "affected_override": null,
      "community_feed": false,
      "detect_feed": false,
      "detect_ok": false,
      "detect_override": null,
      "exposed": true,
      "prevent_feed": false,
      "prevent_ok": false,
      "prevent_override": null,
      "sighted": false
    },
    "group_id": "1632265a-ac31-49a6-9dd2-3127dcc3a39e",
    "id": "a0040965-b3d7-4c91-b247-8d9a5d3d614b",
    "intel_sets": [],
    "last_updated_at": "2016-06-03T10:20:21.768998+00:00",
    "meta": {
      "blob": 3586667,
      "incoming_feed": 237,
      "source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
      "source_name": "Anubis",
      "source_type": "incoming_feed"
    },
    "relevancy": 1,
```

```
    "source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
    "workspaces": [],
    "workspaces_public": []
  }],

  "relevancy": 1,
  "source": "822a4302-0115-42bf-922d-e23ba01fb9c6",
  "type": "entities",
  "workspaces": [],
  "workspaces_public": []
}
```

Observables

View matching observables

When the **Action** for an observable rule is **Ignore**, the platform does not execute any actions on the matching observables. If you want to delete them, you need to manually initiate the action.

You may wish to inspect the matching ignored observables before deleting them. You can do so on the **Matches** tab on the rule detail pane.

To view observable matches for a rule, do the following:

- On the top navigation bar click **⚙ > Rules**.
- On the **Rules > Observable** page click the row corresponding to the rule whose matches you want to view to display the rule detail pane.
- Click the **Matches** tab.

× Ignore city: "San Francisco"

DETAILS
MATCHES
HISTORY

DISABLE
ENABLED

RUN NOW

Last updated	Status	Packages	
17.08.2017 19:36	✔ SUCCESS		

Criteria selection

Extract types	city
Values	San Francisco

Action

- Delete all matching observables

Created at 17.08.2017 19:36

Last updated at 17.08.2017 19:36

- Delete all matching observables
- Edit
- Disable
- Delete

The **Matches** tab shows the matching observables the rule retrieved:

- **Kind:** the matching observable data type; for example, *domain*.
- **Value:** the corresponding observable data value; for example, *www.iphishyourdata.biz*.

On this tab you can or perform actions. For example:

- To view a list of all the entities that share an observable, click the desired observable name on the detail pane.
- To refresh the view, click the refresh icon on the upper-right portion of the pane.
- To edit, disable or delete the rule, or to delete all matching observables when the **Action** of the rule is **Ignore**, select an option from the **Actions** pop-up menu.

Delete matching observables

When the **Action** configuration option of an observable rule is **Ignore**, any observables matching the rule criteria can be disregarded, and they can be deleted.

To delete all observables matching an ignore action rule, do the following:

- On the top navigation bar click **⚙ > Rules**.
- On the **Rules > Observable** page click the row corresponding to the rule whose matches you want to delete to display the corresponding detail pane.
- On the bottom half of the detail pane, click **Actions**.
- From the pop-up menu select **Delete all matching observables**.
- The observables matching the rule are deleted from the platform database, as well as from the platform history.

Alternatively:

- On the top navigation bar click **⚙ > Rules**.
- On the **Rules > Observable** page click the **!** icon on the row corresponding to the rule whose matches you want to delete.
- From the drop-down menu select **Delete all matching observables**.
- The observables matching the rule are deleted from the platform database, as well as from the platform history.

Enrichment rules

Set and run enricher rules to add intelligence value to platform entities by integrating additional context, which helps you draw a more accurate map of the threat landscape under investigation.

Automatically enrich entities

To automatically enrich entities, make sure enricher tasks are active, and the necessary enrichment rules are configured.

Rules give you control over the type of information you want to retrieve or exclude, and what you want to do with it. You can assign one or more enricher sources to specific observable types. You can set multiple filters to cover usage scenarios as needed. You can then examine the returned enrichment observable data, as well as route it to other devices that enforce cyber threat detection or prevention.

Manually enrich entities

To adjust enrichment behavior to manually apply it to the entities you want to enrich, do the following:

- Open an entity in edit mode.
For example, on the top navigation bar click **Browse > Published** to display an overview of the published entities available in the platform.
- On the row corresponding to the entity you want to manually enrich, click the  icon to display the context menu.
- From the drop-down menu select **Edit**.
- At the bottom of the entity editor page click the **Manually enrich** checkbox.
A new input field with a drop-down menu becomes available.
- From the drop-down menu select one or more enrichers you want to apply to the entity.

Workflow

Add to dataset

Manually enrich

Enrichers to apply

Please select one or more options

- Select all options
- CVE Search enricher
- Censys Enricher
- Circl.lu IP's related to SSL Certificate
- Circl.lu SSL Certificate Fetcher

CANCEL SAVE DRAFT PUBLISH

- Click **Save draft** to store your changes without publishing the entity, **Publish** to release the new version of the entity including your changes, or **Cancel** to discard the changes.

Alternatively, you can manually enrich an entity by selecting it; for example, from a dataset, from **Browse** or from **Discovery**.

An overlay slides in from the side of the screen to display the entity detail pane.

- On the entity detail pane, click **Observables**.
- The **Observables** tab shows an overview of the enrichment observables the entity has been augmented with.

To manually enrich the entity observables:

- Click the  refresh icon to trigger a task run that polls all the enrichers configured for the entity.

Alternatively:

- From the **Actions** pop-up menu, select **Enrich > Enrich with all**.
- The platform polls all applicable enrichers for the entity, and it enriches all the entity observables with the retrieved data.

Malicious files detected

Ingested: 06.10.2017 9:20 Incoming feed: TAXII Stand Samples TLP Not Set

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

Type	Value	Relation	Sighted	Conn.	First seen	Maliciousness
hash-sha256	e3b0c44298fc1c149afb4c899...	Related +1		2	06.10.2017 9:20	
hash-sha256	d7a8fbb307d7809469ca9abc...	Related +1		1	06.10.2017 9:20	
file	readme.doc.exe	Related +1		1	06.10.2017 9:20	

Edit
Delete
Add to dataset
Add to graph
Create task
Export
Download original
Enrich

Enrich with all (5)
Censys Enricher
CrowdStrike Enricher
FireEye
Flashpoint AggregINT Enricher
Flashpoint Thresher Enricher

To poll a specific enricher:

- From the **Actions** pop-up menu, select **Enrich**, and then click the specific enricher whose task run you want to trigger.
- The platform polls the specified enricher for the entity, and it enriches all supported entity observables with the retrieved data.

Malicious files detected

Ingested: 06.10.2017 9:20 Incoming feed: TAXII Stand Samples TLP Not Set

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

3 selected Deselect all Enrich Add to

Type	Value	Relation	Sighted	Conn.	Business
<input checked="" type="checkbox"/>	hash-sha256: e3b0c44298fc1c149afb4c899...	Related +1		2	
<input checked="" type="checkbox"/>	hash-sha256: d7a8fbb307d7809469ca9abcb...	Related +1		1	
<input checked="" type="checkbox"/>	file: readme.doc.exe	Related +1		1	

Enrich with all (5)
Censys Enricher
CrowdStrike Enricher
FireEye
Flashpoint AggregINT Enricher
Flashpoint Thresher Enricher

To enrich only specific observables:

- On the **Observables** tab, select the checkboxes corresponding to the observables you want to enrich.
- From the **Enrich** drop-down menu, select **Enrich with all**.
- The platform polls all applicable enrichers for the entity, and it enriches the selected entity observables with the retrieved data.

Malicious files detected

Ingested: 06.10.2017 9:20 Incoming feed: TAXII Stand Samples TLP Not Set

OVERVIEW OBSERVABLES NEIGHBORHOOD JSON VERSIONS HISTORY

2 selected Deselect all Enrich Add to

Type	Value	Relation	Sighted	Conn.	Business
<input checked="" type="checkbox"/>	hash-sha256: e3b0c44298fc1c149afb4c899...	Related +1		2	
<input type="checkbox"/>	hash-sha256: d7a8fbb307d7809469ca9abcb...	Related +1		1	
<input checked="" type="checkbox"/>	file: readme.doc.exe	Related +1		1	

Enrich with all (5)
Censys Enricher
CrowdStrike Enricher
FireEye
Flashpoint AggregINT Enricher
Flashpoint Thresher Enricher

The available enricher tasks in the drop-down menu are automatically filtered to show only the applicable enrichers for the entity.

Enrichers automatically augment all the entities that accept the enricher's content type as an observable. In other words, the observable types an entity supports define the applicable enrichers an entity can use.

View enrichment rules

To view enricher rules, do the following:

- On the left-hand navigation sidebar, click **⚙️ > System jobs > Running**.
- From the drop-down menu select **Rules**.
- On the left-hand navigation sidebar click **Enrichment**.
- The **Rules > Enrichment** page shows an overview of the configured enricher rules. You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.
- To view the details of a specific rule, click an area on the row corresponding to the rule you want to examine. An overlay slides in from the side of the screen to display the rule detail pane.

Add enrichment rules

To add a new enricher rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Enrichment**.
- The **Rules > Enrichment** page shows an overview of the configured enricher rules. You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.
- Click the **+** icon.

✔ Input fields marked with an asterisk are required.

On the **Create enrichment rule** page, fill out the fields to create the new enricher rule:

- **Name:** define a name to identify the rule. It should be descriptive and easy to remember.
- **Description:** additional textual details. If you want, you can add a short description to provide more information and context.
- Click **+ Add** or **+ More** to add a filtering option.
- **Source:** from the drop-down menu select the incoming feed, enricher, or group whose entities and observables you want to augment with additional information.
- **Entity types:** from the drop-down menu select the entity types you want to enrich with additional information.

- **TLP:** from the drop-down menu select the TLP color code you want to use to filter enrichment data.
TLP (<https://www.us-cert.gov/tlp>) provides an intuitive reference to assess how sensitive information is, focusing in particular on how serious it is, and whom it should or should not be shared with.
- Click **+ Add** or **+ More** to add a new filtering option. For example, to include another incoming feed or a different entity type. A filter can take only one source and one entity type at a time, but you can set up rules with as many filters as you need.
- **Enrichers:** from the drop-down menu select one or more enrichers to apply the rule to.
When a rule is applied to one or more enrichers, it filters the enrichment data polled from the enricher source, based on the specified rule filters and criteria.
- Select the **Enabled** checkbox to enable the rule immediately after creating it.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new:** saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate:** saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit enrichment rules

To edit enricher rules, do the following:

- On the top navigation bar click **Data configuration > Rules > Enrichment**.
- The **Rules > Enrichment** page shows an overview of the configured enricher rules.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.

To edit the details of a specific rule, do the following:

- Click an area on the row corresponding to the rule you want to examine. An overlay slides in from the side of the screen to display the rule detail pane.
- On the rule detail pane click **Actions > Edit**.

Alternatively:

- Click the  icon on the row corresponding to the enricher you want to configure or modify.
- From the drop-down menu select **Edit**.



Input fields marked with an asterisk are required.

- **Name:** define a name to identify the rule. It should be descriptive and easy to remember.
- **Description:** additional textual details. If you want, you can add a short description to provide more information and context.
- **Source:** from the drop-down menu select the incoming feed, enricher, or group whose entities and observables you want to augment with additional information.
- **Entity types:** from the drop-down menu select the entity types you want to enrich with additional information.
- **TLP:** from the drop-down menu select the TLP color code you want to use to filter enrichment data.
TLP (<https://www.us-cert.gov/tlp>) provides an intuitive reference to assess how sensitive information is, focusing in particular on how serious it is, and whom it should or should not be shared with.
- Click **+ Add** or **+ More** to add a new filtering option. For example, to include another incoming feed or a different entity type.
- **Enrichers:** from the drop-down menu select one or more enrichers to apply the rule to. They are external data providers that are polled to obtain relevant enricher raw data; for example, whois lookup, reverse DNS, or GeoIP information.
- Select the **Enabled** checkbox to enable the rule immediately after creating it.
- Click **Save** to store your changes, or **Cancel** to discard them.

Delete enrichment rules

To delete an enricher rule, do the following:

- On the top navigation bar click **Data configuration > Rules > Enrichment**.
- The **Rules > Enrichment** page shows an overview of the configured enricher rules. You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.

To delete a specific rule, do the following:

- Click an area on the row corresponding to the rule you want to delete. An overlay slides in from the side of the screen to display the rule detail pane.
- On the rule detail pane click **Actions > Delete**.

Alternatively:

- Click the  icon on the row corresponding to the rule you want to delete.
- From the drop-down menu select **Delete**.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.
- The rule is deleted.

Discovery rules

Use discovery filters and rule-based searches to retrieve specific cyber threat information from selected data sources.

Apply discovery filters to view specific entities

The **Discovery** page returns an overview of selected ingested entities, after applying rule-based search queries. You can refine the results by applying one or more quick filters. They are available on the left-hand navigation sidebar:

- On the top navigation bar click **Discovery**.
- By default, quick filters are switched off.
 - To toggle quick filter visibility click
 - To toggle quick filter visibility click
- On the left-hand navigation sidebar click a filter group name to expand the corresponding sub-nodes:
 - **Entity type**: select one or more checkboxes to include in the filtered results only the specified entity types.
 - **Source**: select one or more checkboxes to include in the filtered results only the specified entity sources.
 - **TLP**: select one or more checkboxes to include in the filtered results only entities flagged with the specified TLP color codes.
 - **Date**: select a time interval to include in the filtered results only entities ingested between the specified start and end dates.
 - **Reliability**: select one or more checkboxes to include in the filtered results only entities with the specified level(s) of reliability.
 - **Discovery rules**: select one or more checkboxes to include in the filtered results only entities matching the specified rule criteria.
 - **Dataset**: select one or more checkboxes to include in the filtered results only entities belonging to the specified datasets.
The **Dataset** filter is not available when the results do not include any entities belonging to at least one dataset.

You can stack and combine filters as you need.

For example, you can create a filter to retrieve only indicators ingested from Hailataxii in the first two weeks of last month, and whose reliability flag is either A (completely reliable) or B (usually reliable).

Apply discovery rules to retrieve specific entities

The **Discovery** service is a rule-based feature looking for cyber threat information that satisfies specific search criteria. You define the search criteria in a search query. The query sets the scope for the discovery rule. If you want, you can further restrict the discovery rule context by selecting one or more workspaces and/or workspace types.

Query task execution is capped: the response can return max. 500 matches.

In the platform discovery rules work like configurable, specialized intel fetchers:

- Configurable because you can define discovery rules as necessary.

- Specialized because the rules use search queries to focus on a specific search scope.

When you execute a discovery rule for the first time, it runs incrementally as a provider: the first run returns matching data, up to a maximum of 500 entities, *since the beginning of time*; that is, there is no start time setting to limit the discovery scope to a specific starting point in the past.

Following runs execute the specified query starting from the previous successful run, and they discover only entities added since the previous successful execution of the same rule. Repeated runs return all discovered entities since the previous successful execution of the same query.

If you want to run a discovery task without this temporal constraint, you need to create a new discovery rule.

Editing a rule does not affect this behavior. If you want a discovery query to go through all available data since the beginning of time, you need to create a new rule, and then you need to run it for the first time.

You can also edit a discovery rule, and then click **Save and re-run for all time**.

This option saves any changes, resets the execution time counter, and then it runs the rule task without applying any time constraint.

The run returns matching data for the rule, up to a maximum of 500 results, *since the beginning of time*; that is, there is no start time setting to limit the discovery scope to a specific starting point in the past.



- When a rule is active, it automatically runs every 15 minutes.
- Query task execution is capped: the response can return max. 500 matches.
- Discovery search queries use the **Elasticsearch query syntax** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html>).

View discovery rules

To view a list of all saved discovery rules, do the following:

- On the top navigation bar click **⚙️ > Rules > Discovery**.
- **Rules > Discovery** shows an overview of the existing discovery rules. You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.

Create discovery rules



Input fields marked with an asterisk are required.

To create a new discovery rule, do the following:

- On the top navigation bar click **⚙️ > Rules > Discovery**.
- Click the **+ Rule** button.

- Fill out the **Rules > Discovery > Create** form with the necessary details to create the new rule:
 - **Name:** enter a name to describe the rule. It should be descriptive and easy to remember.
Example: *China or Russia, 1 year till now*
 - **Description:** enter a short description to briefly explain what the rule does, its purpose, and the type of data it looks for.
Example: *Discovers any indicator data types having either "China" or "Russia" as a tag, and whose creation date falls in the range "one year ago until now".*
 - **Search query:** the search query you want to run when executing the rule. It should do what you explain in the rule description field. Search queries for discovery rules and rules in general use the **Elasticsearch query syntax** (<https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html>).
Example: `data.type:indicator OR entity.tags:China OR entity.tags:Russia AND created_at:[now-1y TO now]`
 - **Correlated workspaces:** you can select one or more workspaces to focus the search only on those entities that are associated with the selected workspaces. To remove a selection from the input field, click the **✕** icon corresponding to the item(s) you want to remove.
Example: *IOCs originating in China and Russia*
 - **Correlated workspaces types:** if you want, you can specify one or more workspace types to focus the search only on those entities that are related to all workspaces of a specific type. To remove a selection from the input field, click the **✕** icon corresponding to the item(s) you want to remove.
Example: *Topic*
 - **Enabled:** select or deselect this checkbox to enable or disable the rule.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new:** saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate:** saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit discovery rules

To edit a rule, do the following:

- On the top navigation bar click **⚙ > Rules > Discovery**.
- On the rule overview, click the row corresponding to the rule you want to modify.
- An overlay slides in from the side of the screen. It displays detailed rule information in a flash-card format.
- On the rule detail view, select **Actions > Edit**.
- On the **Rules > Discovery > Edit** form, you can change the field inputs as appropriate.
- Click **Save** to store your changes, or **Cancel** to discard them.

Alternatively:

- On the top navigation bar click **⚙ > Rules > Discovery**.

- On the rule overview, click the  icon on the row corresponding to the rule you want to modify.
- On the **Rules > Discovery > Edit** form, you can change the field input as appropriate.
- Click **Save** to store your changes, or **Cancel** to discard them.



You can also edit a discovery rule, and then click **Save and re-run for all time**.

This option saves any changes, resets the execution time counter, and then it runs the rule task without applying any time constraint.

The run returns matching data for the rule, up to a maximum of 500 results, *since the beginning of time*; that is, there is no start time setting to limit the discovery scope to a specific starting point in the past.

Delete discovery rules

To delete a rule, do the following:

- On the top navigation bar click  > **Rules > Discovery**.
- On the rule overview, click the  icon on the row corresponding to the rule you want to delete.
- From the pop-up context menu, select **Delete**.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.
- The discovery rule is deleted.

Enable and disable discovery rules

To manually enable and disable an existing rule, do the following:

- On the top navigation bar click  > **Rules > Discovery**.
- On the rule overview, click the row corresponding to the rule you want to run manually.
- An overlay slides in from the side of the screen. It displays detailed rule information in a flash-card format.

On the **Details** tab you can enable and disable the rule.

- If the rule is disabled:
 - Click **Enable**.
 - The button name changes to **Enabled** to notify that the rule is active.
 - A pop-up dialog asks you whether you want to run the rule right away.
 - A notification message confirms enabling the rule.
- If the rule is enabled:
 - Click **Disable**.
 - The button name changes to **Disabled** to notify that the rule is inactive.
 - A notification message confirms disabling the rule.

Manually run discovery rules

You can bypass automatic execution and decide to manually run a rule, for example to test it immediately after creating it.

To manually run a rule, do the following:

- On the top navigation bar click **⚙ > Rules > Discovery**.
- On the rule overview, click the row corresponding to the rule you want to run manually.
- An overlay slides in from the side of the screen. It displays detailed rule information in a flash-card format.
- On the **Details** tab, either click the **Run now** button, or select the **Actions > Run now** menu option.

After completing the run, you can review the outcome on the **Details** tab:

- Under the **Status** column you can check the execution outcome.

 Started	The task run has been initiated, it has been added to the queue, and it is waiting to be executed.
 Success	The task run completed correctly.
 Error	The task run failed. Click the status icon to view an error message and a traceback with more details about the failure. This information can be helpful to troubleshoot the issue.

- Under the **Results** column you can see whether the discovery action yielded any new results matching the rule criteria.

Workspaces

Private and public workspaces are containers that help you structure and organize objects under investigation, tasks, and collaboration efforts with team members and colleagues.

As your daily tasks and your workload expand and become more complex, you may want to create some structure to avoid drowning in a swamp of cryptic scribbled notes and dangling post-it messages. Workspaces help you organize your workload to keep it manageable and efficient.

Workspaces are containers that hold things neat and tidy. They enable you to label and categorize your work, to isolate subsets of tasks and objects so that you can zero in on them with more focus and less noise, and to collaborate with team members and other colleagues by exchanging information, calls for action, requests, and so on.

Workspaces provide an easy access to the tools and the datasets you use in your daily work:

- Threat overviews
- Datasets
- Graphs to visually examine threat relationships
- Any relevant files you may want to check or keep ready at hand for reference
- Comments and feedback from other colleagues, as well as tasks assigned to you or that you can assign to other people.

Workspace types

Workspace types are labels that help you keep your work in order. You can assign types to workspaces to clarify their purpose. This action does not affect workspace features and functionality, and you can change the workspace type at any time.

Generic

A generic workspace is a one-size-fits-all container to collect structured, semi-structured, and unstructured information without any specific focus. It can be handy as a temporary space where you store information and files that you are organizing in a more structured way.

Case

A case workspace is a structured container to organize intelligence on a case basis.

For example, you can create a case workspace to group together entities, datasets, file attachments, and any existing graphs concerning a specific cyber attack, or targeting a specific victim, or suspicious activity originating from an IP address range related to the same email address.

Team

A team workspace stresses collaboration and knowledge sharing. It is a repository for all the intelligence and the tasks related to a team. It helps organize and distribute workload among the team members, and it makes it easier to share comments, files, graphs, and other data among the members of the team.

For example, a team workspace enables you to organize and share information at team level, assign tasks to team members and keep track of progress.

Topic

A topic workspace can be as large and generic or as small and focused as you need. It can help you drill down on a specific threat you want to drill down into; or it can focus on a broader area of interest. For example, by grouping

intelligence related to prevention, detection, or to threat assessment.

Access workspaces

To access workspaces, do the following:

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
 - **Archived** to display archived workspaces that are not in use any longer.

You can sort the workspace overview either alphabetically — **Sort: Alphabetically** — or in reverse chronological order — **Sort: Last changed**.

Use the drop-down filters to show only the selected workspace types: **Generic**, **Team**, **Topic**, or **Case**.

On the overview workspaces are represented as tiles. Each workspace tile provides a flashcard-style summary of the main details:

- Workspace name;
- Date and time of the last change;
- Number of tasks associated with the workspace;
- Number of workspace collaborators;
- Number of saved graphs in the workspace;
- If the workspace is public — any platform user can access it — or personal — only the workspace creator and the workspace collaborators can access it.

Create a workspace

To create a workspace, do the following:

- On the top navigation bar click **Intelligence > All intelligence > Workspaces** .

Alternatively:

- On the top navigation bar click **Workspaces**, and then click the **+ Workspace** button.
- On the **Workspaces > Create** form fill out the input fields to define the new workspace as necessary.

✓ Input fields marked with an asterisk are required.

- **Name**: enter a name to designate the new workspace. It should be descriptive and easy to remember.
Example: *B-R5RB bloodbath*
- **Type**: from the drop-down menu select a workspace type to clarify the purpose and the scope of the workspace.
Example: *Case*

- **Contact info:** enter here the details of a workspace collaborator that can act as the main contact person for the workspace. For example, the designated contact person can be the workspace creator or a team leader.
Example: *Lazarus Telraven*
- **Description:** enter a short description outlining the purpose and scope of the workspace, any specific objects or topics it focuses on, and any relevant information to provide a summary overview of the workspace.
The content of this field is displayed on the workspace **Overview** tab under the **Short description** header, and it is visible only to the workspace collaborators.
Example: *B-R5RB was truly a bloodbath*
- **Public description:** enter a short description of the workspace that is visible to all platform users. It can be the same as **Description**, or a different one.
The content of this field is displayed on the workspace **Front page** tab under the **Short description** header.
Example: *B-R5RB was a walk in the park*
- **Analysis:** work notes and analysis findings to provide more context about and insight into the workspace content and its purpose. Typically, an analysis consolidates the findings of an investigation and it weaves actors, victims, incidents, events, and gathered evidence into a logical and consistent narrative.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new:** saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate:** saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit a workspace

To edit a workspace, do the following:

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
- Either on the **Workspaces > All** or the **Workspaces > Personal** overview page, browse for the workspace you want to edit, and then click the corresponding tile.
- On the left-hand navigation sidebar click **Edit details**.
- Change the content of the input fields as necessary.
- Click **Save** to store your changes, or **Cancel** to discard them.

Archive a workspace

When you do not need a workspace any longer, you can either permanently delete it, or you can archive it. Archiving a workspace makes it available for reference in read-only mode, so that you can look up comments, attachments, saved graphs, tasks, and the workspace history.

To archive a workspace, do the following:

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
- Either on the **Workspaces > All** or the **Workspaces > Personal** overview page, browse for the workspace you want to archive, and then click the corresponding tile.
- On the left-hand navigation sidebar click **Edit details**.
- Under **Archive**, click the **Archive workspace** button.
- On the confirmation pop-up dialog, click **Archive** to confirm the action.

The workspace is archived. It remains accessible in read-only mode.

Restore a workspace

If you need to gain write-access to an archived workspace to edit it, for example, to reopen a previously closed case, you can restore it.

To restore a workspace, do the following:

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
- Either on the **Workspaces > All** or the **Workspaces > Personal** overview page, browse for the workspace you want to restore, and then click the corresponding tile.
- On the left-hand navigation sidebar click **Edit details**.
- Under **Archive**, click the **Restore workspace** button.
- On the confirmation pop-up dialog, click **Restore** to confirm the action.

The workspace is restored. It becomes accessible in read-write mode.

Delete a workspace

When you do not need a workspace any longer, not even for reference, you can permanently delete it.



Warning: You cannot undo deleting a workspace.

When you delete a workspace, the action deletes the artifacts created within the workspace such as saved graphs and comments. It does not delete entities or datasets, which remain available in the platform.

To delete a workspace, do the following:

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
- Either on the **Workspaces > All** or the **Workspaces > Personal** overview page, browse for the workspace you want to delete, and then click the corresponding tile.
- On the left-hand navigation sidebar click **Edit details**.
- Under **Archive**, click the **Delete workspace** button.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.

The workspace is deleted.

Toggle between personal and public workspace

You can change the visibility of a workspace between **Public** and **Personal**:

- When a workspace is *public*, all platform users can access them in read-only mode, and workspace collaborators can access them in read-write mode.
- When a workspace is *personal*, only the workspace creator and the workspace collaborators, if any, can access it in read-write mode. Platform users who are not collaborators of the workspace cannot view it.

To toggle public and personal (private) visibility for a workspace, do the following:

- On the top navigation bar click **Workspaces**.
- On the **Workspaces > All** overview page, browse for the workspace whose visibility you want to change from either personal to public, or from public to personal, and then click it.
- On the workspace **Overview** page, click the  lock icon in the header section.
- From the drop-down, select the **Public workspace** checkbox, and then click **OK** to make the workspace public, and therefore accessible in read-only mode to all platform users.
The lock icon appearance changes from locked to open lock.
- Deselect the **Public workspace** checkbox, and then click **OK** to make the workspace personal, and therefore accessible only to the workspace collaborators.
The lock icon appearance changes from open lock to locked.

Add collaborators

Besides providing a container space for you to hold and catalog data subsets based on specific topics and areas of investigation, workspaces help you promote collaboration.

Whether a workspace is personal, that is, private and accessible only to its collaborators, or publicly open to all platform users, only workspace collaborators can perform actions such as adding and removing datasets and attachments, saving graphs, writing comments, creating tasks, and so on.

To add collaborators to a workspace, do the following:

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
- Either on the **Workspaces > All** or the **Workspaces > Personal** overview page, browse for the workspace you want to add collaborators to, and then click the corresponding tile.
- On the workspace **Overview** tab under the **Collaborators** header, click the **+ 👤** icon.
- From the drop-down menu select one or more users to add to the workspace as collaborators. The **Assign** option on the menu includes a counter with the total amount of added users.
- When you are done, click **Assign** to add the selected users as collaborators, or **Cancel** to abort the operation.

The newly added user avatars are listed under **Collaborators**.

Hover the mouse over any avatar to display the corresponding user name.

Create and review tasks

Tasks help you make intelligence actionable by creating follow-up actions to specific pieces of intelligence. Use tasks to assign and manage activities among your collaborators and to drive transparent workflows in your team.

Tasks describe follow-up actions to carry out as a response to relevant acquired knowledge. The actions can be reactive — for example, a response to a sighting or an incident — or preventive — for example, to prevent intrusions from a known threat actor.

They enable you to manage and distribute workload among collaborators, keep track of progress, and create workflows to document and drive a strategy based on actionable intelligence.

Create a task

You can access the task editor from several starting points in the platform:

From Tasks

- On the left-hand navigation sidebar click  > .
The task editor is displayed.

Alternatively:

- On the top navigation bar click  > **View all tasks**.
- On the **Tasks** overview page, click the  **Task** button.
The task editor is displayed.

From the Create new sidebar

- On the top navigation bar click  > **Workflow** > **Task**
The task editor is displayed.

From an entity overview page

You can display an entity overview page from almost anywhere in the platform.

To do so, go to an entity overview page. For example, by selecting **Intelligence** > **All intelligence** > **Browse**, **Production**, **Discovery** or **Exposure** on the top navigation bar, or by clicking **Intelligence** > **All Intelligence** > **Browse** > **Datasets** > <dataset_name> or **Data configuration** > **Incoming feeds** > <incoming_feed_name> > **Entities** tab on the feed detail pane.

- On the selected entity overview page, click the  icon on the row corresponding to the entity you want to create a task for.
- From the drop-down menu select **Create task**.
The task editor is displayed.

From an entity detail pane

You can access the entity detail pane by clicking an entity.

- To do so, go to an entity overview page. For example, by selecting **Intelligence** > **All intelligence** > **Browse**, **Production**, **Discovery** or **Exposure** on the top navigation bar, or by clicking **Intelligence** > **All Intelligence** > **Browse** > **Datasets** > <dataset_name> or **Data configuration** > **Incoming feeds** > <incoming_feed_name> > **Entities** tab on the feed detail pane.

- On the active view, browse to the entity you want to create a task for, and then click it. The entity detail pane slides in from the side of the screen.
- Click the **Actions** menu on the bottom half of the detail pane.
- From the pop-up menu select **Create task**. The task editor is displayed.

In a workspace

- On the top navigation bar click **Workspaces**.
- On the left-hand navigation sidebar click one of the following options:
 - **All** to display both public and private workspaces.
 - **Personal** to display only private workspaces.
- Either on the **Workspaces > All** or the **Workspaces > Personal** overview page, browse for the workspace you want to create a task in, and then click the corresponding tile.
- On the left-hand navigation sidebar click **Tasks**.
- On the workspace **Tasks** overview page, click the **+ Task** button. The task editor is displayed.

✓ Input fields marked with an asterisk are required.

To create a new task in the task editor, fill out the input fields to provide some details:

- **Name:** assign the task a name. It should be descriptive and easy to remember.
- **Description:** enter a short description of the task to provide a high-level overview.
- **Assigned to:** click the **+ 👤** icon, and from the drop-down menu select a platform user to assign the task to. The selected user becomes the owner of the task. You can also start typing a user name in the search box to filter only user names containing your input string.
- **Due date:** click the **📅** icon to set a deadline for the task. Time and date use the **UTC time standard** (https://en.wikipedia.org/wiki/coordinated_universal_time).
- **Stakeholders:** click the **+ 👤** icon, and from the drop-down menu select one or more stakeholders sponsoring the task. For example, a team leader or a project manager. You can also start typing a user name in the search box to filter only user names containing your input string.
- **Guidance angle:** include requirements, guidelines, and recommendations to instruct the task owner about *what* needs to be done and *why*. Providing an explanation of the expectations or the desired goals along with a rationale helps avoid misinterpretation and miscommunication. You can also add pointers to any relevant reference or context.
- **Workspaces:** from the drop-down menu select one or more workspaces to associate the task to.
- **Entities:** click **+ Add** to add one or more entities to the task. The selected entities are the objects of the task activities.
- On the entity search pop-up dialog, type a search string into the search field or select one or more checkboxes corresponding to the entities you want to associate with the task. Use the context filters to narrow down the pool of available entities to specific entities based on or more selection criteria. You can combine filters to drill down into, for example, a specific entity type ingested from a specific source in a given date range, and included in a specific dataset.

Generic searches can yield noisy results, whereas very specific searched may yield no results.

You can refine the displayed results by specifying a search string in the filter input field.

Alternatively, click one of the available filter options to select and filter by specific:

- **Entity types**

- **Source**
- **Date**
- **Datasets**
- When you are done, click the **Select** button to add them to the task.
- You can add more entities to the task at any time by clicking **+ Add** on the task editor.
- Click **Save** to store your changes, or **Cancel** to discard them..

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new**: saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate**: saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

View tasks

View tasks created by or assigned to the current user

- On the left-hand navigation sidebar click .
- A pop-up dialog lists all open tasks associated with the current user.
- A task counter hides the icon to display the total number of tasks either created by or assigned to the current user.
- Click a task on the list to display the corresponding task detail pane with an overview of all the information related to the task.
- To view all tasks either created by or assigned to the current user, on the task pop-up dialog click **View all tasks**.
- Alternatively, on the top navigation bar click **Tasks**.
The task overview page is displayed.
- To create a new task, on the task pop-up dialog or on the task overview page click **+** to open the task editor.

You can filter tasks to view only tasks created by or assigned to the current user, or only tasks in a specific status.

- On the task overview page, click the **My tasks** filter, select a checkbox, and then click **OK** on the filter drop-down menu to display the following task subsets:
 - **Created by me**: select this checkbox to view tasks created by the current user.
 - **Assigned to me**: select this checkbox to view tasks assigned to the current user.
 - Select both checkboxes to display all tasks created by *and* assigned to the current user.

View tasks by status

- On the task overview page, click the **Status** filter, select one or more checkboxes, and then click **OK** on the filter drop-down menu to display the following task subsets:
 - **Open**: select this checkbox to view open tasks that have not been started yet.
 - **In progress**: select this checkbox to view tasks that are being worked on.
 - **Done**: select this checkbox to view completed tasks.
 - **Canceled**: select this checkbox to view canceled/revoked tasks.

View task details

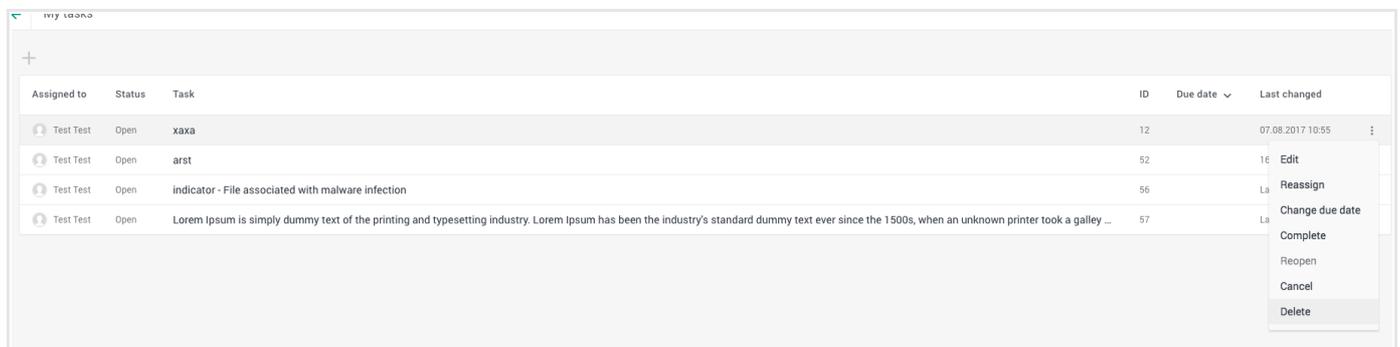
- On the left-hand navigation sidebar click .
- A pop-up dialog lists all open tasks associated with the current user.
- To view all tasks either created by or assigned to the current user, on the task pop-up dialog click **View all tasks**.
- Alternatively, on the top navigation bar click **Tasks**.
The task overview page is displayed.
- On the task overview page click anywhere on the row corresponding to the task you want to review to open the task detail pane.
- On the workspace overview page click the tile corresponding to the workspace you want to open.
- On the top navigation bar click **Tasks**.
- On the task overview page, click anywhere on the row corresponding to the task you want to add a comment to.
The task detail pane slides in from the side of the screen.
- On the task detail pane you can review existing comments related to the task, and you can add new ones to provide additional information:
 - Under **Comments**, type your message in the input field, and then press **ENTER**.

Edit tasks

To edit a task, do the following:

- Go to the task overview page.

- on the row corresponding to the task you want to modify click the  icon to select one of the following options:
 - **Edit**: opens the task editor, where you can change and update the task details.
 - **Reassign**: opens a pop-up dialog with a drop-down menu. From the drop-down menu, select the new task owner, and then click **Save** to confirm your selection.
 - **Change due date**: opens a pop-up dialog with a calendar menu. Click  to display a calendar where you can choose a new deadline for the task, and then click **Save** to confirm your selection.
 - **Complete**: available for **Open** and **In progress** tasks. Fast-forwards the task status to **Done**.
 - **Reopen**: available for **Done** and **Canceled** tasks. Reopens the selected tasks and sets the corresponding status to **Open**.
 - **Cancel**: available for **Open** and **In progress** tasks. Cancels the selected task, without deleting it. You can reopen a canceled task at any time.
 - **Delete**: deletes the selected task.
On the confirmation pop-up dialog, click **Delete** to confirm the action.
You cannot undo deleting a task.



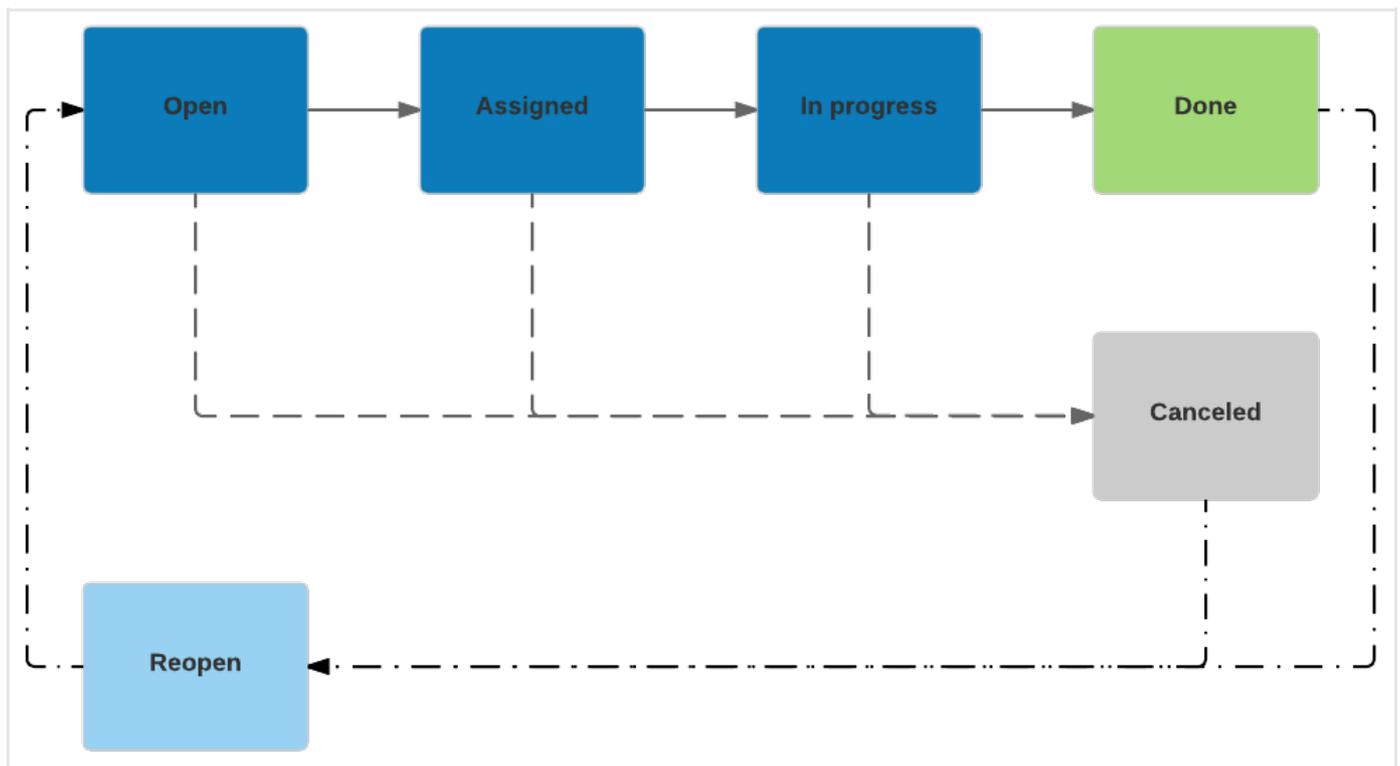
Grayed-out options in the menu are disabled for the selected item.

You can access the same action options by clicking the **Actions** the pop-up menu on the task detail pane.

Edit task status

Tasks go through different statuses during their lifecycle. Statuses give a snapshot of a task at a given point in the workflow. They allow you to monitor task progress and to take action when necessary, for example, by reassigning a task, or by changing the deadline.

Status	Description
Open	The default status a task takes upon its creation.
Assigned	The task was assigned to a (workspace) collaborator who owns it, but who has not yet started working on it.
In progress	The task owner has started working on the assigned task.
Done	The task was completed.
Canceled	The task was canceled.



A standard task status flow

To change a task status, do the following:

- Select a task and open it in edit mode.
- Under **Status**, the new status you want to assign to the task.
- Click **Save** to store your changes, or **Cancel** to discard them.

Write and review comments

Tasks and workspaces support leaving and replying to comments. Use comments to provide additional context and guidance, to ask for missing information, and to share knowledge with collaborators.

Tasks and workspaces enable you to write, edit, reply to, and delete comments.

Comments represent a way to share knowledge, to request for missing information, to discuss issues before tackling them, or to point out potential problems.

A workspace or a task comment thread can build a strong narrative providing context, reference, and relevant background information to a team of analysts working together.

Add a comment to a workspace

- On the workspace overview page click the tile corresponding to the workspace you want to open.
- On the top navigation bar click **Comments**.
The workspace comment area displays an overview of all existing comments in the current workspace.
- To add a new comment, type your message in the comment input field, and then press **ENTER**.

Edit and delete workspace comments

Comments are published under the input field in reverse chronological order (most recent first).

- Hover the mouse over a published comment line to display the icons to edit and delete it:
 - Click  to edit the selected comment.
 - When you are done, click the **Edit** button to apply your changes and to republish the updated comment.
 - Click  to delete the selected comment.
 - On the confirmation pop-up dialog, click **Delete** to confirm the action..

Add a comment to a task

- On the workspace overview page click the tile corresponding to the workspace you want to open.
- On the top navigation bar click **Tasks**.
- On the task overview page, click anywhere on the row corresponding to the task you want to add a comment to.
The task detail pane slides in from the side of the screen.

- On the task detail pane you can review existing comments related to the task, and you can add new ones to provide additional information:
 - Under **Comments**, type your message in the input field, and then press **ENTER**.

Edit and delete task comments

Comments are published under the input field in reverse chronological order (most recent first).

- Hover the mouse over a published comment line to display the icons to edit and delete it:
 - Click  to edit the selected comment.
 - When you are done, click the **Edit** button to apply your changes and to republish the updated comment.
 - Click  to delete the selected comment.
 - On the confirmation pop-up dialog, click **Delete** to confirm the action..

Work with entities and attachments

Quickly review entities in a workspace, or inspect them, load them on the graph, and analyze them.

When you work inside a workspace, you are probably focusing your activities and your attention on a specific case or scenario — for example, a recent security breach at company X by threat actor Y — on a topic — for example, investigating the latest WannaCry strands — or you may be performing your tasks in a team context.

Get an overview on the workspace dashboard

To get an idea of the threat actors and the threat targets inside a workspace, you can have a quick look at the workspace dashboard:

- On the workspace overview page click the tile corresponding to the workspace you want to open.
- On the top navigation bar click **Dashboard**.
- On the workspace dashboard, scroll to the **Entities** section listing all entities added to the workspace.
- Click an entity name to open the entity detail view, where you can carry out actions such as editing it, viewing or adding observables, or load it on the graph for analysis.

Work with entities, datasets, and attachments

Browse is a feature-rich area of the workspace with a more powerful set of tools to work with entities:

- On the top navigation bar click **Browse**.
- The **Entities** section enables you to review and modify entities, add them to a dataset, and load them on the graph to analyze them.
- You can search for specific entities in the workspace by defining a search string in the filter input field. Alternatively, click one or more quick filters to select and filter by specific:
 - Entity types
 - Data source types
 - TLP color code
 - Date ranges
 - Source reliability
 - Datasets

The number and type of available quick filters may vary: quick filters are enabled and they become available only when there are selectable values, and therefore filtering options, for the data points the filters process.

- To manage the workspace datasets, click **Datasets**.
- To manually upload file attachments, click **Uploads**.
- Drag and drop files on the upload area, or click it to browse to the location of the files you want to upload to select them.

- To download an uploaded attachment, click the file name.

The workspace upload section look and feel is very similar to the entity upload feature. However, the end result is different: any files you upload to a workspace are not ingested as entities; they are available inside the workspace as attachments, for example, to provide additional information, context, and reference.

To give more visibility to an attachment, you can pin it to the workspace front page, so that other workspace collaborators can find it more easily:

- On the row corresponding to the attachment you want to pin click the  icon, and then select **Pin to front page**.
- To remove a pinned attachment from the front page, on the row corresponding to the attachment you want to unpin click the  icon, and then select **Unpin from front page**.
- To delete an attachment from the workspace, on the row corresponding to the attachment you want to delete click the  icon, and then select **Delete**.
The attachment is immediately removed.

Work with saved graphs

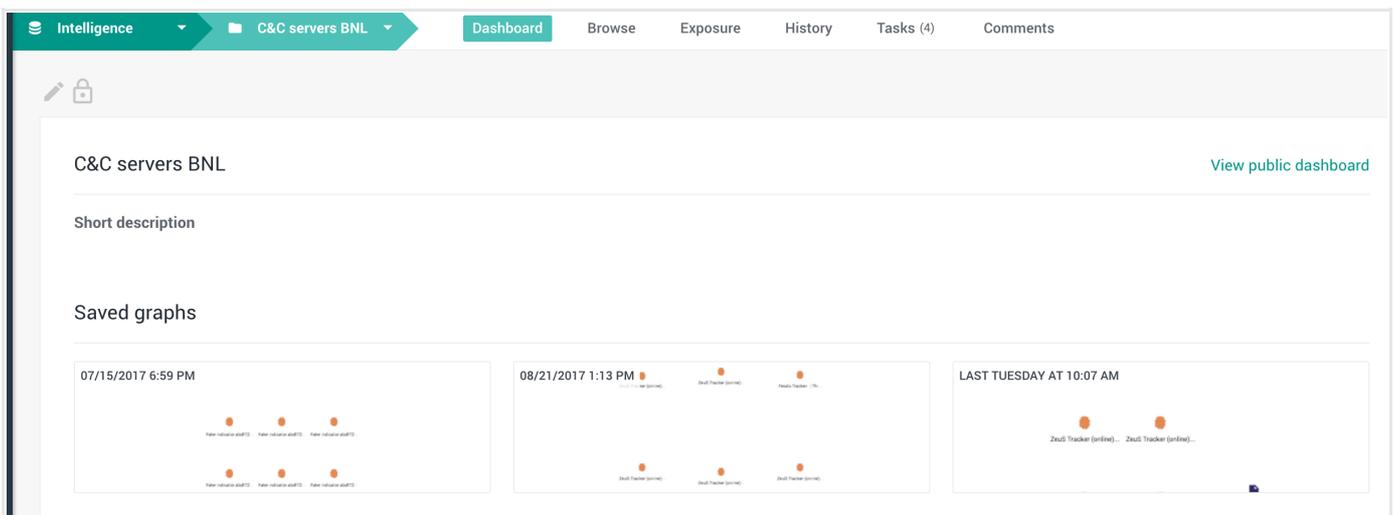
You can load entities on the graph for analysis.

After saving a graph, you can move it to a workspace.

Under **Saved graphs** you can find all the graphs that have been assigned to and moved to the workspace.

To open a graph from the **Saved graphs** section in a workspace, do the following:

- On the top navigation bar click **Browse > Saved graphs**.
- Click the  icon on the graph tile you want to open.
- From the drop-down menu select select **Load**.



The screenshot displays the workspace interface for 'C&C servers BNL'. The top navigation bar includes 'Intelligence', 'C&C servers BNL', 'Dashboard', 'Browse', 'Exposure', 'History', 'Tasks (4)', and 'Comments'. Below the navigation bar, there is a 'Short description' section and a 'Saved graphs' section. The 'Saved graphs' section contains three graph tiles. The first tile is dated '07/15/2017 6:59 PM' and shows a network graph with several nodes and edges. The second tile is dated '08/21/2017 1:13 PM' and shows a network graph with several nodes and edges. The third tile is dated 'LAST TUESDAY AT 10:07 AM' and shows a network graph with several nodes and edges. A 'View public dashboard' link is visible in the top right corner of the workspace.

Review exposure

Analyze and act on exposure and potential security issues at workspace level to pinpoint specific cases or threats.

When you work inside a workspace, you can analyze exposure at workspace level.

Exposure configuration is platform-wide. However, when you analyze it within the more specific context of a workspace, you can more easily detect any potential security issues related to the active workspace that may be threatening or affecting your environment.

To view exposure details for a workspace, do the following:

- In the active workspace, on the top navigation bar click **Exposure**.
- You can refine the displayed results by specifying a search string in the filter input field. Alternatively, click one or more quick filters to select and filter by specific:
 - Entity types
 - Workspaces
 - Datasets
 - Date ranges

The number and type of available quick filters may vary: quick filters are enabled and they become available only when there are selectable values, and therefore filtering options, for the data points the filters process.

Apart from a more limited scope, workspace exposure behaves like platform-wide exposure.

View the workspace history

The workspace history displays an overview in reverse chronological order of the actions applied to the workspace and the users who performed them.

To view history details for a workspace, do the following:

- In the active workspace, on the top navigation bar click **History**.

Click **History** to display an overview in reverse chronological order of the actions that the workspace collaborators carried out and saved in the workspace since its creation.

This reference view enables you to inspect *what happened* to the entity (the action), *who did it* (the user), and *when it happened* (the date and time).

Manage users

Configure and manage users to control access to platform resources.

EclecticIQ Platform manages and controls resource access and consumption by defining access profiles with the following characteristics:

- **Users:** individual platform consumers who can access the platform by signing in with their designated account credentials: user name and password.
Example: *mhamilton*
- **Roles:** the expected functions of users. Roles define typical tasks and behaviors of the functions they are related to.
Example: *Team Lead*
- **Permissions:** rules and policies constraining user scope. Permissions delimit scope by defining:
 - *What* users are authorized to do.
 - *Where* they can carry out the allowed actions, by setting areas in the platform where users can perform the tasks and behaviors that comply with their assigned roles.
Example: *modify files*
- **Groups:** several users brought together under a common umbrella, sharing the same roles and permissions.
Example: *Threat Analysts*

Write access to user profiles depends on the permissions assigned to a user role. Usually, admin roles include the **modify users** permission, and they have read and write access to user profiles. Non-admin roles should not need to be granted this permission: they should be able to edit their own user profiles, and they should access other user profiles in read-only mode.

View users

To view a list of platform users, do the following:

- On the left-hand navigation sidebar click **⚙ > User management**.
- The default **User management** view is **Users**, which shows an overview of the registered platform users.

By default, the overview displays only active users. To view inactive users, do the following:

- Click the filter icon on the top-left corner of the page, from the drop-down menu select **Disabled**.

To view details about a specific user, on the user overview page click anywhere on the row corresponding to the user whose profile you want to review.

The user detail pane slides in from the side of the screen.

- The default user detail pane view is **Overview**, where you can view all the configured options for the current user profile.
- Click **History** to display an overview in reverse chronological order of the actions performed on the user profile since its creation.

This reference view enables you to inspect *what happened* to the user profile (the action), *who did it* (the user who carried out the action), and *when it happened* (the date and time).

Create users

To add a new user, do the following:

- On **User management > User** click the **+** icon on the top-left corner of the page. The user editor is displayed. In the user editor define the following configuration settings:
 - **First name**: enter the user's first/given name.
 - **Last name**: enter the user's last/family name.
 - **User name**: enter the designated user name to identify the user, when signed in to the platform.
 - **Email**: enter the user's valid email address.
 - **Active**: select this checkbox to enable the user immediately after saving the newly created user profile. Active users can sign in to the platform and carry out actions, based on their user profile and their permissions.
 - **Administrator**: select this checkbox to elevate the user's role to administrator. When the checkbox is selected, the user has full administrator rights and permissions.
 - **Contact info**: the user's contact details such as address or phone number.
 - **PGP public key**: the user's **PGP public key** (<https://ssd.eff.org/en/module/introduction-public-key-cryptography-and-pgp>).
 - **Locale**: from the drop-down menu select the appropriate **locale** ([https://en.wikipedia.org/wiki/locale_\(computer_software\)](https://en.wikipedia.org/wiki/locale_(computer_software))) **settings for the user interface**.
 - **Use system timezone**: select this checkbox to override any locale-specific time zone setting with the system-defined time zone.
 - **Groups**: from the drop-down menu select one or more groups to assign the new user to.
 - Alternatively, search for a group by starting typing a group name in the autocomplete text input field.
 - To remove the user from one or more groups, remove the relevant entries by clicking the **✕** corresponding to the group you want to remove the user from.
 - **Roles**: it works like **Groups**, the only difference being that instead of adding the user to one or more groups, this option assigns one or more roles to the user.
- Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new**: saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate**: saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit users

- To edit user details, click the  icon on the row corresponding to the user profile you want to modify, and then from the drop-down menu select **Edit** to open the user editor.
- Alternatively, click anywhere on the row corresponding to the user profile you want to modify. The user detail pane slides in from the side of the screen.
- Click **Edit** at the bottom of the detail pane to open the user editor.
- Change the user details as necessary.
- Click **Save** to store your changes, or **Cancel** to discard them.



The screenshot shows a 'User management' interface with a navigation bar containing 'USERS', 'GROUPS', 'ROLES', and 'PERMISSIONS'. Below the navigation bar is a table of users. The table has columns for 'User', 'First name', and 'Last name'. The 'alexey' user row has an 'Edit' button next to it.

User	First name	Last name
admin	Admin	Test
alexey	alexey	al Edit
analyst	Analyst	Test

Edit your user profile

- To edit your user profile on the left-hand navigation sidebar click your avatar picture, and then select **My profile** to open your user profile.

Change avatar image

- To upload a custom image or to change the existing avatar image, click  or the current avatar image you want to replace.
- Browse to the location where the replacement image is stored, and double-click it to upload it as the new avatar image.

Edit your user details

- Click **Edit** at the bottom of your user profile overview to open the user editor.
- Change your user details as necessary.
- Click **Save** to store your changes, or **Cancel** to discard them.

Change your password

- Click **Change password** at the bottom of your user profile overview to open the user editor.
- On the pop-up dialog, enter the new password, and then enter it again to confirm it.
- Click **Save** to store your changes, or **Cancel** to discard them.

Change user password

To change a user password, do the following:

- Open the user editor.
- Click **Change password** at the bottom of the detail pane.
- On the pop-up dialog, enter the new password, and then enter it again to confirm it.
- Click **Save** to store your changes, or **Cancel** to discard them.

Revoke user access

To revoke a user's ability to access the platform, do the following:

- Open the user editor.
- On the user editor page, deselect the **Active** checkbox.
- Click **Save** to store your changes, or **Cancel** to discard them.

Manage groups

Configure and manage user groups to control access rights and permissions at team or department level.

EclecticIQ Platform manages and controls resource access and consumption by defining access profiles with the following characteristics:

- **Users:** individual platform consumers who can access the platform by signing in with their designated account credentials: user name and password.
Example: *mhamilton*
- **Roles:** the expected functions of users. Roles define typical tasks and behaviors of the functions they are related to.
Example: *Team Lead*
- **Permissions:** rules and policies constraining user scope. Permissions delimit scope by defining:
 - *What* users are authorized to do.
 - *Where* they can carry out the allowed actions, by setting areas in the platform where users can perform the tasks and behaviors that comply with their assigned roles.
Example: *modify files*
- **Groups:** several users brought together under a common umbrella, sharing the same roles and permissions.
Example: *Threat Analysts*

View groups

To view a list of platform user groups, do the following:

- On the left-hand navigation sidebar click **⚙ > User management > Groups**.
The **Groups** view shows the existing user groups.
- To view details about a specific user group, on the **Groups** overview page click anywhere on the row corresponding to the group you want to review.
The user group detail pane slides in from the side of the screen.
- On the user group detail pane, click **Overview** to see a list of the intelligence data sources the group, and therefore the users that belong to it, have access to. Besides the name of the data source you can see if it is an enricher, a feed, or a group, and optionally a TLP color code providing information handling and sharing guidelines.
 - To remove a data source from the list, click the **⋮** icon on the row corresponding to the data source you want to make unavailable to the group and its users, and then select **Remove**.
- Click **Users** to view a list of the users belonging to the group.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.
- To filter group users by role, from the **Roles** drop-down filter select one or more checkboxes corresponding to the roles you want to include in the filter.
- To remove the filter, deselect the checkboxes.

- Click **History** to display an overview in reverse chronological order of the actions performed on the user group since its creation.
This reference view enables you to inspect *what happened* to the user group (the action), *who did it* (the user who carried out the action), and *when it happened* (the date and time).

Create groups

To add a new user group, do the following:

- On the left-hand navigation sidebar click **⚙ > User management**.
- Under **User management > Groups**, click the **+** icon.
The user group editor is displayed.

✓ Input fields marked with an asterisk are required.

- Under **Create group**, define the following configuration settings:
 - **Name**: a descriptive name for the user group.
Example: *Fraud analysts*
 - **Description**: a short description of the user group and its purpose.
Example: *Groups together fraud analysts from the Black, Red, and Pale Fuchsia teams*
 - **Allowed sources**: click **+ Add** or **+ More** to add new rows as needed, where you can enter additional criteria.
 - **Sources**: from the drop-down menu select one or more data sources the user group and its members can access to fetch data from. The data sources can be existing incoming feeds, enrichers, as well as other user groups.
 - **TLP**: from the drop-down menu select a **Traffic Light Protocol** (<https://www.us-cert.gov/tlp>) color to filter data accordingly.
 - Click **+ Add** or **+ More** to add new rows as needed, where you can enter additional criteria.
 - **Source reliability**: from the drop-down menu select a value to filter data source reliability, so as to allow the user group to access only data from reliable sources, based on the value you set here.
 - Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new**: saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate**: saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Edit groups

- To edit user group details, click the  icon on the row corresponding to the user group you want to modify, and then from the drop-down menu select **Edit** to open the user group editor.
- Alternatively, click anywhere on the row corresponding to the user group you want to modify. The user group detail pane slides in from the side of the screen.
- Click **Edit** at the bottom of the detail pane to open the user group editor.
- Change the user group details as necessary.
- Click **Save** to store your changes, or **Cancel** to discard them.

Add users to a group

To add one or more existing platform users to a groups, do the following:

- On the left-hand navigation sidebar click  > **User management** > **Groups**.
- On the **Groups** overview page click anywhere on the row corresponding to the group you want to add users to. The user group detail pane slides in from the side of the screen.
- On the user group detail pane click **Users**, and then the **Add user** button on the top half of the pane.
- On the pop-up dialog start typing the name of an existing platform user in the autocomplete text input field.
- Click **Assign** to add the user to the group, or **Cancel** to close the dialog.

This is a handy option to add users to a group on the fly.

To add a user to one or more groups you can edit the user profile, where you can select groups and roles .

Remove users from a group

To remove one or more users from a group, do the following:

- On the left-hand navigation sidebar click  > **User management** > **Groups**.
- On the **Groups** overview page click anywhere on the row corresponding to the group you want to remove users from. The user group detail pane slides in from the side of the screen.
- On the user group detail pane click **Users**.
- On the **Users** overview click the  icon on the row corresponding to the user you want to remove from the group.
- From the drop-down menu select **Remove**.

This is a handy option to remove users from a group on the fly.

To remove a user from one or more groups you can edit the user profile, where you can deselect multiple groups and roles at once.

Edit users in a group

To edit one or more users in a group, do the following:

- On the left-hand navigation sidebar click **⚙ > User management > Groups**.
- On the **Groups** overview page click anywhere on the row corresponding to the group whose users you want to edit. The user group detail pane slides in from the side of the screen.
- On the user group detail pane click **Users**.
- On the **Users** overview click the **⋮** icon on the row corresponding to the user you want to edit.
- From the drop-down menu select **Edit**.
The user editor is displayed.
Modify the user profile as necessary.

This is a handy option to edit users in a group on the fly.

Delete groups

- To delete a user group, click the **⋮** icon on the row corresponding to the user group you want to delete, and then from the drop-down menu select **Delete**.
- Alternatively, click anywhere on the row corresponding to the user group you want to delete. The user group detail pane slides in from the side of the screen.
- Click **Delete** at the bottom of the detail pane.
- On the confirmation dialog, click **Delete** to confirm the action.
The user group is deleted from the platform.



Warning:

Before deleting a group, check that it is not an authorized group in an outgoing feed configuration. Deleting a group that is currently selected as an authorized group to access the outgoing feed content breaks the outgoing feed functionality.

If you need to remove such a group:

- First, remove it from the **Authorized group** selection in the relevant outgoing feed(s).
- Then, proceed to delete the group.

Manage roles

Configure and manage user roles and their sets of permissions to control user access to platform resources and their ability to modify them.

EclecticIQ Platform manages and controls resource access and consumption by defining access profiles with the following characteristics:

- **Users:** individual platform consumers who can access the platform by signing in with their designated account credentials: user name and password.
Example: *mhamilton*
- **Roles:** the expected functions of users. Roles define typical tasks and behaviors of the functions they are related to.
Example: *Team Lead*
- **Permissions:** rules and policies constraining user scope. Permissions delimit scope by defining:
 - *What* users are authorized to do.
 - *Where* they can carry out the allowed actions, by setting areas in the platform where users can perform the tasks and behaviors that comply with their assigned roles.
Example: *modify files*
- **Groups:** several users brought together under a common umbrella, sharing the same roles and permissions.
Example: *Threat Analysts*

View roles

To view a list of the available roles in the platform, do the following:

- On the left-hand navigation sidebar click **⚙ > User management > Roles**.
The **Roles** view shows the existing roles.
- To view details about a specific role, on the **Roles** overview page click anywhere on the row corresponding to the role you want to review.
The role detail pane slides in from the side of the screen.
- On the role detail pane, click **Overview** to see a list of permissions associated with the role.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.

The permissions on the list should map the typical tasks normally associated with the role.

For example, a system administrator role should be granted a broader range of permissions and access rights than a standard user.

- Click **History** to display an overview in reverse chronological order of the actions performed on the role since its creation.
This reference view enables you to inspect *what happened* to the role (the action), *who did it* (the user who carried out the action), and *when it happened* (the date and time).

Create roles

To add a new role, do the following:

- On the left-hand navigation sidebar click **⚙ > User management**.
- Under **User management > Roles**, click the **+** icon.
The role editor is displayed.

✓ Input fields marked with an asterisk are required.

- Under **Create role**, define the following configuration settings:
 - **Name**: a descriptive name for the role.
Example: *Analyst*
 - **Description**: a short description of the user group and its purpose.
Example: *Analyzed platform data to identify and prioritize potential threats that could impact the organization.*
 - **Permissions**: from the drop-down menu select one or more permissions for the role.
Each permission grants the role the right to access a platform resource, either in read-only mode, or in write-mode.
A role should have the necessary set of permissions to perform the tasks normally associated with it.
 - Click **Save** to store your changes, or **Cancel** to discard them.

About permissions

- Permissions are associated with roles. Roles act as containers for sets of permissions defining the scope of actions of the corresponding roles.
- Permissions are predefined in the platform, and they are not editable or configurable. You can either grant them to roles, or revoke them.
- Permission names strive to be self-explanatory:
Format: *<type of action> <object of the action>*
Example: *modify entities*
- Permissions allow two types of action:
 - **modify**: a modification permission that allows write operations.
 - **read**: a read permission that grants access to data without allowing any modifications.

To get an overview of the available permissions available on the platform, do the following:

- On the left-hand navigation sidebar click **⚙ > User management**.
- Under **User management > Permissions**, the permission overview is displayed as a table, where each permission is assigned a row.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.

Edit roles

- To edit role details, click the  icon on the row corresponding to the role you want to modify, and then from the drop-down menu select **Edit** to open the role editor.
- Alternatively, click anywhere on the row corresponding to the role you want to modify. The role detail pane slides in from the side of the screen.
- Click **Edit** at the bottom of the detail pane to open the role editor.
- Change the role details as necessary.
- Click **Save** to store your changes, or **Cancel** to discard them.

Delete roles

- To delete a role, click the  icon on the row corresponding to the role you want to delete, and then from the drop-down menu select **Delete**.
- Alternatively, click anywhere on the row corresponding to the role you want to delete. The role detail pane slides in from the side of the screen.
- Click **Delete** at the bottom of the detail pane.
- On the confirmation dialog, click **Delete** to confirm the action. The role is deleted from the platform.

Manage automation users

Configure and manage dedicated users and groups for automation tasks that interact with external components or systems, such as in platform integration implementations.

It is a good idea to have one or more dedicated users and user groups, as necessary, to handle automation tasks that interact with external products or components of your system.

Automation groups bring together automation users, and they act as global controllers of the permissions the automation users require to operate.

Automation users handle automation and integration tasks such as data transmission through feeds and enrichers, or automatic entity creation as a follow-up action on a specific event.

Create an automation group



Warning: The automation group should include all the data sources — incoming feeds, enrichers, and groups — the automation users in the group need to access.

To add a new automation user group, do the following:

- On the left-hand navigation sidebar click **⚙ > User management**.
- Under **User management > Groups**, click the **+** icon.
The user group editor is displayed.



Input fields marked with an asterisk are required.

- Under **Create group**, define the following configuration settings:
 - **Name:** a descriptive name for the automation user group.
Example: *DomainTools integration automation group*
 - **Description:** a short description of the automation user group and its purpose.
Example: *Automation group for DomainTools automation users*
 - **Allowed sources:** click **+ Add** or **+ More** to add new rows as needed, where you can enter additional criteria.
 - **Sources:** from the drop-down menu select one or more data sources the automation user group and its members can access to fetch data from. The data sources can be existing incoming feeds, enrichers, as well as other user groups.
 - **TLP:** from the drop-down menu select a **Traffic Light Protocol** (<https://www.us-cert.gov/tlp>) color to filter data accordingly.
 - Click **+ Add** or **+ More** to add new rows as needed, where you can enter additional criteria.
 - **Source reliability:** from the drop-down menu select a value to filter data source reliability, so as to allow the user group to access only data from reliable sources, based on the value you set here.
 - Click **Save** to store your changes, or **Cancel** to discard them.

Save options

Besides committing current data by clicking **Save**, you can also click the downward-pointing arrow on the **Save** button to display a context menu with additional save options:

- **Save and new**: saves the current data for the active item, and it allows you to start creating a new item of the same type right away. For example, a dataset, a feed, a rule, a workspace, or a task.
- **Save and duplicate**: saves the current data for the active item, and it creates a pre-populated copy of the same item, which you can use as a template to speed up manual creation work.

Create an automation role

To add a new automation role, do the following:

- On the left-hand navigation sidebar click **⚙ > User management**.
- Under **User management > Roles**, click the **+** icon. The role editor is displayed.

✓ Input fields marked with an asterisk are required.

- Under **Create role**, define the following configuration settings:
 - **Name**: a descriptive name for the automation role.
Example: *Rule modifier*
 - **Description**: a short description of the automation role and its purpose.
Example: *This role allows a third-party product to modify platform rules.*
 - **Permissions**: from the drop-down menu select the actions the role is allowed to perform.

Alternatively:

- Start typing a permission name in the autocomplete text input field.
- Select one or more filtered permissions from the list.
- To revoke one or more permissions for the role, click the **✕** icon corresponding to the permission you want to remove, or the **✕** icon next to the drop-down arrow in the input field to remove all permissions at once.
- Click **Save** to store your changes, or **Cancel** to discard them.

About permissions

- Permissions are associated with roles. Roles act as containers for sets of permissions defining the scope of actions of the corresponding roles.

- Permissions are predefined in the platform, and they are not editable or configurable. You can either grant them to roles, or revoke them.
- Permission names strive to be self-explanatory:
Format: *<type of action> <object of the action>*
Example: *modify entities*
- Permissions allow two types of action:
 - **modify**: a modification permission that allows write operations.
 - **read**: a read permission that grants access to data without allowing any modifications.

To get an overview of the available permissions available on the platform, do the following:

- On the left-hand navigation sidebar click **⚙️ > User management**.
- Under **User management > Permissions**, the permission overview is displayed as a table, where each permission is assigned a row.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing **▲** or a downward-pointing **▼** arrow in the header indicates ascending and descending sort order, respectively.

Create an automation user

To add an automation user, do the following:

- On the left-hand navigation sidebar click **⚙️ > User management**.
- Under **User management > User**, click the **+** icon.
The user editor is displayed.

✓ Input fields marked with an asterisk are required.

In the user editor define the following configuration settings:

- **First name**: enter a name that provides a short description of the automation role and its purpose.
- **Last name**: enter a name that provides a short description of the automation role and its purpose.
- **User name**: enter the designated user name to identify the user, when signed in to the platform.
Choose a name that helps understand what the automation user does.
Example: *DomainTools integration connector*
- **Email**: an email address associated with the automation user. You can use this address to send and receive automated notifications.
- **Active**: select this checkbox to enable the user immediately after saving the newly created user profile.
Active users can sign in to the platform and carry out actions, based on their user profile and their permissions.
- **Administrator**: select this checkbox to elevate the user's role to administrator.
When the checkbox is selected, the user has full administrator rights and permissions.
- **Contact info**: n/a
- **PGP public key**: the user's **PGP public key** (<https://ssd.eff.org/en/module/introduction-public-key-cryptography-and-pgp>).

- **Locale:** from the drop-down menu select the appropriate **locale** ([https://en.wikipedia.org/wiki/locale_\(computer_software\)](https://en.wikipedia.org/wiki/locale_(computer_software))) settings for the user interface.
- **Use system timezone:** select this checkbox to override any locale-specific time zone setting with the system-defined time zone.
- **Groups:** from the drop-down menu select one or more groups to assign the new user to. Alternatively, search for a group by starting typing a group name in the autocomplete text input field.
- To remove the user from one or more groups, remove the relevant entries by clicking the ✕ corresponding to the group you want to remove the user from.
- **Roles:** it works like **Groups**, the only difference being that instead of adding the user to one or more groups, this option assigns one or more roles to the user.
- Click **Save** to store your changes, or **Cancel** to discard them.

Get the automation group meta.source ID

Platform entities include a `meta.source` property key/value pair to identify the platform group as a data source.

If you want to programmatically create entities in the platform, you need to pass a group `meta.source` ID value when you make the corresponding calls to the platform API.

Likewise, if you want to identify the platform source group an entity comes from when the platform transmits data to an external product or system, you can retrieve the `meta.source` property key/value pair.

To retrieve the correct `meta.source` ID value related to an automation group, do the following:

- Get the automation group ID.
- Pass the automation group ID to get the `meta.source` ID.

Step 1 of 2: get the group ID

To retrieve the automation group ID value you need, so that you can retrieve the `meta.source` ID you pass in the calls to the platform API, do the following:

- On the left-hand navigation sidebar click ⚙️ > **User management**.
- Under **User management**, click **Groups**.
- On the platform group overview page, click the row corresponding to the automation group associated with the data source(s) you want to use as input *and* to the automation user making the API calls.
- The action returns a URL with the following format:
`https://<platform_host>/user-management/groups?detail=<integer>`
Example: `https://platform.host/user-management/groups?detail=30`

In the example, the `detail` value is 30. This is the group ID.

You need to pass this value in a call to a specific platform API endpoint to retrieve the `meta.source` ID.

Step 2 of 2: get the group source ID

To retrieve the `meta.source` ID to make calls to the platform API to programmatically create entities, do the following:

- Make an authentication call to the platform API to validate your user credentials and to receive a Bearer token.

- Make a call to the `/api/groups/<group_ID>` endpoint:
 - Include the Bearer token in a `Bearer` header in the call.
 - Include the group ID you previously retrieved as a trailing element in the URL.
Example: `https://platform.host/api/groups/30`
- In the JSON response, look for the group object with the `"id" : <integer>` key/value pair matching the group ID you previously retrieved.
Example: `"id" : 30,`
- In the same group object, look for the `"source" : "<UUID_string>"` key/value pair.
This is the group `meta.source` ID you need to pass in API calls to programmatically create entities.

Get the automation group meta.source ID example

Get the group ID

- On the platform group overview page, click the row corresponding to the automation group associated with the data source(s) you want to use as input.

```
https://platform.example.com/user-management/groups?detail=30
```

cURL API request — fetches the group meta.source

```
$ curl -X GET
-v
--insecure
-i
-H "Content-Type: application/json"
-H "Accept: application/json"
-H "Authorization: Bearer <token>"
https://platform.host/api/groups/30

# copy-paste version:
$ curl -X GET -v --insecure -i -H "Content-Type: application/json" -H "Accept: application/json"
-H "Authorization: Bearer <token>" https://platform.host/api/groups/30
```

API response — returns the group meta.source

```
{
  // Number of returned user groups
  "count": 18,

  "data": [

    ...

    {
      "allowed_sources": [

        // Lists all allowed data sources configured in the group editor
        ...

      ],

      // Group id, same value as the 'detail=' URL param for the group
      "id": 30,

      // Group 'meta.source' ID you need to pass in API calls
      "source": "42c051f8-9f5b-4696-a629-b86c2ead955f",

      // Group 'meta.source_name', the group name defined in the group editor
      "name": "DomainTools automation group",

      "type": "groups",
      "users": [

        // Lists all users that are part of the group
        ...

      ]
    },

    ...

  ]
}
```

Authentication

The authentication mechanism is based on **JSON web tokens** (<http://jwt.io/>).

By default, the token expires 30 minutes after successfully signing in. The corresponding session is terminated, and you need to sign back in to the platform.

When human interaction is detected — for example, keystrokes or mouse activity — the token is automatically refreshed every 2 minutes. This prevents the system from signing out users who may be working or saving data at that time.

Therefore, the default maximum amount of idle time without any human interaction before being automatically signed out equals to *session token validity - 2 minutes*.

To authenticate and access the platform, do the following:

- Make a `POST` call.
- In the call, pass your authentication credentials as a JSON object to the `/auth` endpoint. The credential data is used to generate a token that is returned with the response.

You need to include the generated bearer token in the `Authorization` HTTP header with each subsequent API call. The `Authorization` HTTP header has the following format: `Authorization: Bearer <token>`

Auth request

API endpoint	/auth
Auth method	POST
HTTP headers	"Content-Type: application/json", "Accept: application/json"
API request	POST + "Content-Type: application/json" + "Accept: application/json" + { "username": "<valid_user_name>", "password": "<valid_password>" } + <platform_host>/api/auth
API response	{ "expires_at": "<expiry timestamp>", "token": "<token>" }

The following example uses cURL to authenticate:

```
$ curl -X POST
  --insecure
  -H "Content-Type: application/json"
  -d '{ "username" : "<valid_user_name>", "password" : "<valid_password>" }'
  https://platform.host/api/auth

# copy-paste version:
$ curl -X POST --insecure -H "Content-Type: application/json" -d '{ "username" : "
<valid_user_name>", "password" : "<valid_password>" }' https://platform.host/api/auth
```

Auth response

When the user name and password credential are valid, the `POST` call returns a JSON web token:

```
{
  "expires_at": "2016-03-30T12:11:40.078219+00:00",
  "token"      :
  "abHpYXQiOjE0NTkzMzI3MDAsIm4TcCI6MTQ1OTMzOTkwMCwiYWxnIjoisSFMyNTYifQ.oyY1c2VyX2lkIjolfQ.LQQ3NdUHp4s-
  QCXsXq3feI0Dy6tf5XQX9DOML1RNIzQ"
}
```

You need to include the bearer token value in each subsequent API call. You pass the token by including an `Authorization` HTTP header in the API request.

The `Authorization` HTTP header has the following format: `Authorization: Bearer <token>`

In the following example, you make a `GET` request to the `/api/` endpoint to retrieve a list of the available API endpoints and the corresponding methods:

```
$ curl -X GET
-v
--insecure
-i
-H "Content-Type: application/json"
-H "Accept: application/json"
-H "Authorization: Bearer <token>"
https://platform.host/api/

# copy-paste version:
$ curl -X GET -v --insecure -i -H "Content-Type: application/json" -H "Accept: application/json"
-H "Authorization: Bearer <token>" https://platform.host/api/
```

**Warning:****About cURL calls**

- If you make HTTPs cURL calls to the API *and* you have a self-signed or an invalid certificate, include the `-k` or the `--insecure` parameter in the cURL call to skip the SSL connection CA certificate check.
- Always append a `/` trailing slash at the end of an API URL endpoint. The only exception is `/auth`, which does not take a trailing forward slash.
- In the cURL call, the `-d` data payload with the entity information always needs to be flat JSON, not hierarchical JSON.
If you want to pass a hierarchical JSON object, include the `--data-binary` parameter, followed by the path to the JSON file, for example `@/path/to/entity_file.json`.

View the help

The online help is a click away when you need assistance with the platform.

Sometimes it might happen that you get bogged down while working with the platform or when configuring functionality. The help documentation is there to offer hands-on assistance to get you out of there and back up and running as soon and as smoothly as possible.

To open the help, do the following:

- On the left-hand navigation sidebar click  > .
- From the pop-up menu select a help area.
For example:
 - The user guide for end-users looking for help about platform features and platform behavior.
 - The installation and configuration guide for system administrators looking for help about platform installation and configuration..
 - The release notes for business and technical platform users looking for updates about the latest released features.

Set host name and timeout

Configure host name, time zone, and default user session timeout for the platform.

Configure general server settings

To configure host name, time zone, and default user session timeout for the platform, do the following:

- On the left-hand navigation sidebar click **⚙ > System settings**.
- The default system settings view is **General**, where you can view the current values of the main settings.
- Click **Edit settings**.

✓ Input fields marked with an asterisk are required.

On **Edit general settings** you can set values for the following configuration options:

- **Hostname:** enter the host name of the machine hosting the platform instance to identify it on the network. The platform host name should match an existing `server_name` literal or pattern value defined in the Nginx `nginx.conf` file. See the official **Nginx documentation** (<http://nginx.org/en/docs/>) and the **Nginx wiki** (<https://www.nginx.com/resources/wiki/>) for further details.

Example: *prod.platform.host.com*

⚠ Warning: You need to specify a valid, working host name for the platform. If you enter an incorrect platform host name or no host name at all, *outgoing feeds and TAXII links to TAXII services won't work*.

- **Timezone:** from the drop-down menu select the appropriate time zone to display local time on the platform.

i While you can set a local or a custom time zone value for the platform, the host environment needs to be consistently on **UTC time**. This includes OS, databases, as well as any other products or components that allow setting a time zone.

- **Session timeout:** enter an integer to set the user session timeout interval, that is, after how long inactive user sessions are automatically signed out.
- **Timeout unit:** from the drop-down menu select the time measurement unit for the timeout value: in *seconds, minutes, hours, or days*.
- Click **Save** to store your changes, or **Cancel** to discard them.

To remove a selection from the input field, click the **✕** icon corresponding to the item(s) you want to remove.

Delete server settings

To delete the current server settings and server configuration for the platform do the following:

- **⚙ > System settings**.
- The default system settings view is **General**, where you can view the current values of the main settings.
- Click **Edit settings**.
- Under **Delete server settings**, click the **Delete settings** button.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.

The platform server settings are deleted.

Configure the proxy

Configure one or more proxies, if necessary, to channel incoming and outbound platform traffic through them.

Configure proxy settings

If your Internet connection setup includes a proxy server, specify the configuration in this section.

- On the left-hand navigation sidebar click **⚙ > System settings > Proxy**.
- The **Proxy** view shows the active proxy settings, if they are configured.
- Click **Edit settings**.

✓ Input fields marked with an asterisk are required.

Under **Edit proxy settings**, and depending on the protocol in use — non-secure, secure, or both — under **Web proxy (HTTP) settings**, **Secure web proxy (HTTPS) settings**, or both define the following configuration settings:

- **Server**: enter the proxy server IP address.
Example: *10.0.2.148*.
- **Port**: enter the proxy server access port.
Example: *8118*.
- **Username**: enter valid user name credentials to authenticate and to receive authorization to access the resource(s).
Example: *nigeltufnel*.
- **Password**: enter valid password credentials to authenticate and to receive authorization to access the resource(s).
Example: *thesegoto11*.
- **Bypass settings for the following hosts and domains (all protocols)** : enter here any domains and/or IP addresses that should communicate without going through the proxy server.
For example, you may want to specify here local network addresses or LAN subdomains.
If you enter multiple values, separate them with a comma.
Example: *localhost, 127.0.0.1*.
- If you use a proxy for both non-secure and secure connections, and if the proxy settings for both protocols are the same, populate the **Web proxy (HTTP) settings** section first, and then select the **Keep in sync with web proxy (HTTP) settings** checkbox under **Secure web proxy (HTTPS) settings**.
- Click **Save** to store your changes, or **Cancel** to discard them.

System settings

GENERAL PROXY EMAIL LICENSE AUDIT INTEL REPORT PRIVATE KEY TRUSTED KEYS SYSTEM

Edit proxy settings

Web proxy (HTTP) settings

Server *

Port *

Proxy server requires password

Bypass settings for the following hosts and domains (all protocols) ⓘ

Add protocol

CANCEL SAVE

Update proxy settings

When you modify or update proxy settings, the following notification message is displayed:



Proxy configuration updated. The process needs to be restarted in order for these settings to be applied.

You need to restart all Supervisor-managed processes for the changes to become effective. To do so, run the following command(s):

Delete proxy settings

To delete the current proxy server settings and proxy server configuration for the platform do the following:

- > System settings > Proxy.

- The **Proxy** view shows the active proxy settings, if they are configured.
- Click **Edit settings**.
- Under **Delete settings**, click the **Delete settings** button.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.

The proxy server settings are deleted.

Configure email

Configure email to enable features such as automatic email notification sending to registered platform users.

Configure email settings

This section configures sent-from and reply-to email addresses for the platform. These email addresses are used to send automatic notification messages to platform users, and to receive message replies, where applicable.

To fully configure the platform email server to enable email-based platform features like email notifications, you also need to configure Postfix to handle email traffic.

- On the left-hand navigation sidebar click **⚙️ > System settings > Email**.
- The **Email** view shows the active sent-from and reply-to email addresses for the platform, if they are configured.
- Click **Edit settings**.

✓ Input fields marked with an asterisk are required.

Under **Edit email settings**, define the following configuration settings:

- **From email**: enter the email address used to send automatic email notifications from the platform.
- **Reply to**: enter the return email address used to receive email messages to the platform from third-parties such as users or external applications and systems.
- Click **Save** to store your changes, or **Cancel** to discard them.

The screenshot shows the 'System settings' page with the 'EMAIL' tab selected. The 'Edit email settings' form contains two input fields: 'From email *' with the value 'platform-dev@eclecticiq.com' and 'Reply to *' with the value 'platform-dev@eclecticiq.com'. Below the form are 'CANCEL' and 'SAVE' buttons. At the bottom, there is a 'Delete email settings' section with a 'DELETE SETTINGS' button.

Test email

After configuring email settings and Postfix, you can test the email configuration by sending a test email to the current user's email address.

- **⚙ > System settings > Email.**
- The **Email** view shows the active sent-from and reply-to email addresses for the platform, if they are configured.
- Click **Test email**.
- A notification message confirms sending the test email.
- Check the email address associated with the currently signed in user: the inbox should include a test message from the platform.

Delete email settings

To delete the current email settings for the platform do the following:

- **⚙ > System settings > Email.**
- The **Email** view shows the active sent-from and reply-to email addresses for the platform, if they are configured.
- Click **Edit settings**.
- Under **Delete settings**, click the **Delete settings** button.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.

The email settings are deleted.

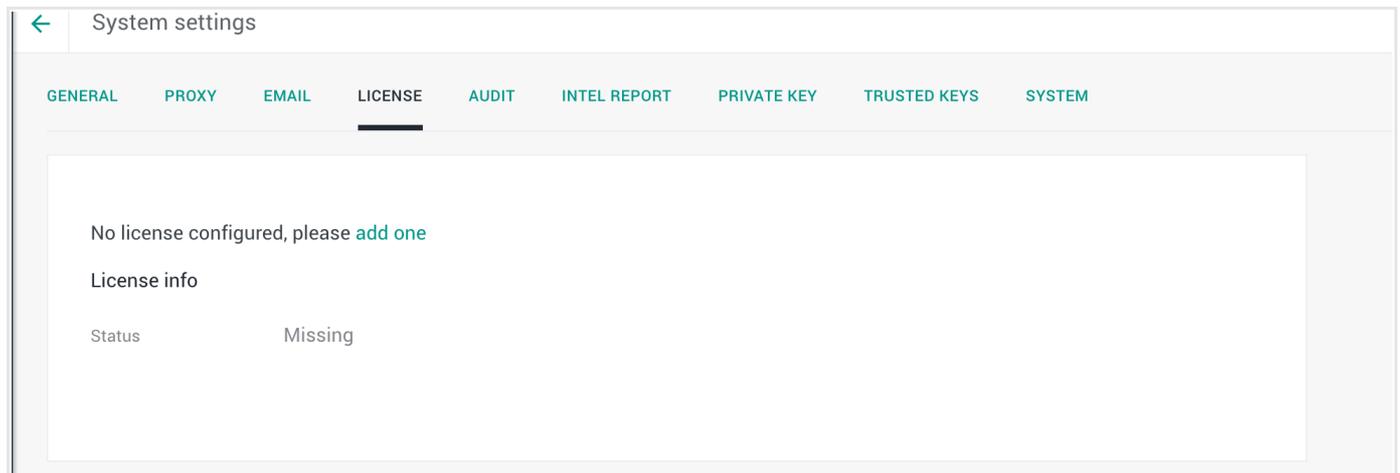
Register the license

Register the platform by entering the license details.

Update license information

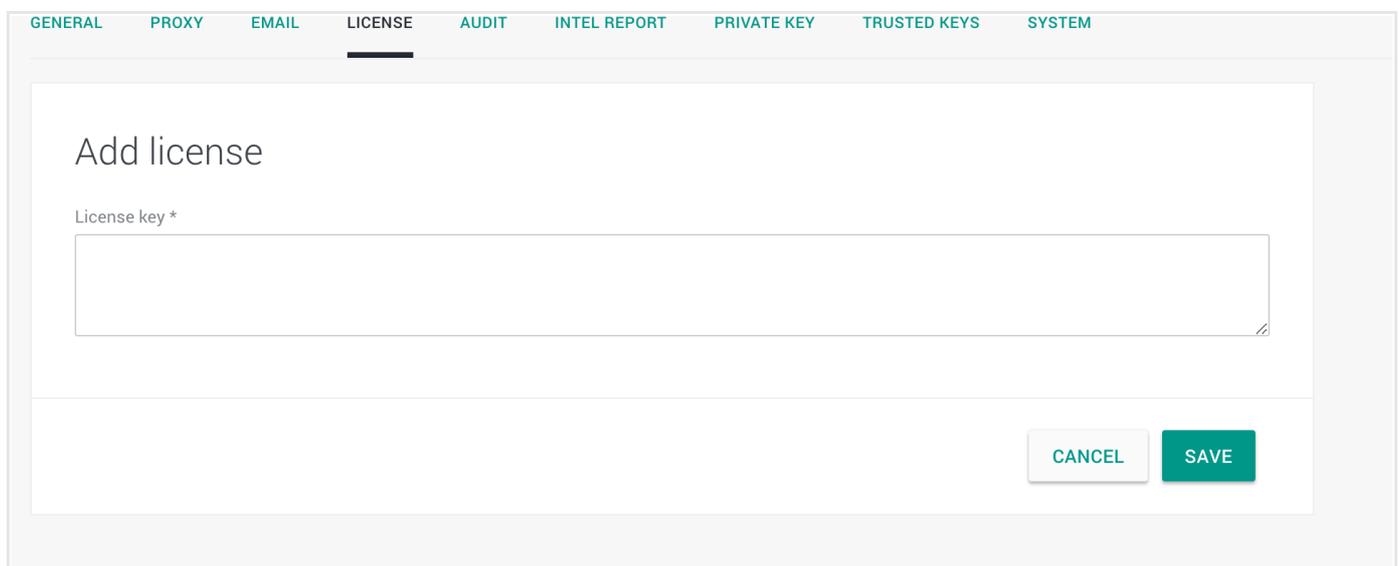
When you purchase a copy of EclecticIQ Platform you receive a license key, which you can use to register the product. To add a license key, do the following:

- On the left-hand navigation sidebar click **⚙️ > System settings > License**.
- The **License** view shows information about the active license, if it is specified.
- Click **Update license key**.
- If no license keys are registered, click **Add one** under **License key**.



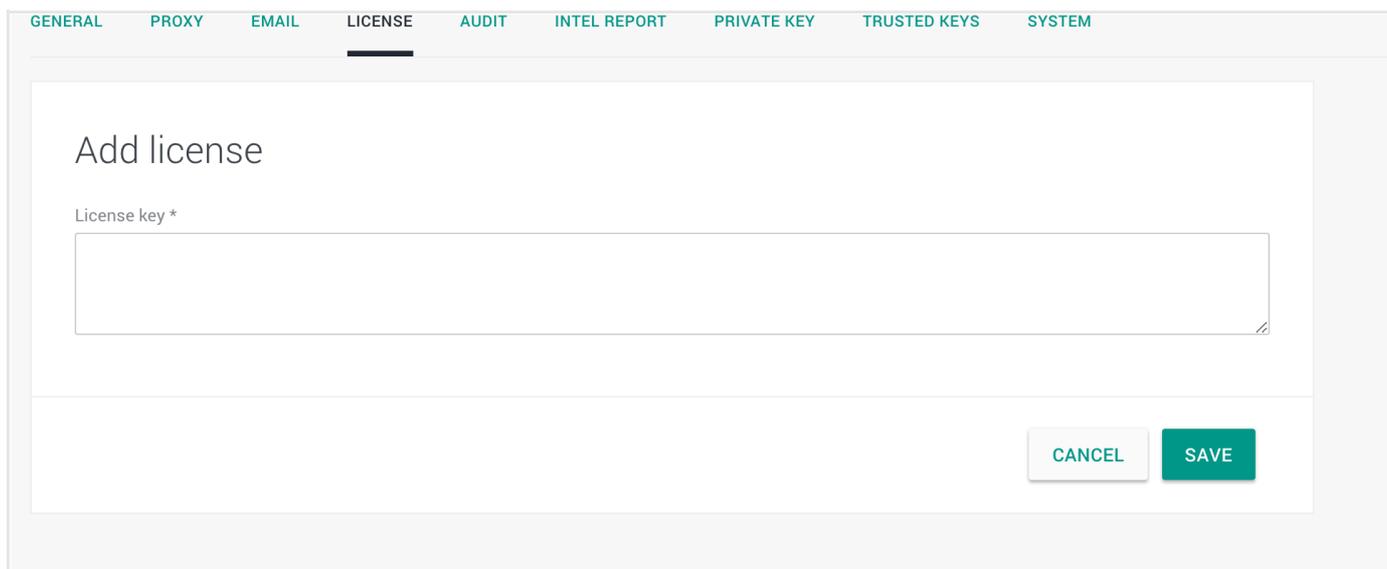
The screenshot shows the 'System settings' page with the 'LICENSE' tab selected. The main content area displays the message 'No license configured, please [add one](#)'. Below this, there is a 'License info' section with a table showing the 'Status' as 'Missing'.

- In the input field under **Edit license**, copy-paste your license details.
- Click **Save** to store your changes, or **Cancel** to discard them.



The screenshot shows the 'System settings' page with the 'LICENSE' tab selected. The main content area displays the 'Add license' form. It features a text input field labeled 'License key *' and two buttons at the bottom right: 'CANCEL' and 'SAVE'.

Valid license information populates the license information section:



The screenshot shows a web interface with a top navigation bar containing tabs: GENERAL, PROXY, EMAIL, LICENSE (selected), AUDIT, INTEL REPORT, PRIVATE KEY, TRUSTED KEYS, and SYSTEM. Below the tabs is a form titled "Add license". The form has a label "License key *" above a large text input field. At the bottom right of the form are two buttons: "CANCEL" and "SAVE".

The license type is displayed on the status bar:

Developer license

If the platform is unlicensed, a notification message is displayed on the status bar, instead of the license type.

Delete license information

To delete the current license key and license details do the following:

- **⚙ > System settings > License.**
- The **License** view shows the active sent-from and reply-to email addresses for the platform, if they are configured.
- Click **Update license key**.
- Under **Delete license key**, click the **Remove settings** button.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.

The license information is deleted.

Configure STIX

Configure STIX to enable STIX format data parsing and analysis in the platform.

Configure STIX settings

To configure STIX settings for the platform, do the following:

- On the left-hand navigation sidebar click **⚙ > STIX and TAXII**.
- The **STIX** view shows the currently configured STIX options for the platform.
- Click **Edit settings**.

✓ Input fields marked with an asterisk are required.

Under **Edit STIX settings**, define the following configuration settings:

- **Alias**: enter an alternative name, usually reader-friendly, to identify the namespace declared for the organization.
Example: *Weyland-Yutani*
Allowed characters for the alias:
 - Alphanumeric [A-Z, a-z, 0-9]
 - Underscore [_]
 - Dash [-]
 - Caseline dot/Period [.,].
The first character in the alias name needs to be either alphabetic or underscore. In other words, the STIX alias cannot start with a dash or a baseline dot.
- **Namespace**: the designated STIX namespace for your organization.
Example: *http://stix.weyland-yutani.com/stix-1*
- **Producer**: optionally, you can enter here a name to identify your organization as the producer, that is, the creator and/or the publisher of the STIX data.
- Click **Save** to store your changes, or **Cancel** to discard them.

← STIX and TAXII

STIX TAXII

Add STIX settings

Alias ⓘ

Namespace *

Producer

CANCEL SAVE

- Platform settings
- User management
- System Settings
- System jobs
- STIX and TAXII**
- System status Development
 - ❗ System partially running [VIEW](#)
 - ❗ The software is not licensed [VIEW](#)

Delete STIX settings

To delete the current STIX settings for the platform do the following:

- On the left-hand navigation sidebar click **⚙ > STIX and TAXII**.
- The **STIX** view shows the currently configured STIX options for the platform.
- Click **Edit settings**.
- Under **Delete STIX settings**, click the **Delete settings** button.
- On the confirmation pop-up dialog, click **Delete** to confirm the action.

The STIX settings are deleted from the platform.

Configure TAXII

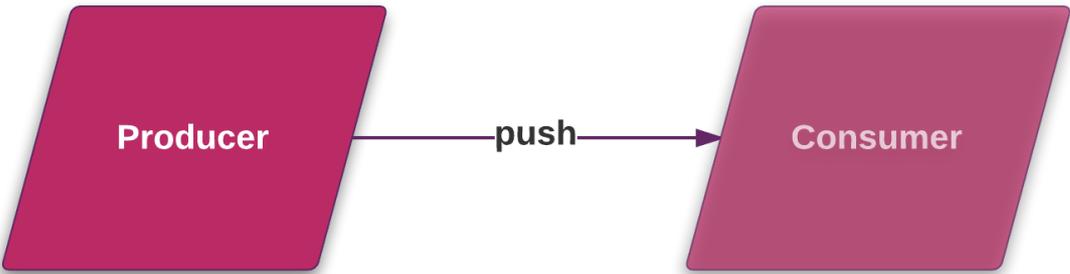
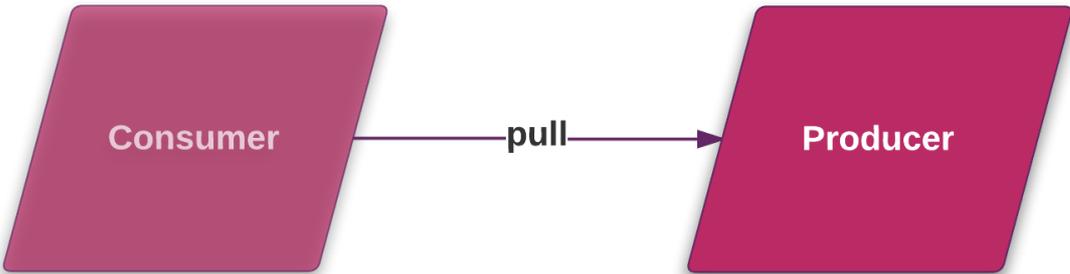
Configure TAXII to enable discovering, ingesting and distributing cyber threat intelligence using the TAXII transport type.

The TAXII server is the designated transport type for STIX data traffic.

About TAXII services

After configuring a TAXII server, you can set up TAXII services. A TAXII service is a specialized data handler that implements a specific TAXII capability.

The platform supports the following TAXII services:

Service type	Description
Collection management service	TAXII consumers can use a TAXII collection management service to request information about, subscribe to, and cancel subscriptions to TAXII data collections (TAXII outgoing data feeds and TAXII datasets). Binding: TAXII HTTP 1.0, TAXII HTTPS 1.0
Discovery service	A TAXII discovery service allows TAXII consumers to obtain information about the availability and use of TAXII services like collection management, inbox, and polling. Binding: TAXII HTTP 1.0, TAXII HTTPS 1.0
Inbox service	<p>The TAXII inbox service allows TAXII consumers to accept push messages initiated by a TAXII producer. This service can be based on a subscription model, or it can be an unsolicited payload a producer pushes to a consumer. Binding: TAXII HTTP 1.0, TAXII HTTPS 1.0</p>  <pre> graph LR Producer[Producer] -- push --> Consumer[Consumer] </pre>
Polling service	<p>The TAXII poll service allows TAXII consumers to request TAXII data collection content from a TAXII producer, usually through TAXII outgoing feeds. Binding: TAXII HTTP 1.0, TAXII HTTPS 1.0</p>  <pre> graph LR Consumer[Consumer] -- pull --> Producer[Producer] </pre>

TAXII data collections (TAXII data feeds and TAXII datasets) are a typical example of inbox and polling service content.

View TAXII services

To access an overview of the existing and configured TAXII services in the platform, do the following:

- On the left-hand navigation sidebar click **⚙ > STIX and TAXII > TAXII**.
- The **TAXII** view shows the currently configured TAXII services for the platform. You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.
- To view the configuration information of a TAXII service, click the **⋮** icon on the row corresponding to the TAXII service whose configuration you want to review.
- From the drop-down menu select **View**.
The TAXII settings page shows the current configuration of the specified service.

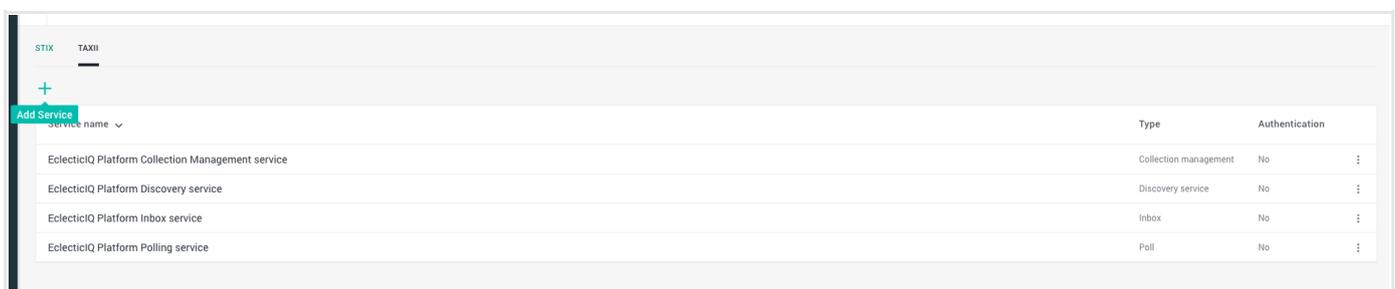
Alternatively:

- On the **TAXII** view click anywhere on the row corresponding to the TAXII service whose configuration you want to review.
The TAXII settings page shows the current configuration of the specified service.

Add a TAXII service

Configure the general options

- On the left-hand navigation sidebar click **⚙ > STIX and TAXII > TAXII**.
- On the **TAXII** view click the **⊕ Add service** icon.



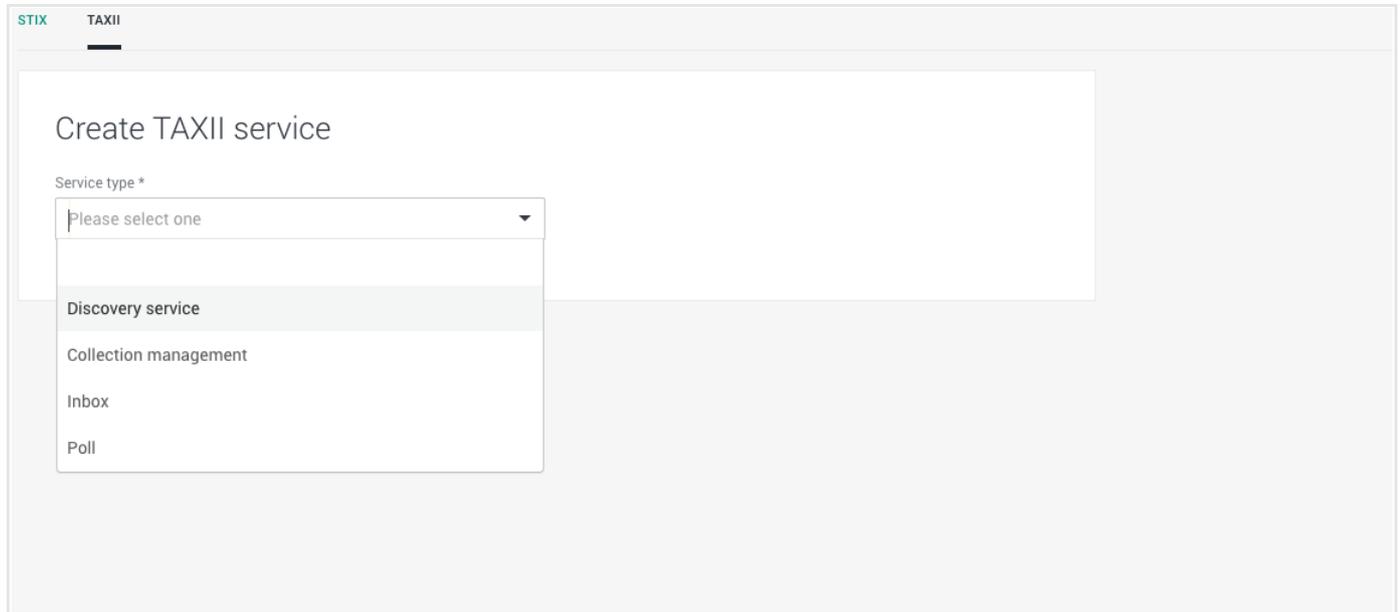
The screenshot shows the 'TAXII' view in the platform. On the left, there is a navigation sidebar with a '+' icon and a button labeled 'Add Service'. Below this is a table listing existing TAXII services. The table has columns for 'Service name', 'Type', and 'Authentication'. Each row also has a vertical ellipsis icon on the right side.

Service name	Type	Authentication
EclecticIQ Platform Collection Management service	Collection management	No
EclecticIQ Platform Discovery service	Discovery service	No
EclecticIQ Platform Inbox service	Inbox	No
EclecticIQ Platform Polling service	Poll	No

Under **Create TAXII service** define the following configuration settings:

- **Service type**: from the drop-down menu select the TAXII service type you want to add.
- **Description**: a free-text description of the service. It should be descriptive and easy to remember.
Example: *Polling from Ecorp threat db*.
- **Address**: the public endpoint the service can be reached at.
It should be relative to the configured domain name identifying the machine hosting the TAXII server.
Example: */taxii/services/polling*.

- **Protocol bindings:** the data exchange transport protocol.
Allowed values: *HTTP*, *HTTPS*.
- **Authentication required:** select the checkbox to enable user authentication, or deselect it to allow anonymous/guest access.
- Click **Save** to store your changes, or **Cancel** to discard them.



The screenshot shows a web interface for creating a TAXII service. At the top, there are tabs for 'STIX' and 'TAXII'. Below the tabs is a form titled 'Create TAXII service'. The form contains a single field labeled 'Service type *' with a dropdown menu. The dropdown menu is open, displaying four options: 'Discovery service', 'Collection management', 'Inbox', and 'Poll'. The 'Discovery service' option is currently selected and highlighted.

Configure specific options per service

Besides the general options applicable to all TAXII services, each service type has additional specific configuration options.

Discovery service

- **Advertised services:** when you set up a new *discovery service*, you need to define the TAXII services you want to advertise and make discoverable, so that consumers can access them as data sources. From the drop-down menu select one or more services.

Collection management

- **Outgoing feeds:** when you set up a new *collection management service*, you need to define the outgoing feeds you want to associate with and be managed by the service. From the drop-down menu select one or more outgoing feeds.



Warning: You first need to configure outgoing feeds before making them available through this drop-down menu.

Inbox

This service has no extra configuration options besides the common settings for all services.

Poll

- **Max result count:** if you set this option to `-1`, a poll request also counts how many entities are available in the feed(s). If you set **Max result count** to a positive integer value, and if the total amount of available entities in the feed(s) exceeds this value, a poll request informs the consumer that the total entity count in the feed(s) is higher than the set maximum result count value. You can set this option if you prefer to not disclose the total amount of entities accessible through the polling service.
- **Max result size:** this option controls pagination, that is, the number of results per page. We recommend keeping the amount of pages limited by setting a relatively large result size value. For example, `200`.
- **Outgoing feeds:** when you set up a new polling service, you need to define the outgoing feeds you want to associate with and be managed by the service. From the drop-down menu select one or more outgoing feeds.



Warning: You first need to configure outgoing feeds before making them available through this drop-down menu.

View system jobs

Monitor the platform load by checking running, completed, and aborted system jobs.

The **System jobs** view provides a complete overview of the platform workload. It enables you to review currently running, successfully completed, failed, and revoked/aborted system jobs.

Each view — **Running**, **Succeeded**, **Failed**, and **Revoked** — shows details such as the name of the job, its unique ID, the data objects it touches, and the time/date when it was executed.

You can click anywhere on the row corresponding to the job you want to inspect to open the corresponding detail pane.

View jobs

To display an overview of the platform jobs, do the following:

- On the left-hand navigation sidebar click **⚙ > System jobs**.
- The default **All** view shows all system jobs.
You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.
- You can filter jobs by clicking one of the following views:

Filter	Description
Running	Displays currently running, that is, active and not yet completed, jobs.
Success	Displays successfully completed jobs.
Failure	Displays failed jobs, that is, jobs that failed to successfully complete because one or more errors occurred.
Revoked	Displays revoked jobs, that is, jobs that were manually terminated before completion.

ID	Name	Related objects	Triggered by	Executed
0b7d1d11-0a9...	eiq.discovery.search_discovery	Discovery Test Rule		Today at 14:45
c2d2100b-b8e...	eiq.discovery.search_discovery	nameeeeeee		Today at 14:45
d2045f90-a2c...	eiq.discovery.search_discovery	RULEEEEE		Today at 14:45
41fea490-911...	eiq.discovery.search_discovery	RM 3		Today at 14:45
a74747cb-488...	eiq.discovery.search_discovery	RM2		Today at 14:45
20ca53f4-8be...	eiq.discovery.search_discovery	Rm1		Today at 14:45
458dacc4-389...	eiq.discovery.search_discovery	catch SoftTimeLimit		Today at 14:45
ff9e6d12-b78...	eiq.discovery.search_discovery	fox		Today at 14:45
e5aae21a-aft...	eiq.discovery.search_discovery	stress		Today at 14:45
d10e408f-fe3...	eiq.incoming-transports.crowdstrike_indicators	Gecko Indicator feeds	Test Test	Today at 14:41

1 - 10 of 22 560

Items per page: 10 20 50 100

- Platform settings
- User management
- System Settings
- System jobs**
- STIX and TAXII
- System status
 - System running [VIEW](#)
 - The software is not licensed [VIEW](#)

- Related objects:** shows the platform objects the system job acts on. In this context, the objects are the channels the platform uses to ingest and to publish information: incoming and outgoing feeds, discovery, and entity or observable rules.
 On the job detail pane you can click a related object name to go to the corresponding detail pane, where you can inspect the selected feed or rule in more detail.
- Triggered by:** shows the user who manually initiated a specific task run of the selected job.
 On the job detail pane you can click a triggering actor's name to go to the corresponding detail pane, where you can inspect the selected user in more detail.
- To inspect a job more closely, click anywhere on the row corresponding to the job you want to review. An overlay slides in from the side of the screen.
- The job detail pane shows job details such as the job/task name, any related platform objects such as feeds or rules that the task acts upon, and the result of the task execution.
 The **Result** section on the detail pane of failed jobs can help system administrators identify the cause of the failure by providing a descriptive error message, and a stack trace.

Terminate jobs

You can terminate a running task in one of the following ways:

- On the **Running** jobs view, click anywhere on the row corresponding to the job you want to manually terminate.
- On the job detail pane, click **Terminate**.
- On the confirmation pop-up dialog, click **Yes** to confirm the action.

Or:

- On the **Running** jobs view click the solid color, square icon on the far right on the row corresponding to the job you want to manually terminate.
 When you terminate a job in this way, no confirmation dialog is displayed. The job is terminated upon clicking the termination icon.

Audit the system

Configure audit logging to monitor the platform, and to examine system events.

Enable auditing

By default, audit logging is not configured in the platform.

To enable/disable audit logging, *add* and *set* `AUDIT_TRAIL_ENABLED` to `True` or `False` in the `/opt/eclecticiq/etc/eclecticiq/platform_settings.py` platform configuration file:

Variable	Description
<code>AUDIT_TRAIL_ENABLED = True</code>	Enables audit logging.
<code>AUDIT_TRAIL_ENABLED = False</code>	Disables audit logging.
<code>LOG_LEVEL</code>	Sets the logging level. Example: <code>'warning,eiq:info'</code>

`LOG_LEVEL` allows the following values:

- `'debug'`
- `'info'`
- `'warning'`
- `'error'`

Example: `LOG_LEVEL = 'warning'`

You can define `LOG_LEVEL` values in any of the following formats:

- String
Example: `LOG_LEVEL = 'info'`
- Key/value pairs: `<platform_component>:<log_level_value>`
Example: `LOG_LEVEL = 'eiq:warning'`
- List (comma-separated)
Example: `LOG_LEVEL = 'info,eiq:warning'`

View audit logs

You can view audit logs in the platform web-based interface, as well as in Kibana.

View audit logs in the web interface

To view audit logs in the platform web interface, do the following:

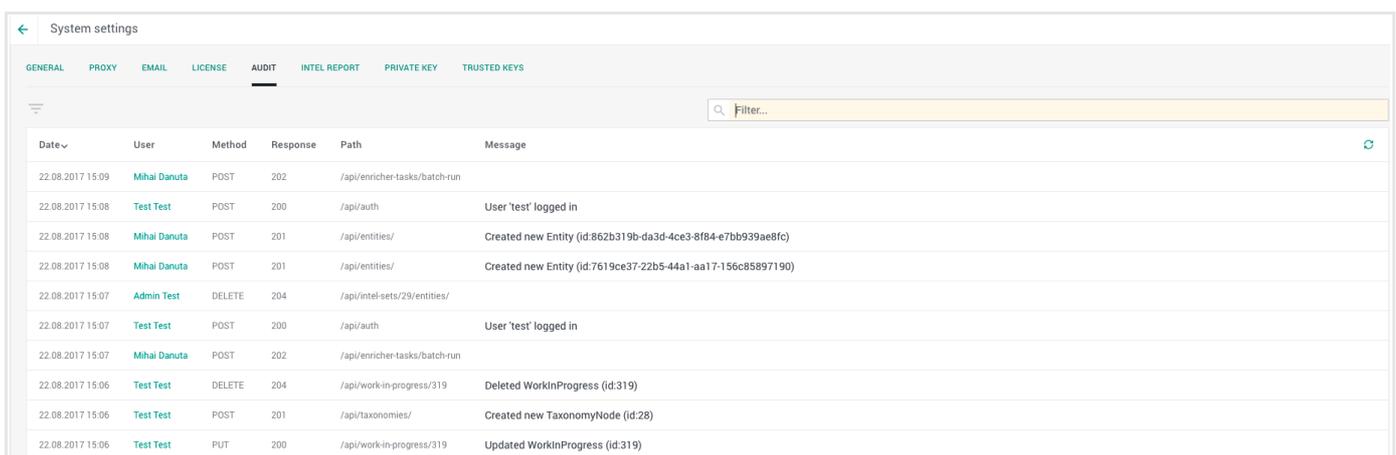
- On the left-hand navigation sidebar click  > **System settings** > **Audit**. The **Audit** tab is displayed. The audit information on this tab relies on the Elasticsearch **audit** index.
- If audit logging is enabled, and if the audit log file is populated, audit log records are returned. Use the quick filters to look for data subsets based on a date range, or on one or more specific users, HTTP methods, or HTTP response status codes.

Click  to display the quick filters.

You can sort the items on the view by column header. To do so, click the column header you want to base the data sorting on. An upward-pointing ▲ or a downward-pointing ▼ arrow in the header indicates ascending and descending sort order, respectively.

- Click the filter icon to configure and apply filters to narrow down the search scope:

Filter	Description
Date	Displays only the search result items included in the specified time range.
User	Displays only the search result items with the selected user name(s).
Method	Displays only the search result items with the selected HTTP method(s): Delete, Post, Put.
Response	Displays only the search result items with the selected HTTP response status code(s): 2xx, 4xx, or 5xx.



Date	User	Method	Response	Path	Message
22.08.2017 15:09	Mihai Danuta	POST	202	/api/enricher-tasks/batch-run	
22.08.2017 15:08	Test Test	POST	200	/api/auth	User 'test' logged in
22.08.2017 15:08	Mihai Danuta	POST	201	/api/entities/	Created new Entity (id:862b319b-da3d-4ce3-8f84-e7bb939ae8fc)
22.08.2017 15:08	Mihai Danuta	POST	201	/api/entities/	Created new Entity (id:7619ce37-22b5-44a1-aa17-156c85897190)
22.08.2017 15:07	Admin Test	DELETE	204	/api/intel-sets/29/entities/	
22.08.2017 15:07	Test Test	POST	200	/api/auth	User 'test' logged in
22.08.2017 15:07	Mihai Danuta	POST	202	/api/enricher-tasks/batch-run	
22.08.2017 15:06	Test Test	DELETE	204	/api/work-in-progress/319	Deleted WorkInProgress (id:319)
22.08.2017 15:06	Test Test	POST	201	/api/taxonomies/	Created new TaxonomyNode (id:28)
22.08.2017 15:06	Test Test	PUT	200	/api/work-in-progress/319	Updated WorkInProgress (id:319)

View audit logs in Kibana

To view audit logs in Kibana, do the following:

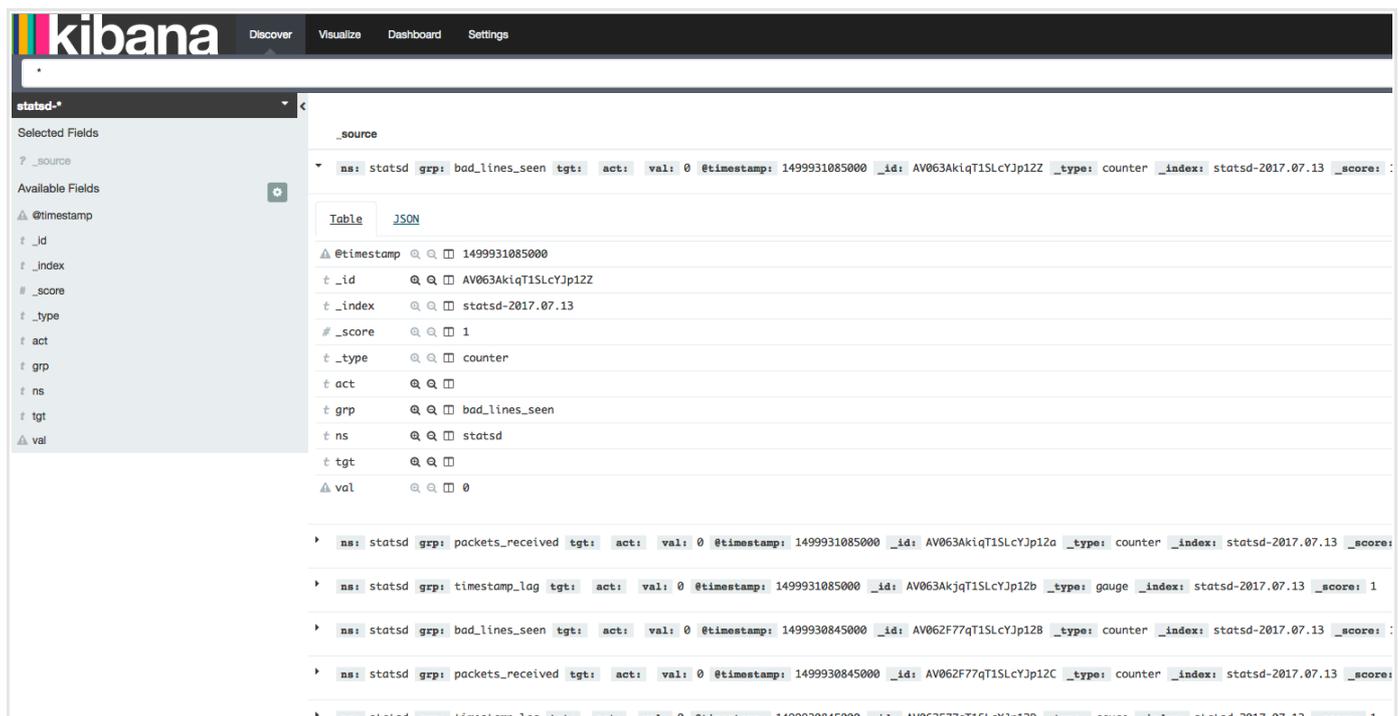
- To access Kibana, in the web browser address bar enter a URL with the following format:
*https://<platform_host>/api/kibana/app/kibana#/
Keep the trailing /
Example: *https://platform.host.com/api/kibana/app/kibana#/
In Kibana, select the **Discover** view.**
- Select one of the following indexes:
 - **audit**: audit trail index. It records events related to entities, datasets, enrichers, incoming and outgoing feeds, rules and tasks.
You can search for specific subsets by entering key/value pairs in the search input field.
Example: *method:DELETE; username:kmitnick; message:Deleted**

A user-friendlier way to view this information is by selecting **⚙ > System settings > Audit** on the left-hand navigation sidebar click to display the **Audit** tab.

- **logstash**: it records log information related to ingestion, tasks, and task scheduling.
You can search for specific subsets by entering key/value pairs in the search input field.
Example: **event:entities.stored; level:error; logger:eiq.platform.ingestion; tags:ingestion*
- **statsd**: it collects metrics about received packets and detected invalid or not well-formed lines in the ingested packets.
You can search for specific subsets by entering key/value pairs in the search input field.
Example: *grp:packets_received; grp:bad_lines_seen*
- If necessary, adjust the time interval by clicking the clock icon on the top-right corner, and by choosing an appropriate time range for the search.



- If audit logging is enabled, and if the audit log file is populated, audit log records are returned.



i The request payload for an audit log entry cannot exceed 4 KB. If the request body in the request payload for an audit log entry is larger than 4 KB, the following message is displayed:
Request body too large to log .

Search for audit logs in Kibana

To search for audit logs in Kibana, do the following:

- To access Kibana, in the web browser address bar enter a URL with the following format:
`https://<platform_host>/api/kibana/app/kibana#/`
 Keep the trailing /
 Example: `https://platform.host.com/api/kibana/app/kibana#/`

Adjust the update interval

By default, the log data update interval is set to 15 minutes. You can adjust it to match your requirements by clicking it, and then by choosing the appropriate value among the available options.

The screenshot shows the Kibana navigation bar with tabs for Discover, Visualize, Dashboard, and Settings. On the right, there is an 'Auto-refresh' button and a dropdown menu currently set to 'Last 15 minutes'. The dropdown menu is open, showing a list of time intervals categorized under 'Quick', 'Relative', and 'Absolute'.

Category	Option
Quick	Today
	This week
	This month
	This year
	The day so far
Relative	Week to date
	Month to date
	Year to date
	Yesterday
	Day before yesterday
Absolute	This day last week
	Previous week
	Previous month
	Previous year
	Last 15 minutes
	Last 30 minutes
	Last 1 hour
	Last 4 hours
	Last 12 hours
	Last 24 hours
	Last 7 days
	Last 30 days
	Last 60 days
	Last 90 days
	Last 6 months
Last 1 year	
Last 2 years	
Last 5 years	

Search by level

Log records are assigned a severity level. In this way, you can filter searches to retrieve only errors, warnings, or informative/notification log records.

To run a search by level in Kibana, do the following:

- In Kibana, select the **Discover** view.
- Make sure **logstash-*** is the active index:



Enter this in the Kibana search bar...	...to obtain this result
level:"error"	Returns all logging errors.
level:"warning"	Returns all logging warnings.
level:"info"	Returns all logging information and notifications.

Search by tag

Log records are tagged, so that you can identify the component a log refers to. To run a search by tag in Kibana to look for issues or information about a specific component, do the following:

- In Kibana, select the **Discover** view.
- Make sure **logstash-*** is the active index:



In the Kibana search bar enter this...	...to get this result
tags.raw:"ingestion"	Returns logs related to the intel ingestion module controlling how the platform ingests source data.
tags.raw:"opentaxii"	Returns logs related to the EclecticIQ OpenTAXII server implementing the TAXII services.
tags.raw:"platform-api"	Returns logs related to the EclecticIQ Platform API.
tags.raw:"task-workers"	Returns logs related to the services and processes carrying out complementary tasks.

Search with Boolean operators

You can use Boolean operators to refine your search and zero in on specific log types or components.

For example you can search for Redis and/or Neo4j issues by entering in the Kibana search bar the following query:

```
level:"error" AND tags.raw:"graph-ingestion"
```

You can search for Redis and/or Neo4j issues, or for issues concerning how data is ingested, by entering in the Kibana search bar the following query:

```
level:"error" AND tags.raw:"graph-ingestion" OR tags.raw:"intel-ingestion"
```

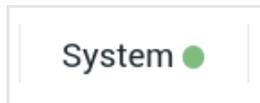
Check system health

System health provides a clear overview of the general health status of the platform and its core components.

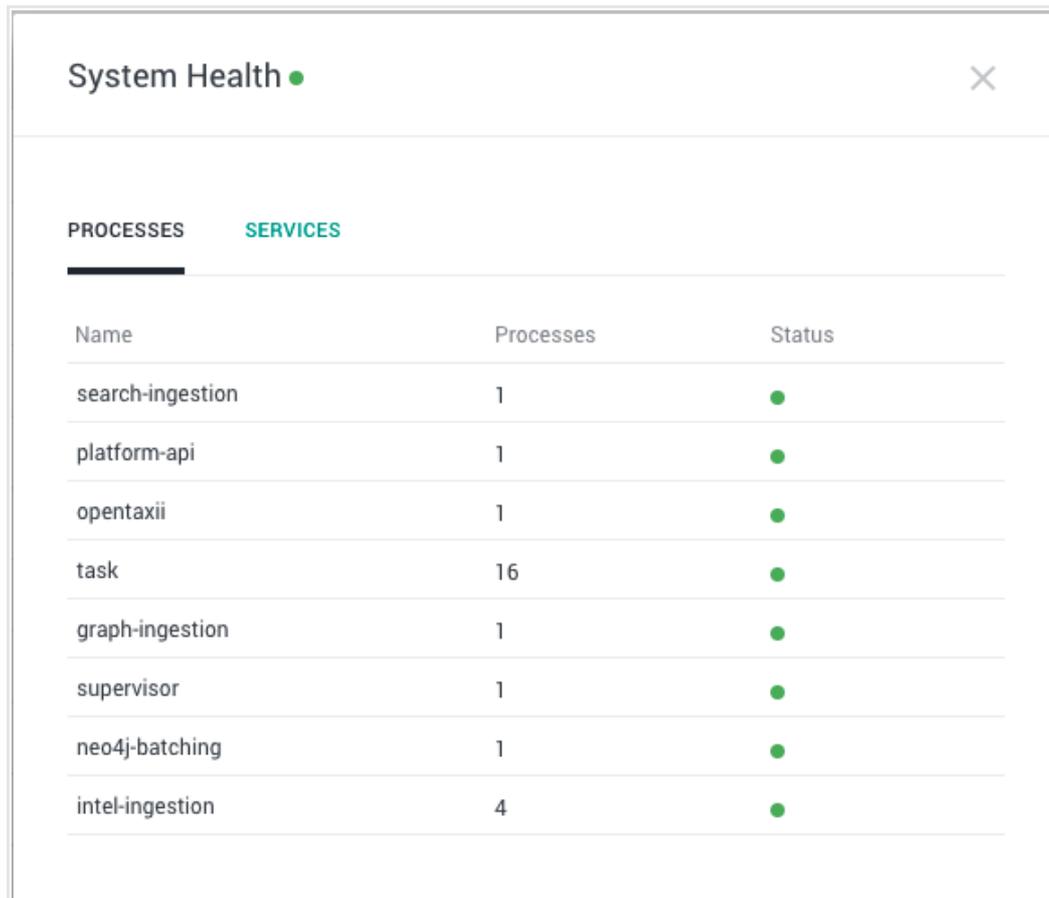
Check system health via the GUI

The easiest way to check platform status is by inspecting the system health tool:

- On the left-hand navigation sidebar click **⚙ > System running — View**.
- The **System health** pop-up view shows the current status of each component.



A green or a red dot next to the **System health** pop-up header indicates the global system health status. Click **Processes** or **Services** to review the status of the normal platform processes or of the platform core services.

A screenshot of the "System Health" pop-up window. The window has a title bar with "System Health" and a green dot, and a close button (X) in the top right. Below the title bar, there are two tabs: "PROCESSES" (selected with a thick black underline) and "SERVICES" (in blue text). The main content is a table with three columns: "Name", "Processes", and "Status". The table lists several components, all with a green dot in the "Status" column, indicating they are healthy.

Name	Processes	Status
search-ingestion	1	●
platform-api	1	●
opentaxii	1	●
task	16	●
graph-ingestion	1	●
supervisor	1	●
neo4j-batching	1	●
intel-ingestion	4	●

System Health ✕

PROCESSES
SERVICES

Name	Status
elasticsearch	● active
logstash	● active
postfix	● active
neo4j	● active
postgresql-9.5	● active
redis	● active
statsd	● active
kibana	● active

Process	Description
graph-ingestion	Data funnel to the Neo4j graph database. Handles data updates for Neo4j.
intel-ingestion	Intel ingestion through feeds and enrichers. Consumes incoming data and saves it PostgreSQL, Neo4j, and Elasticsearch. The platform executes one <code>intel-ingestion</code> per processor core. Running tasks are sequentially numbered starting from 0. For example, a platform instance running on a quad core machine normally executes 4 such processes, progressively numbered from <code>intel-ingestion:0</code> to <code>intel-ingestion:3</code> .
neo4j-batching	Neo4j graph database batch processing application. It lives on the same server hosting the Neo4j database. It prepares data for ingestion into the Neo4j database.
kibana	Generates dashboard graphs.
opentaxii	TAXII server responsible for STIX data transport.
platform-api	The web application implementing the platform API and the API endpoints. The endpoints expose services that can be consumed by making API calls and by passing arguments.
search-ingestion	Search indexer. Handles Elasticsearch data updates.
supervisor	Service monitoring and management
task	Celery-managed tasks such as enrichers, feed integrations, incoming feed data providers, and utilities.

Service	Description
elasticsearch	Elasticsearch search and indexing database.

Service	Description
logstash	Log and data aggregation, data pipeline and funneling.
neo4j	Neo4j graph database. CentOS/RHEL: neo4j; Ubuntu: neo4j-service.
nginx	Web server
postfix	Email server
postgresql-9.5	PostgreSQL (intel database). CentOS/RHEL: postgresql-9.5; Ubuntu: postgresql@9.5-main.
redis	Redis (message broker).

Check system health via the API

You can retrieve system health information also by making an API call to the `/status` API endpoint.



- First, make an authentication call to receive the bearer token you need to pass with all subsequent API calls.
- If the endpoint returns an HTTP 500 error, verify that the Supervisor configuration file includes the `[rpcinterface:supervisor]` section.

API endpoint

API endpoint	<code>/status</code>
Create method	GET
HTTP headers	"Content-Type: application/json", "Accept: application/json", "Authorization: Bearer <token>"
API request	GET + "Content-Type: application/json" + "Accept: application/json" + "Authorization: Bearer <token>" + <platform_host>/status
API response	{ "data" : { <status_reponse> } }

When you make a `GET` request and obtain a JSON object with entity data in the response, the entity object is wrapped in a data wrapper: { "data" : { ... } }

Status request

```
$ curl -X GET
-v
--insecure
-i
-H "Content-Type: application/json"
-H "Accept: application/json"
-H "Authorization: Bearer <token>"
https://platform.host/api/status

# copy-paste version:
$ curl -X GET -v --insecure -i -H "Content-Type: application/json" -H "Accept: application/json"
-H "Authorization: Bearer <token>" https://platform.host/api/status
```



Warning:

About cURL calls

- If you make HTTPs cURL calls to the API *and* you have a self-signed or an invalid certificate, include the `-k` or the `--insecure` parameter in the cURL call to skip the SSL connection CA certificate check.
- Always append a `/` trailing slash at the end of an API URL endpoint. The only exception is `/auth`, which does not take a trailing forward slash.
- In the cURL call, the `-d` data payload with the entity information always needs to be flat JSON, not hierarchical JSON.
If you want to pass a hierarchical JSON object, include the `--data-binary` parameter, followed by the path to the JSON file, for example `@/path/to/entity_file.json`.

Status response

```
{
  "data": {
    "celery_nodes_state": {

      "health": "GREEN",
      "nodes_down": [],

      "nodes_expected": [
        "discovery",
        "discovery-priority",
        "enrichers",
        "enrichers-priority",
        "entity-rules-priority",
        "extract-rules-priority",
        "incoming-transports",
        "incoming-transports-priority",
        "outgoing-feeds",
        "outgoing-feeds-priority",
        "outgoing-transports",
        "outgoing-transports-priority",
        "reindexing",
        "utilities",
        "utilities-priority"
      ]
    },

    "health": "GREEN",
```

```
"process_states": [  
  {  
    "first_down_at": null,  
    "first_up_at": "2017-06-28T06:44:44+00:00",  
    "health": "GREEN",  
    "n_processes": 1,  
    "n_processes_down": 0,  
    "name": "graph-ingestion"  
  },  
  {  
    "first_down_at": null,  
    "first_up_at": "2017-06-26T12:25:48+00:00",  
    "health": "GREEN",  
    "n_processes": 1,  
    "n_processes_down": 0,  
    "name": "opentaxii"  
  },  
  {  
    "first_down_at": null,  
    "first_up_at": "2017-06-28T06:41:08+00:00",  
    "health": "GREEN",  
    "n_processes": 1,  
    "n_processes_down": 0,  
    "name": "platform-api"  
  },  
  {  
    "first_down_at": null,  
    "first_up_at": "2017-06-28T06:44:31+00:00",  
    "health": "GREEN",  
    "n_processes": 4,  
    "n_processes_down": 0,  
    "name": "intel-ingestion"  
  },  
  {  
    "first_down_at": null,  
    "first_up_at": "2017-06-28T06:44:49+00:00",  
    "health": "GREEN",  
    "n_processes": 1,  
    "n_processes_down": 0,  
    "name": "search-ingestion"  
  },  
  {  
    "first_down_at": null,  
    "first_up_at": "2017-06-28T06:44:52+00:00",  
    "health": "GREEN",  
    "n_processes": 1,  
    "n_processes_down": 0,  
    "name": "neo4j-batching"  
  },  
  {  
    "first_down_at": null,  
    "first_up_at": null,  
    "health": "GREEN",  
    "n_processes": 1,  
    "n_processes_down": 0,  
    "name": "supervisor"  
  },  
],
```

```
{
  "first_down_at": null,
  "first_up_at": "2017-06-28T06:41:11+00:00",
  "health": "GREEN",
  "n_processes": 16,
  "n_processes_down": 0,
  "name": "task"
}
],
"service_states": [
  {
    "health": "GREEN",
    "name": "elasticsearch",
    "state": "active"
  },
  {
    "health": "GREEN",
    "name": "logstash",
    "state": "active"
  },
  {
    "health": "GREEN",
    "name": "postfix",
    "state": "active"
  },
  {
    "health": "GREEN",
    "name": "neo4j",
    "state": "active"
  },
  {
    "health": "GREEN",
    "name": "postgresql-9.5",
    "state": "active"
  },
  {
    "health": "GREEN",
    "name": "redis",
    "state": "active"
  }
]
}
```

Monitor system health

System administrators can use tools like Celery and Supervisor to monitor platform tasks to check day-to-day operations, and to investigate in case of issues.

Monitor the platform to ensure normal operation, to research and identify the root cause of an issue, and to inspect the status of key platform processes such as incoming and outgoing feeds, enrichers, ingestion queues, and tasks.

In the current context, monitoring covers on/off status only: the tasks and the commands described here allow you to verify whether a task, a process, or a component is running or not. Metrics and other types of measurements are outside the scope of the topic.

System administrators and DevOps engineers can run quick checks to inspect platform operation, to identify issues and review errors, so that they can address them in a timely manner.

Tools

Celery	The task runner. It manages task execution and scheduling.
Redis	The message broker. It handles background task processing by managing message queues based on the pub-sub pattern (https://en.wikipedia.org/wiki/publish%e2%80%93subscribe_pattern).
Supervisor	The process controller to start and stop processes.
systemd	The initialization system to bootstrap processes and start services.

Core components

Component	Address	Port
platform	localhost	8008
nginx	<server_name> (http://nginx.org/en/docs/http/nginx_http_core_module.html#server_name)	80; 443
postfix	<myhostname> (http://www.postfix.org/postconf.5.html#myhostname)	25; 587
postgresql	localhost	5432
redis	localhost	6379
neo4j	localhost	7474; 7473
elasticsearch	localhost	9200
kibana	localhost	5601

Component	Address	Port
logstash	localhost	6755
statsd	127.0.0.1:8125 (https://github.com/etsy/statsd)	8125

The default PostgreSQL database name is `platform`. You can change it as necessary.

Monitoring

Platform monitoring covers two main areas:

Components	
platform-api	The web application implementing the platform API and the API endpoints. The endpoints expose services that can be consumed by making API calls and by passing arguments.
nginx	Web server
postfix	Email server
opentaxii	TAXII server responsible for STIX data transport.
postgresql-9.5	PostgreSQL (intel database). CentOS/RHEL: <code>postgresql-9.5</code> ; Ubuntu: <code>postgresql@9.5-main</code> .
redis	Redis (message broker).
elasticsearch	Elasticsearch search and indexing database.
kibana	Generates dashboard graphs.
logstash	Log and data aggregation, data pipeline and funneling.
neo4j	Neo4j graph database. CentOS/RHEL: <code>neo4j</code> ; Ubuntu: <code>neo4j-service</code> .
statsd	Gather stats such as counters, timers, discovered entities and so on, and it sends aggregates to Kibana through Elasticsearch.

Processes	
graph-ingestion	Data funnel to the Neo4j graph database. Handles data updates for Neo4j.
intel-ingestion	Intel ingestion through feeds and enrichers. Consumes incoming data and saves it PostgreSQL, Neo4j, and Elasticsearch. The platform executes one <code>intel-ingestion</code> per processor core. Running tasks are sequentially numbered starting from 0. For example, a platform instance running on a quad core machine normally executes 4 such processes, progressively numbered from <code>intel-ingestion:0</code> to <code>intel-ingestion:3</code> .
neo4j-batching	Neo4j graph database batch processing application. It lives on the same server hosting the Neo4j database. It prepares data for ingestion into the Neo4j database.
search-ingestion	Search indexer. Handles Elasticsearch data updates.

Processes	
task	Celery-managed tasks such as enrichers, feed integrations, incoming feed data providers, and utilities.
Feeds	Incoming and outgoing feeds
Enrichers	Enricher tasks
Celery tasks	Other/Misc. Celery tasks

To successfully execute several commands in the command line or in the terminal, you may need root-level access rights.

To obtain admin rights, run the following command(s):

```
$ sudo su -
```

Alternatively:

- Grant admin rights to a specific user, who can then log in with their password to perform admin tasks:

```
$ su - {user_name}
```

Or:

- Prefix `sudo` to the command you want to run:

```
$ sudo {command}
```

Monitor components with Supervisor

Tool: *Supervisor* is a useful tool to inspect platform components to verify if they are operating normally.

These are the configuration files of the applications Supervisor manages:

File	Location
platform-api.ini	/opt/eclecticiq/etc/supervisord.d/
graph-ingestion.ini	/opt/eclecticiq/etc/supervisord.d/
intel-ingestion.ini	/opt/eclecticiq/etc/supervisord.d/
search-ingestion.ini	/opt/eclecticiq/etc/supervisord.d/
neo4j-batching.ini	/opt/eclecticiq/etc/supervisord.d/
opentaxii.ini	/opt/eclecticiq/etc/supervisord.d/

File	Location
task-workers.ini	/opt/eclecticiq/etc/supervisord.d/

Use it to check the following components:

Component	Description	If it is not running...
graph-ingestion	Data funnel to the Neo4j graph database. Handles data updates for Neo4j.	The graph database may go out of sync and miss data. The <code>queue:graph:inbound</code> queue increases in size.
intel-ingestion	Intel ingestion through feeds and enrichers. Consumes incoming data and saves it PostgreSQL, Neo4j, and Elasticsearch. The platform executes one <code>intel-ingestion</code> per processor core. Running tasks are sequentially numbered starting from 0. For example, a platform instance running on a quad core machine normally executes 4 such processes, progressively numbered from <code>intel-ingestion:0</code> to <code>intel-ingestion:3</code> .	Provider tasks keep running, but no data is displayed in the system. The <code>queue:ingestion:inbound</code> queue increases in size.
search-ingestion	Search indexer. Handles Elasticsearch data updates.	Search indexing stops working correctly. Elasticsearch may go out of sync and miss data. The <code>queue:search:inbound</code> queue increases in size.
neo4j-batching	Neo4j graph database batch processing application. It lives on the same server hosting the Neo4j database. It prepares data for ingestion into the Neo4j database.	Graph ingestion stops working, queries may return results that are not up to date.
opentaxii	TAXII server responsible for STIX data transport.	The TAXII transport type becomes unavailable.
platform-api	The web application implementing the platform API and the API endpoints. The endpoints expose services that can be consumed by making API calls and by passing arguments.	Platform services become unavailable and the API endpoints return an HTTP 502 error.
task	Celery-managed tasks such as enrichers, feed integrations, incoming feed data providers, and utilities.	Tasks may or may not start or stop, resulting in unexpected task execution behavior.
task:beat	Task scheduler	Task execution order may be affected, and tasks may or may not start or stop, resulting in unexpected task execution behavior.
task:enrichers	Enricher tasks	Enricher data may become only partially available or completely unavailable.

Component	Description	If it is not running...
<code>task:integrations</code>	Outgoing feed integrations	Outgoing feed transport types are affected, and they may stop working correctly or become unavailable.
<code>task:providers</code>	Incoming feed data provider tasks	Incoming feed transport types are affected, and they may stop working correctly or become unavailable.
<code>task:utilities</code>	The little platform elves that keep things tidy while working in the background.	The platform may behave unexpectedly. For example data updates and discovery may hang or stop working.

`supervisorctl` is Supervisor's command line interface utility. The commands you can pass are called *actions*, and they can optionally take *arguments*:

```
$ supervisorctl <action> <argument>
```

To check if Supervisor is installed, run the following command(s):

```
$ yum info supervisor
```

To check if the `supervisord` daemon is running, run the following command(s):

```
$ systemctl status supervisord
```

Launch Supervisor to load the configuration to start the tasks and processes other platform components depend on:

```
$ systemctl start supervisord
```

If you need to stop Supervisor, run the following command(s):

```
$ systemctl stop supervisord
```

To check the statuses of the tasks managed by Supervisor, run the following command(s):

```
$ supervisorctl status
```

The response should return `RUNNING` for all relevant tasks to confirm that all Supervisor tasks are being executed normally.

The following example serves as a guideline:

graph-ingestion	RUNNING	pid 19527, uptime 0:00:03
intel-ingestion:0	RUNNING	pid 19071, uptime 0:00:51
intel-ingestion:1	RUNNING	pid 19070, uptime 0:00:51
intel-ingestion:2	RUNNING	pid 19073, uptime 0:00:51
intel-ingestion:3	RUNNING	pid 19072, uptime 0:00:51
neo4j-batching	RUNNING	pid 19268, uptime 0:00:43
opentaxii	RUNNING	pid 19330, uptime 0:00:36
platform-api	RUNNING	pid 19077, uptime 0:00:51
search-ingestion	RUNNING	pid 19075, uptime 0:00:51
task:beat	RUNNING	pid 19061, uptime 0:00:51
task:discovery	RUNNING	pid 19068, uptime 0:00:51
task:discovery-priority	RUNNING	pid 19065, uptime 0:00:51
task:enrichers	RUNNING	pid 19056, uptime 0:00:51
task:enrichers-priority	RUNNING	pid 19062, uptime 0:00:51
task:entity-rules-priority	RUNNING	pid 19063, uptime 0:00:51
task:extract-rules-priority	RUNNING	pid 19055, uptime 0:00:51
task:incoming-transport	RUNNING	pid 19053, uptime 0:00:51
task:incoming-transport-priority	RUNNING	pid 19054, uptime 0:00:51
task:outgoing-feeds	RUNNING	pid 19066, uptime 0:00:51
task:outgoing-feeds-priority	RUNNING	pid 19057, uptime 0:00:51
task:outgoing-transport	RUNNING	pid 19060, uptime 0:00:51
task:outgoing-transport-priority	RUNNING	pid 19058, uptime 0:00:51
task:reindexing	RUNNING	pid 19064, uptime 0:00:51
task:utilities	RUNNING	pid 19059, uptime 0:00:51
task:utilities-priority	RUNNING	pid 19067, uptime 0:00:51

To check the statuses of specific tasks managed by Supervisor, run the following command(s):

```
# Retrieve the status of a specific process
$ supervisorctl status <process_name>

# Retrieve the status of multiple specific processes
$ supervisorctl status <process_name> <process_name> <process_name> ...

# Retrieve the status of all processes
# whose name contains the specified search string
$ supervisorctl status | grep "<search_string>"
```

To reload the Supervisor configuration and to restart all Supervisor-managed processes run the following command(s):

```
$ supervisorctl reload
```



Warning: When you edit or update Supervisor configurations, run `systemctl restart supervisor` and `supervisorctl reload`, so that Supervisor can pick up and reload any updated configurations to the platform with the latest changes.

Supervisor actions

Examples of `supervisorctl` actions:

Run this...	...to do this
<code>\$ supervisorctl start all</code>	Start all the processes defined in the Supervisor configuration

Run this...	...to do this
<code>\$ supervisorctl start <process_name></code>	Start the specified process defined in the Supervisor configuration
<code>\$ supervisorctl start <process_name> <process_name></code>	Start the specified processes defined in the Supervisor configuration
<code>\$ supervisorctl stop all</code>	Stop all the processes defined in the Supervisor configuration file
<code>\$ supervisorctl stop <process_name></code>	Stop the specified process
<code>\$ supervisorctl stop <process_name> <process_name></code>	Stop the specified processes
<code>\$ supervisorctl status</code>	Retrieve the statuses of the processes managed by Supervisor. For further information, see the official ocumentation on the return state values (http://supervisord.org/subprocess.html#process-states)
<code>\$ supervisorctl status <process_name></code>	Retrieve the status of the specified process
<code>\$ supervisorctl status <process_name> <process_name></code>	Retrieve the status of the specified processes
<code>\$ supervisorctl reload</code>	Reload the Supervisor configuration and restart all tasks and processes. If you modify or update the Supervisor configuration file, you need run this command to reload the latest Supervisor configuration file
<code>\$ supervisorctl reload all</code>	Reload all the processes managed by Supervisor. This command works just like <code>\$ supervisorctl reload</code>
<code>\$ supervisorctl update</code>	Restart the applications whose configuration has changed. This command affects existing configurations that were modified. If new application configurations become available, they start automatically only after restarting Supervisor or after rebooting the system. Typically, you run <code>reread</code> and then <code>update</code> .
<code>\$ supervisorctl tail <log_file_name></code>	Retrieve the most recent lines of the specified log file. To follow the log file as it updates with new information and new lines, run <code>\$ supervisorctl tail -f <log_file_name></code>
<code>\$ supervisorctl status grep "<search_string>"</code>	Retrieve the status of all processes whose name contains the specified search string.
<code>\$ supervisorctl reread</code>	Update the changed configurations and reload them. It does not automatically (re)start any applications. Typically, you run <code>reread</code> and then <code>update</code> .

Monitor components with systemd

Tool: *systemd* helps you inspect platform components to verify if their services are running normally. Use it to check the following components:

Component	Description	If it is not running...
postgresql-9.5	PostgreSQL (intel database). CentOS/RHEL: postgresql-9.5; Ubuntu: postgresql@9.5-main.	It is not possible to access platform data.
redis	Redis (message broker).	Tasks and processes may hang and/or behave unexpectedly.
neo4j	Neo4j graph database. CentOS/RHEL: neo4j; Ubuntu: neo4j-service.	Graph data queries stop working, it is not possible to poll the graph database.
elasticsearch	Elasticsearch search and indexing database.	No data searching and indexing capabilities are available.
kibana	Generates dashboard graphs.	The dashboard does not load correctly, and the /kibana/ API endpoint returns a HTTP 502 error.
logstash	Log and data aggregation, data pipeline and funneling.	No data aggregation, deduplication, and normalization.
statsd	Gather stats such as counters, timers, discovered entities and so on, and it sends aggregates to Kibana through Elasticsearch.	No metrics about discovered entities, feed updates, and so on.
postfix	Email server	No automatic email notifications.

`systemctl` is `systemd`'s command line interface utility. The commands can optionally take options:

```
$ systemctl <options> <command> <component_name>
```

For a complete list of supported commands and options, see the **systemd documentation**

(<https://www.freedesktop.org/software/systemd/man/systemctl.html>).

To obtain a list of all running services, run the following command(s):

```
$ systemctl
```

The response is displayed in the following format:

```
UNIT          LOAD          ACTIVE        SUB           JOB DESCRIPTION
<service_name> <loaded_or_not> <active_or_not> <running_or_not> <description_of_the_job>
```

To verify if Nginx is running, run the following command(s):

```
$ systemctl status -l nginx
```

To verify if PostgreSQL is running, run the following command(s):

```
$ systemctl status -l postgresql-9.5
```

To verify if Redis is running, run the following command(s):

```
$ systemctl status -l redis
```

To verify if Logstash is running, run the following command(s):

```
$ systemctl status -l logstash
```

To verify if Elasticsearch is running, run the following command(s):

```
$ systemctl status -l elasticsearch
```

To retrieve status information about all these services at once, run the following command(s):

```
$ systemctl status -l nginx postgresql-9.5 redis logstash elasticsearch
```

To retrieve status information about all the systemd-managed services whose name contains a specific search string, run the following command(s):

```
$ systemctl | grep "<search_string>"
```

Monitor processes

Monitor ingestion queues with Redis

Tool: *Redis* acts as a message broker for Celery-managed tasks.

`redis-cli` is Redis's command line interface utility:

```
# Launch redis-cli
$ redis-cli

$ > <command> <item_name>
```

For a complete list of supported commands and options, see the ***redis-cli* command reference**

(<https://redis.io/topics/rediscli>).

Within the platform Redis manages task queues. Possibly the only command you need is the one that allows you to check queue length: `llen`.

To inspect the platform data ingestion queue length, run the following command(s):

```
# Launch redis-cli
$ redis-cli

$ > llen "queue:ingestion:inbound"
```

To inspect the graph database queue length, run the following command(s):

```
$ > llenn "queue:graph:inbound"
```

To inspect the Elasticsearch data update queue length, run the following command(s):

```
$ > llenn "queue:search:inbound"
```

Monitor running tasks with Celery

Tool: *Celery* is the task runner that manages task execution and scheduling.

To use Celery to request task information, you need to pass the following environment variable(s) with your request:

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
```

Append Celery commands after the environment variable(s). Celery commands have the following format:

```
$ celery -A <module_name> <command>
```

Ping Celery to see which tasks are up and listening. This is the easiest way to check task running status. All active tasks reply with pong.

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
$ /opt/eclecticiq/platform/api/bin/celery -A eiq.platform.taskrunner.app inspect ping
```

To inspect active tasks, run the following command(s):

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
$ /opt/eclecticiq/platform/api/bin/celery -A eiq.platform.taskrunner.app inspect active
```

To inspect active tasks queues, run the following command(s):

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
$ /opt/eclecticiq/platform/api/bin/celery -A eiq.platform.taskrunner.app inspect active_queues
```

To inspect scheduled tasks, run the following command(s):

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
$ /opt/eclecticiq/platform/api/bin/celery -A eiq.platform.taskrunner.app inspect scheduled
```

To inspect overall task status, run the following command(s):

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
$ /opt/eclecticiq/platform/api/bin/celery -A eiq.platform.taskrunner.app status
```

To request task statistics (exhaustive, but it can be verbose), run the following command(s):

```
$ export EIQ_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
$ /opt/eclecticiq/platform/api/bin/celery -A eiq.platform.taskrunner.app inspect stats
```

(For further details, see the documentation on **Celery ping**

(<http://docs.celeryproject.org/en/latest/userguide/workers.html#ping>), **Celery workers**

(<http://docs.celeryproject.org/en/latest/userguide/workers.html>), **Celery worker statistics**

(<http://docs.celeryproject.org/en/latest/userguide/workers.html#worker-statistics>), and **Celery monitoring** (<http://docs.celeryproject.org/en/latest/userguide/monitoring.html>)