Getting started with EclecticIQ Platform

User guide for cyber threat analysts

Last generated: July 21, 2017

Table of contents

Table of contents 2
Getting started with EclecticIQ Platform 5
Scope 5
Goal 5
Audience 5
Feedback 5
Launch the platform 6
RPM install 6
Access the platform 6
Virtual appliance 6
Start the VM 6
Get the VM IP address 7
Go to the platform 7
Configure the system 8
Configure the system 8
Server 8
Proxy 9
Proxy update 10
Email 11
STIX 12
TAXII 13
TAXII services 14
Add a TAXII service 15
View TAXII services 17
License 18
Monitor the system 19
Exposure 20
Audit 21
Jobs 22
View jobs 22
Terminate jobs 23
Configure users and roles 24
Configure users 24
Save options 25
View users 25
Configure roles 26
Configure permissions 28
Configure groups 28
Save options 29
View groups 29
Incoming feeds 31
Content types 31
Transport types 32
Configure incoming feeds 34
Set up incoming feeds 34
Save options 35
Run feeds 36
Manually run a feed task 36
Check task results 36
Configure outgoing feeds 37
Set up outgoing feeds 37
Set a schedule 37
Set a TLP override 38
Set reliability and relevancy 38
Set observable filters 38
Save options 39
Enrichment 40
Enriching entities via observables 40

Enrich entities 41
Automatically enrich entities 42
Manually enrich entities 42
Enricher rules 45
View enricher rules 45
Add enricher rules 46
Save options 47
Edit enricher rules 47
Delete enricher rules 48
Enricher tasks 48
View enricher tasks 48
Edit enricher tasks 49
Taxonomy 51
The Taxonomy feature 51
Predefined taxonomies 51
Admiralty code 51
Kill chain 53
Create a taxonomy entry 55
Save options 56
Edit a taxonomy entry 56
Delete a taxonomy entry 57
Rules 58
Rule types 58
Entity rules 58
Add an entity rule 58
Save options 62
Observable rules 62
Add an observable rule 62
Save observable rules 65
Derivation and levels 65
Original + level 1 66
Original + level 2 66
Derived + level 1 66
Derived + level 2 67
Edit rules 68
Delete rules 68
Filter rules 68
Example 69
Observables 73
View matching observables 73
Delete matching observables 75
Get observable types via API 75
Observable types 79
Dashboard 81
Customize the dashboard 81
Search the platform 83
Search 83
Search cheatsheet 83
Search query fields 84
Search timeout 86
Upload data files 87
Upload data 87
Discovery 89
View discovery rules 89
Create discovery rules 90
Save options 90
Edit discovery rules 91
Delete discovery rules 92
Run rules manually 92
Workspaces 93

Workspace types 93
Access workspaces 94
Create a workspace 94
Workspace types 96
Workspace tiles 96
Manage workspaces 96
Workspace Overview tab 98
Add and remove collaborators 98
Workspace Tasks tab 99
View tasks 99
View tasks created by or assigned to the current user 99
View tasks by status 100
View task details 101
View task status 101
Edit tasks 102
Workspace Comments tab 104
Workspace Saved graphs tab 105
Workspace Entities tab 106
Actions 106
View entity details 107
Create a new entity 108
Filter entities 110
Entity reliability 112
Workspace Files tab 114
Upload a file 114
Workspace Edit details tab 116
Archive and delete a workspace 116
Use the graph 117
Add entities to the graph 117
View the graph 118
Examine entities on the graph 120
Change visualization layout 124
Move around on the graph 125
Datasets 127
Create a dataset 127
Save options 128
Edit or delete a dataset 128
Tasks 130
Create a task 130
Save options 131
View tasks 131
View tasks created by or assigned to the current user 131
View tasks by status 132
View task details 132
View task status 133
Editor 135
Go to the entity editor 135
View entity details 135
Overview 136
Enrichments 139
Neighborhood 139
JSON 145
Versions 145
History 145
Entity types 145
Create an entity 146
Filter entities 153
Manage entities 155

Getting started with EclecticIQ Platform

This Getting Started guide helps you set up, configure and start working with EclecticIQ Platform.

Scope

This document guides you through the steps you need to carry out to complete the following tasks:

- Configure the platform
- Use it to carry out typical tasks, like ingesting data, analyze threat data, and share the results.

Goal

After completing these tasks, you'll be able to use the platform to perform standard threat analysis tasks, like:

- Acquire cyber threat data through feeds
- Visualize the data
- Analyze the data to extract actionable intelligence.

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

Launch the platform

Launch Google Chrome (recommended web browser), enter the platform address, and sign in to start working.

RPM install

After correctly installing and configuring the EclecticIQ Platform, you can sign in to start working.

Access the platform

To access the platform, do the following:

- Launch a web browser (recommended: Google Chrome).
- Go to the configured platform address, for example: *https://platform.host*

> **Warning:** The browser may display an untrusted connection warning: add it as an exception, and then proceed to the platform.

Virtual appliance

After setting up the EclecticIQ Platform as a virtual appliance, you can sign in to start working.

Start the VM

> To access the VM, you may need to enter valid login credentials. If you do not have these details, contact us.

- Launch VirtualBox or VMWare Player, and then start/play the VM.
- If you are prompted for login credentials at startup, enter the provided user name and password.
- The VM should be up and running.

Get the VM IP address

By default, our VM images run **CentOS Linux 7 (1511)** (https://lists.centos.org/pipermail/centos-announce/2015-december/021518.html).

- Run the following command(s):

```
$ ifconfig
```

- Press **ENTER**.
- Look for the following entry to identify the VM IP address:

```
inet <IP_address>
```

The `inet` IP address is the one you need to use to access the platform.
Example: `inet 10.0.2.148`.
Go to *https://10.0.2.148* to sign in to the platform web-based UI.

Go to the platform

- In your host machine, launch a web browser (recommended: Google Chrome).
- In the address bar, enter the VM IP address.
  Example: *https://10.0.2.148*
- The platform login screen is displayed.
- Sign in with the appropriate credentials.

> **Warning:** The browser may display an untrusted connection warning: add it as an exception, and then proceed to the platform.

Configure the system

Configure the system, define the STIX namespace, and set up TAXII services.

Before you start using the platform to ingest and analyze cyber threat data, you need to configure it.

There are two main configuration areas:

- System: set up the platform so that it can correctly fetch data from external sources.
- Users: set up the platform with well-defined user roles, permissions, and groups to create a secure and structured user interaction environment.

> ✔ On the forms, input fields marked with an asterisk are required.

Configure the system

You can configure the system by defining options in the following sections:

- Platform server address
- Proxy, if applicable
- The email address used to send automatic platform notifications from
- STIX namespace and its corresponding alias
- TAXII services to ingest and publish intelligence

Server

- On the left-hand navigation sidebar, click **System**.
- Click **Server > General**.
- Under **Server settings**, click **Edit settings**.
- Under **Hostname**, enter the platform host name, for example *platform.host*.
  The platform host name should match an existing server name 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 specific instructions.
  If you enter an incorrect platform host name, outgoing feeds and TAXII links won't work.
- Click **Save** to store your changes, or **Cancel** to discard them.

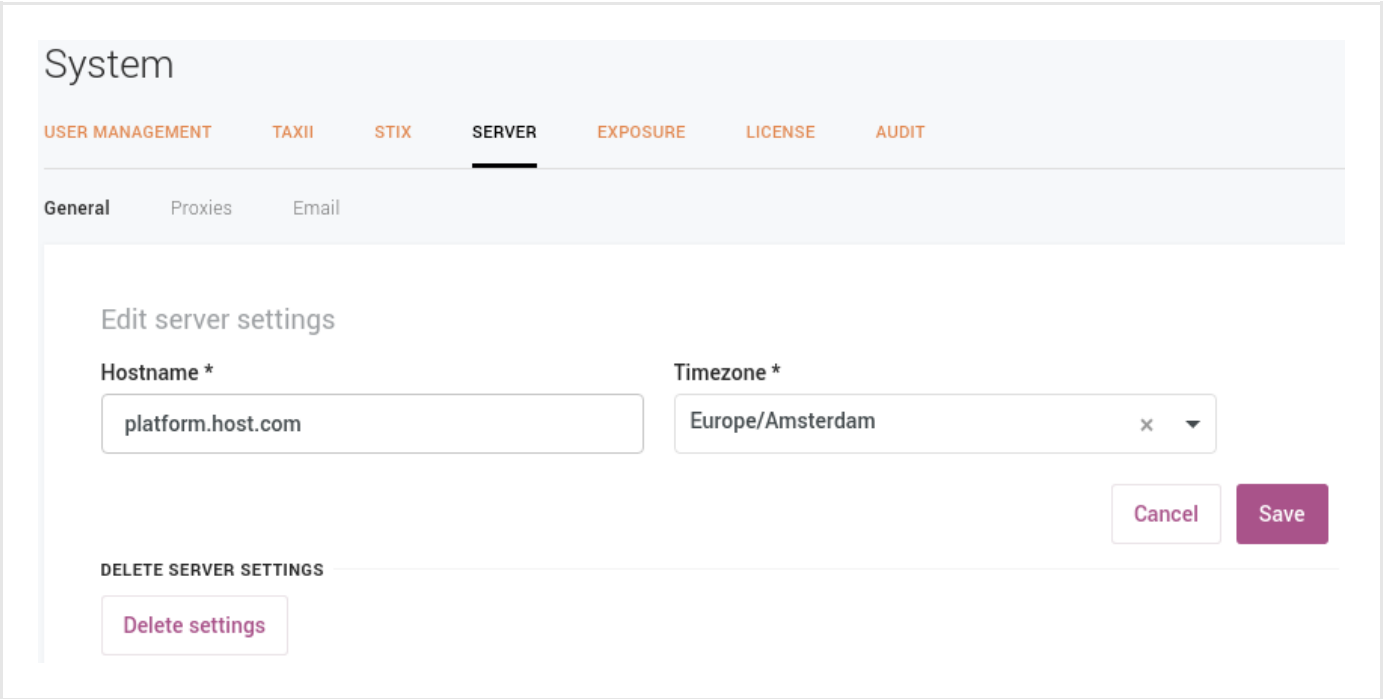

Proxy

If your Internet connection setup includes a proxy server, specify the configuration in this section.

- On the left-hand navigation sidebar, click **System**.
- Click **Server > Proxies**.
- Under **Add 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**: the proxy server domain name, for example *host.com*.
  - **Port**: the proxy server access port, for example *9999*.
  - **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. You may want to specify here local network addresses or LAN subdomains, for example. When 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.

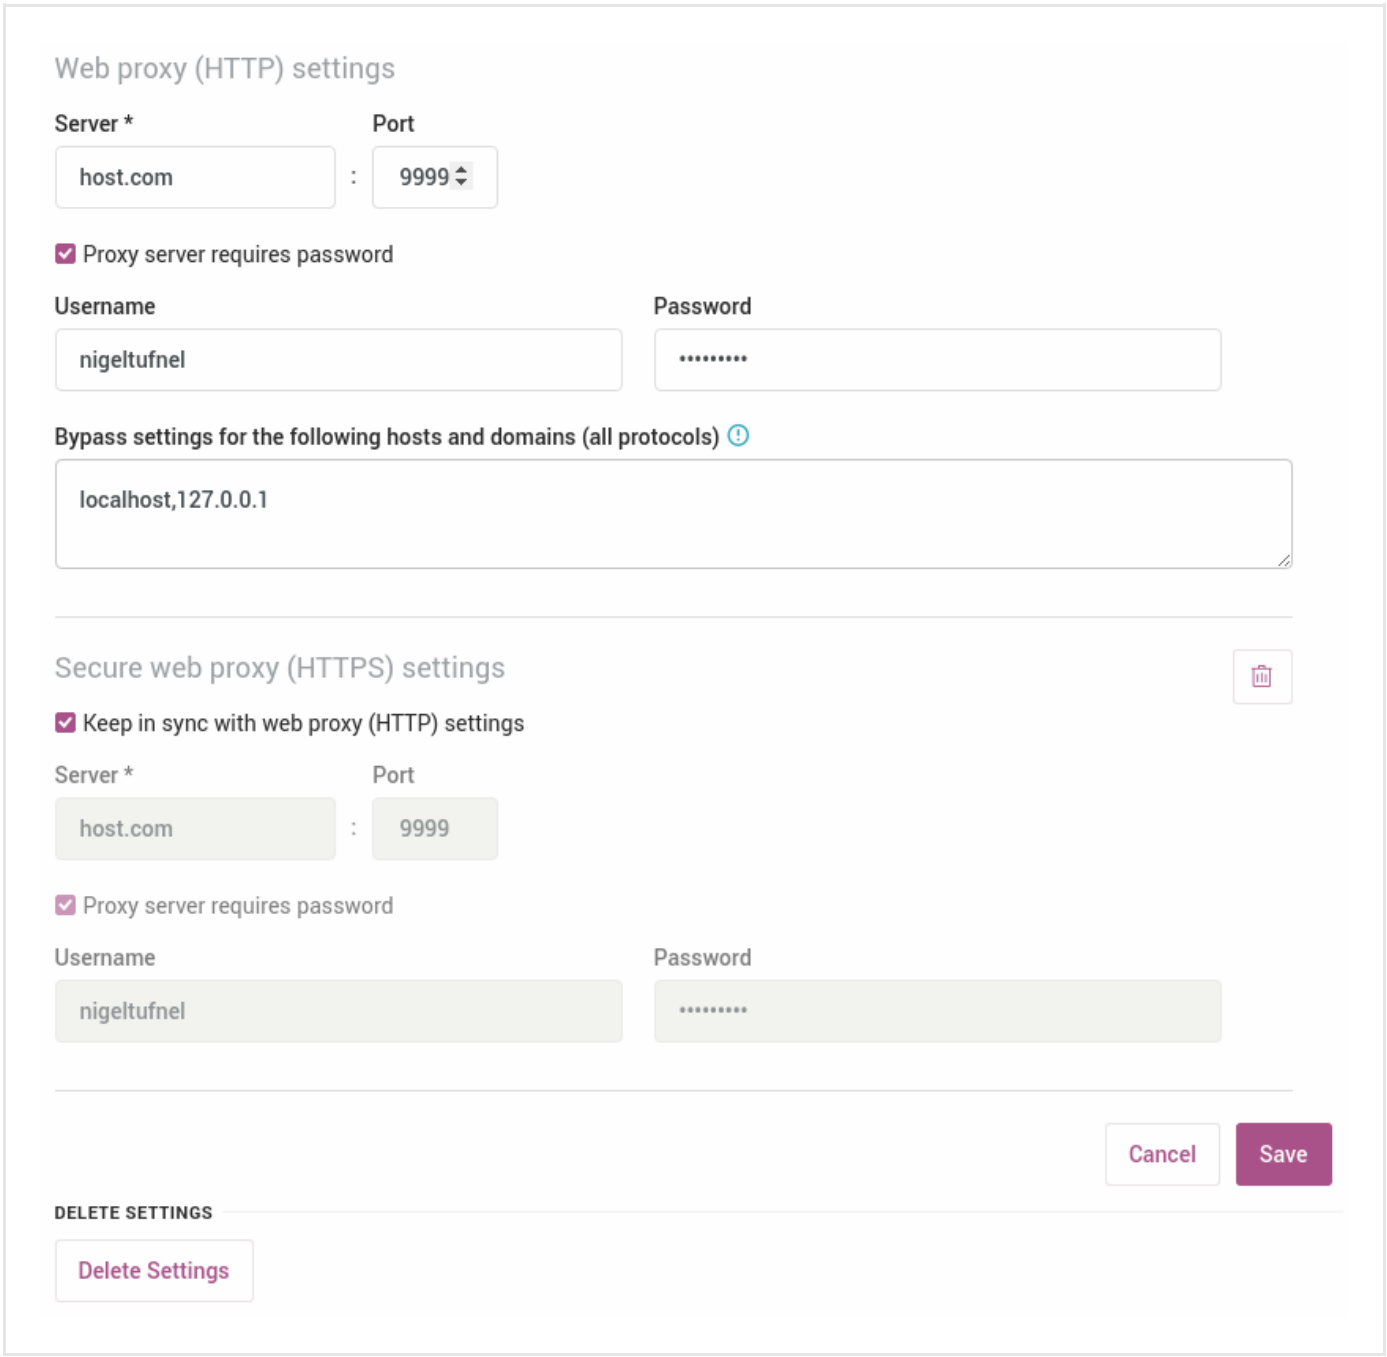

Proxy update

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):

```
$ systemctl restart supervisor
```

or:

```
$ systemctl stop supervisor

$ systemctl start supervisor
```

Email

This section configures the email address automatic platform notifications originate from.

- On the left-hand navigation sidebar, click **System**.
- Click **Email**, and then **Edit settings**.
- Under **Edit email settings**, define the following configuration settings:
  - **From email**: the email address used to send automatic email notifications from.
  - **SMTP host**: the outgoing email service address, for example *smtp.emailserver.com*.
  - **SMTP port**: the outgoing email service port.
    The standard **submission port number** `(https://tools.ietf.org/html/rfc6409)` for email services is *587*.
  - **SMTP username**: usually, this value corresponds to the email address.
  - **SMTP password**: enter valid password credentials to authenticate and to receive authorization to access the resource(s).
    Example: *thesegoto11*. .
  - **SMTP connection type**: the cryptographic data transport protocol.
    Allowed values: *default* (based on system configuration), *SSL*, *TLS*.
  - Click **Save** to store your changes, or **Cancel** to discard them.

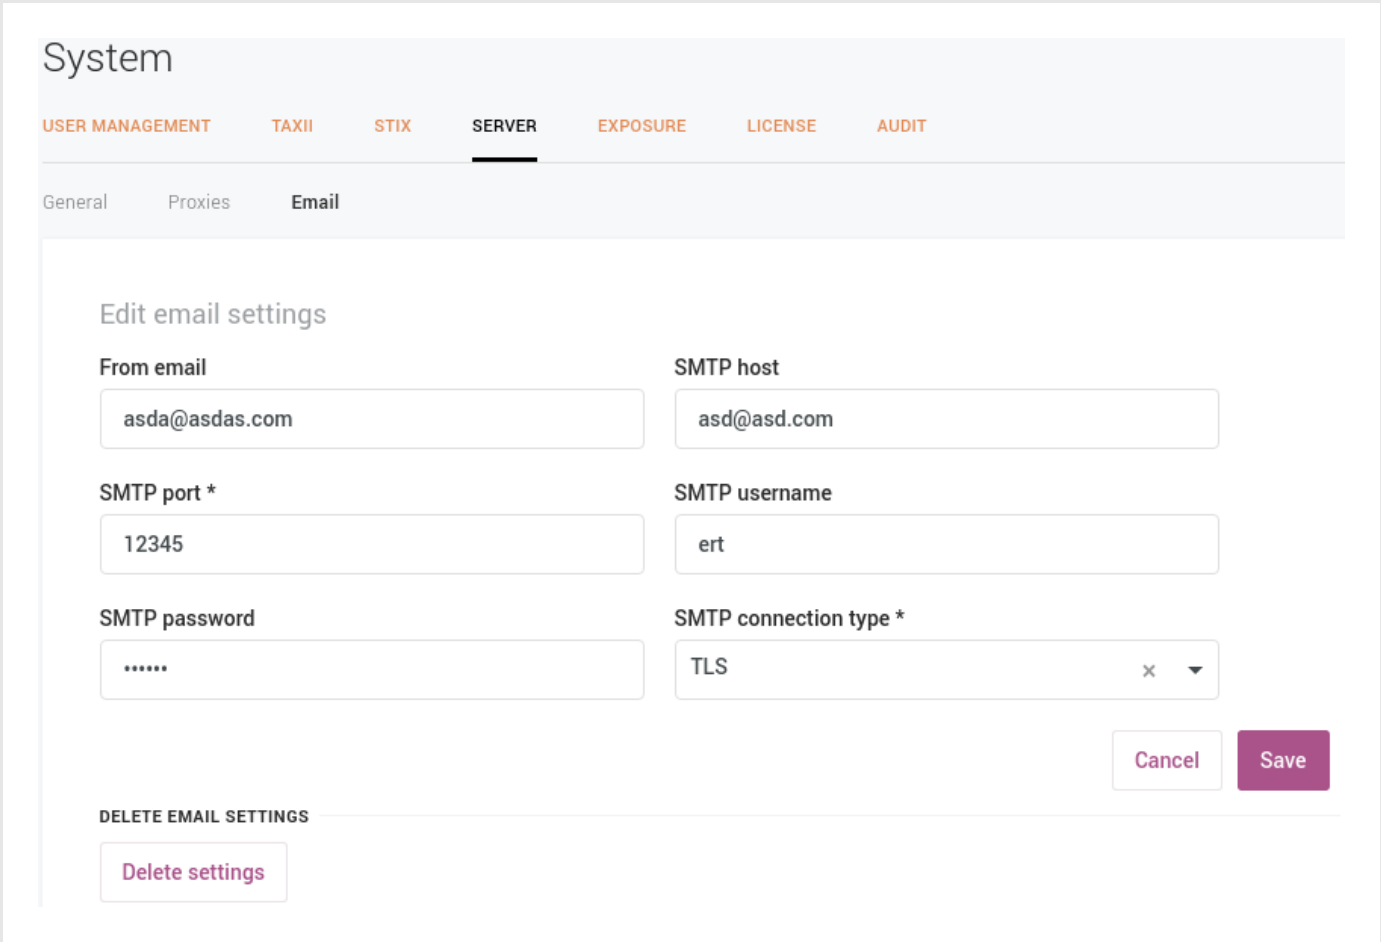

STIX

- On the left-hand navigation sidebar, click **System**.
- Click **STIX**, and then **Edit settings**.
- Under **Edit STIX settings**, define the following configuration settings:
  - **Alias**: the alias of the namespace you declare for your organization, for example *mymightyorganization*.
    Allowed characters for the alias:
    alphanumeric [*A-Z, a-z, 0-9*], underscore [_], dash [-], baseline 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, for example *http://stix.mymightyorganization.com/stix-1*.
  - **Producer**: optionally, you can enter here a name to identify your organization as the producer, i.e. the creator and/or the publisher of the STIX data.
  - Click **Save** to store your changes, or **Cancel** to discard them.

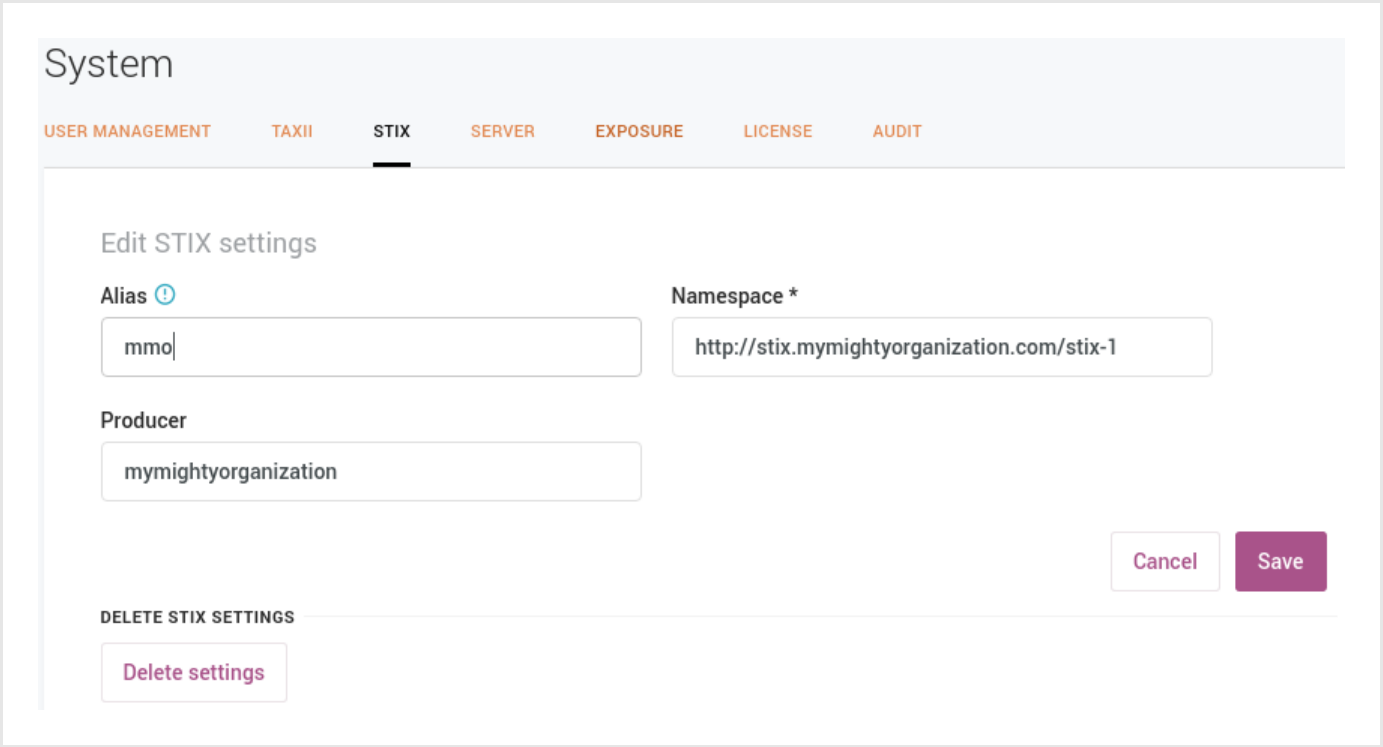

TAXII

The TAXII server is the designated transport handler for STIX data traffic. To set up and configure a TAXII server, do the following:

- On the left-hand navigation sidebar, click **System**.
- Click **TAXII**, **Settings**, and then **Edit settings**.
- Under **Edit TAXII server settings**, specify the domain of the TAXII server handling data traffic, for example *taxii.myserver.com*.
- Click **Save** to store your changes, or **Cancel** to discard them.

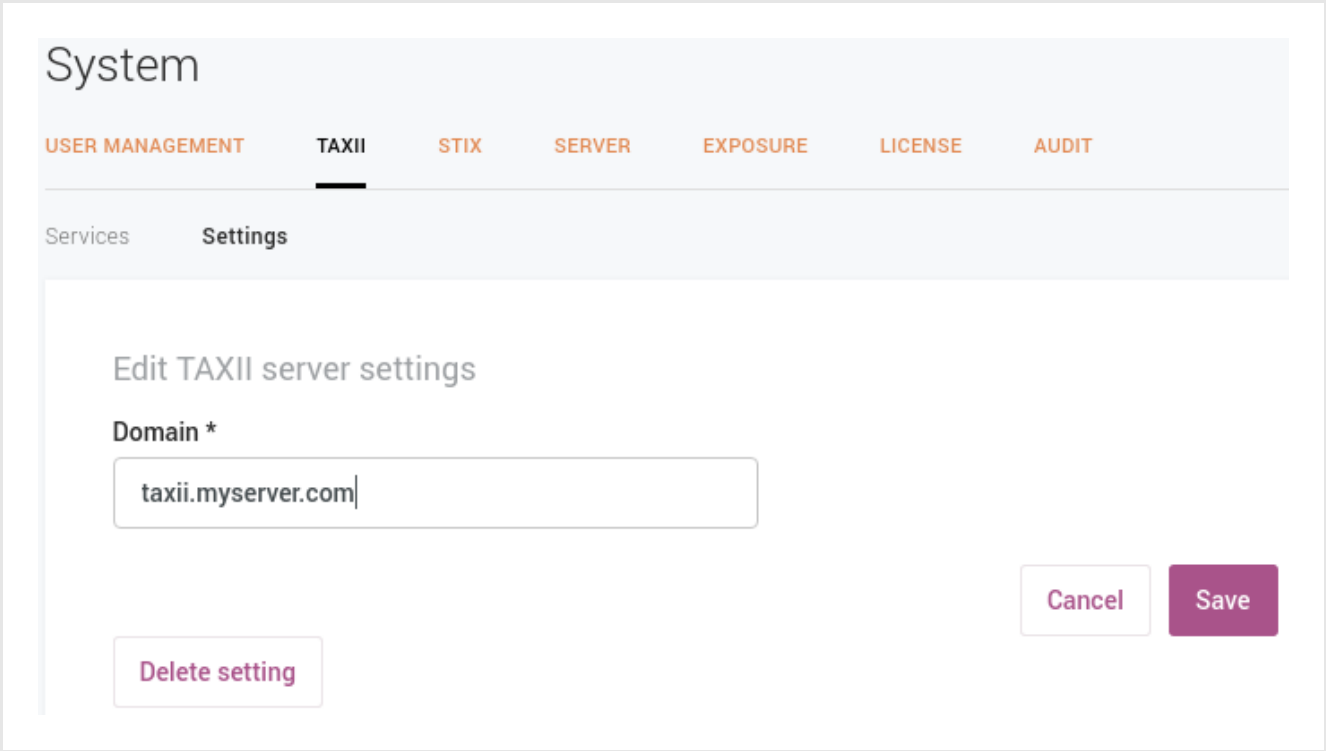

TAXII services

After configuring the 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	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 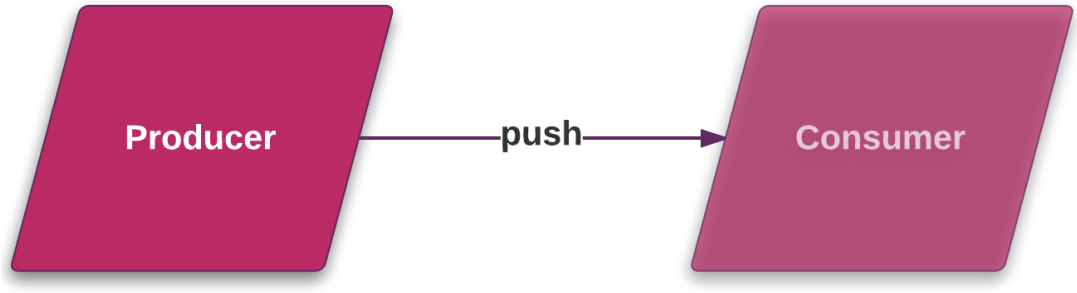

Service type	Description
Polling service	The TAXII poll service allows TAXII consumers to request TAXII data collection content from a TAXII producer, usually through TAXII outgoing feeds. Binding: TAXII HTTP 1.0, TAXII HTTPS 1.0 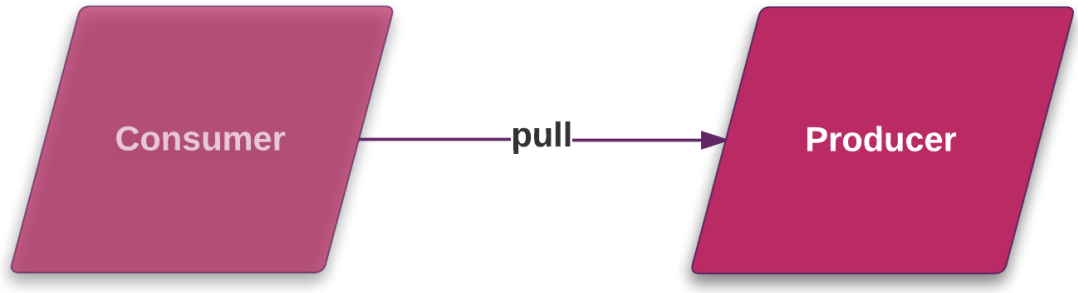

TAXII data collections (TAXII data feeds and TAXII datasets) are a typical example of inbox and polling service content.

Add a TAXII service

- On the left-hand navigation sidebar, click **System**.
- Click **TAXII**.
- Click the **+ Service** button.

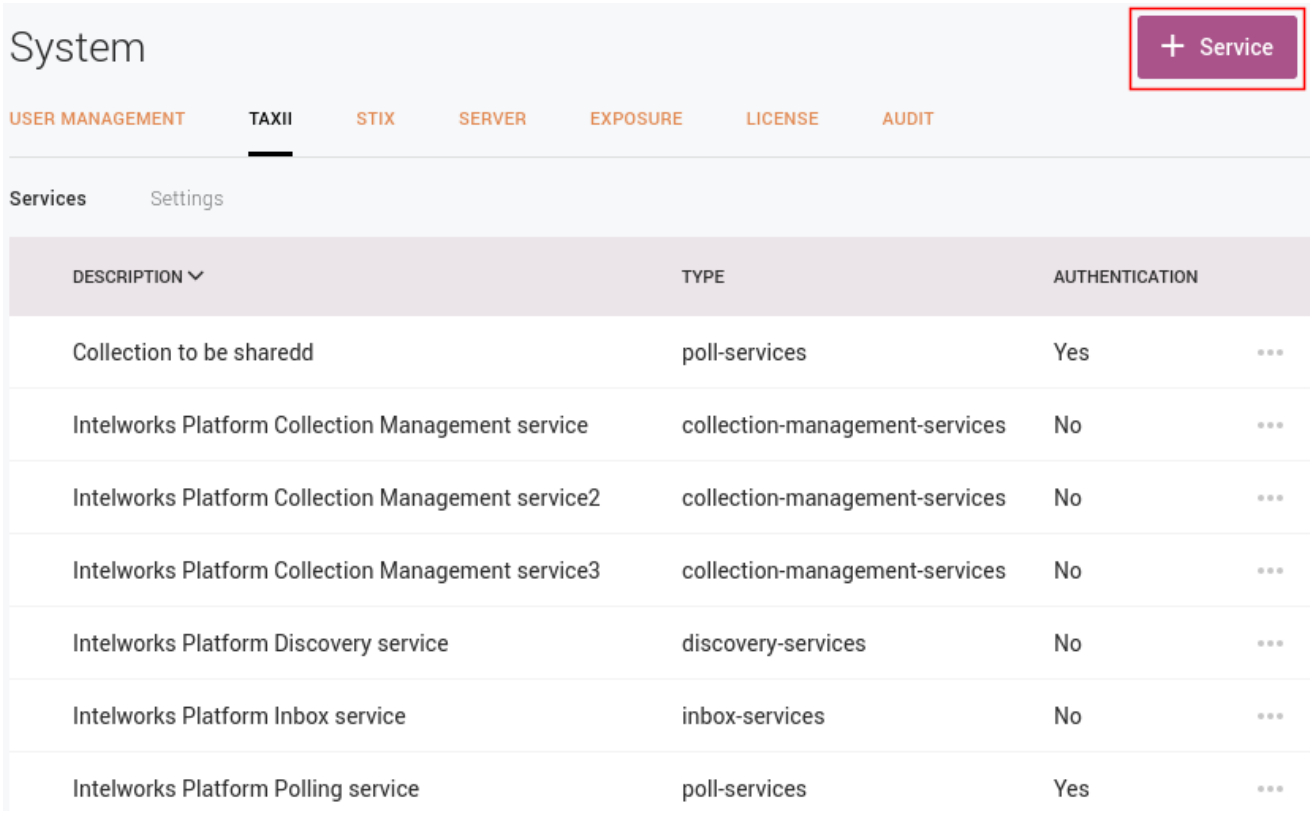

- Under **Add taxii service > Service type** , from the drop-down menu select the TAXII service type you want to add.

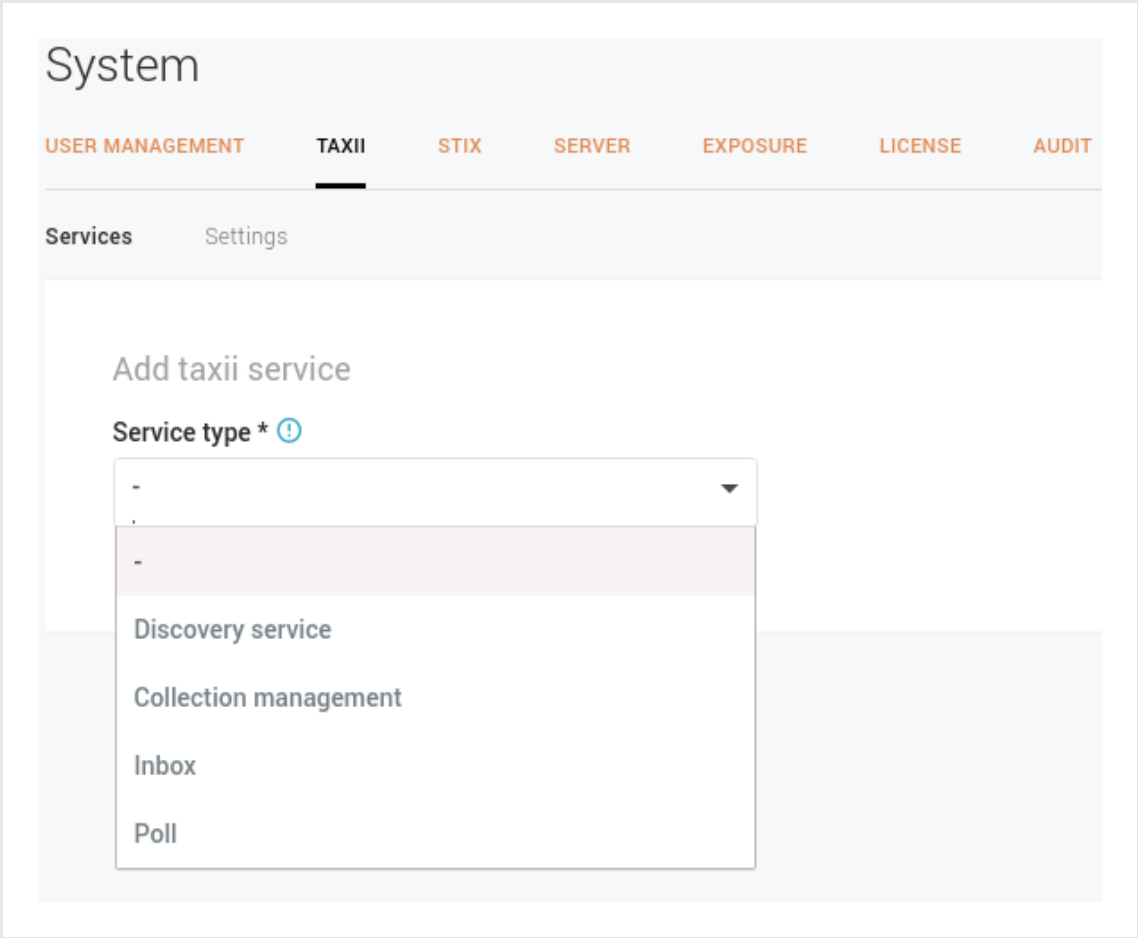

- Fill out the required fields:
  - **Description**: a free-text description of the service. It should be descriptive and easy to remember. Example: *Polling from XYZ*.
  - **Address**: the public endpoint the service can be reached at. Example: */taxii/services/polling*.
  - **Protocol bindings**: the data exchange transport protocol. Allowed values: *HTTP*, *HTTPS*.
  - **Authentication required**: select the checkbox to enable authentication, or deselect it to allow anonymous/guest access.

Besides these common settings for all services, each service type has 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. 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 you 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 available to the polling service.
  - **Max result size**: this option controls page size, so how many results each page can hold. We recommend keeping the amount of pages limited; therefore, we suggest 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 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 **System**.
- Click **TAXII**.

It shows an overview of the available TAXII services.

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 or edit the configuration information for a service, click the ⋮ icon on the row corresponding to the service.

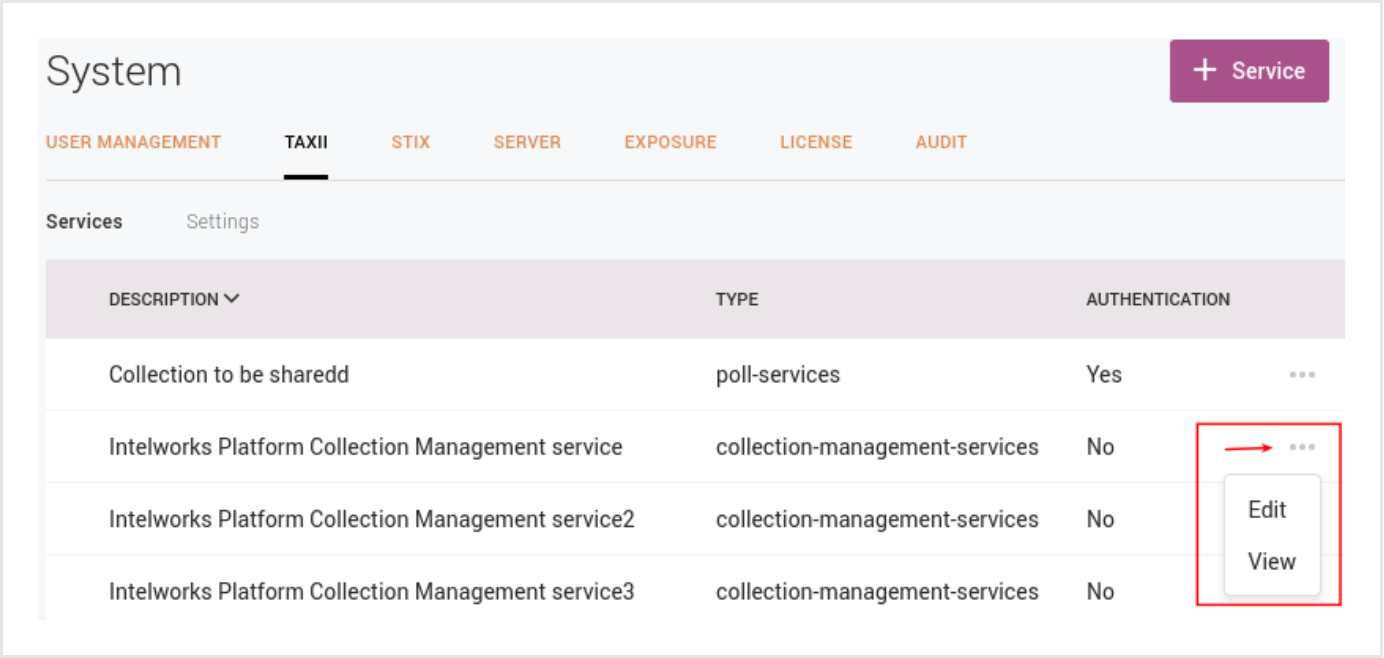

License

When you purchase a copy of the EclecticIQ Platform you receive a license key, which you use to register the product. To add a license key, do the following:

- On the left-hand navigation sidebar, click **System**.
- Click **License**.
- If there no license keys are registered, click the **add one** link under **License key**.

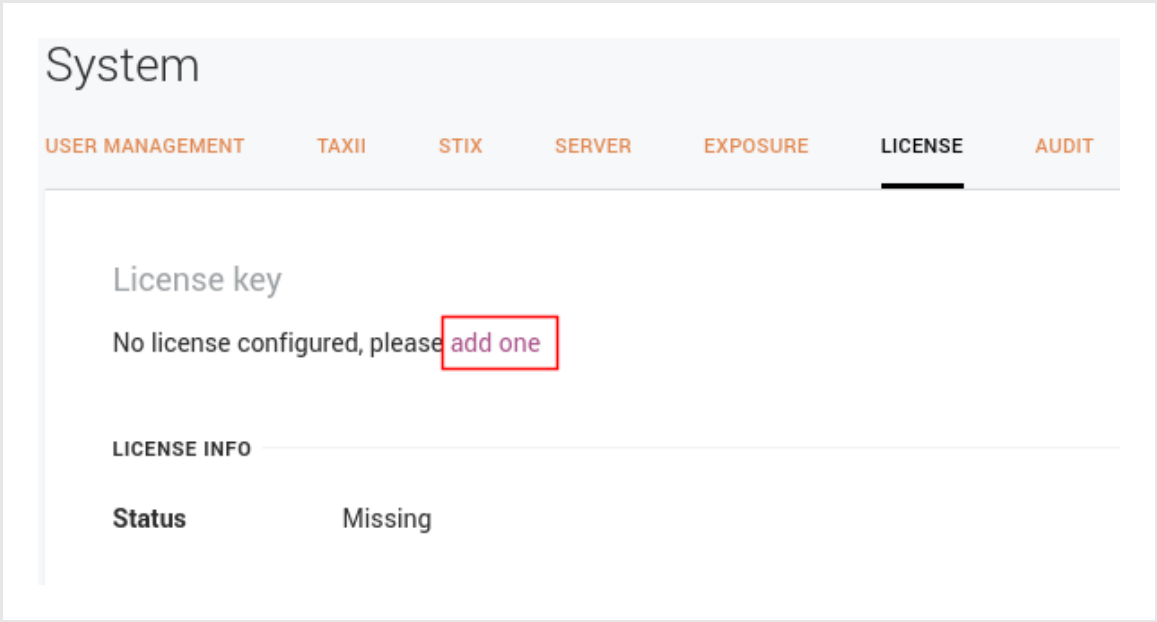

- In the input field under **Add license key**, copy-paste your license details.
- Click **Save**.

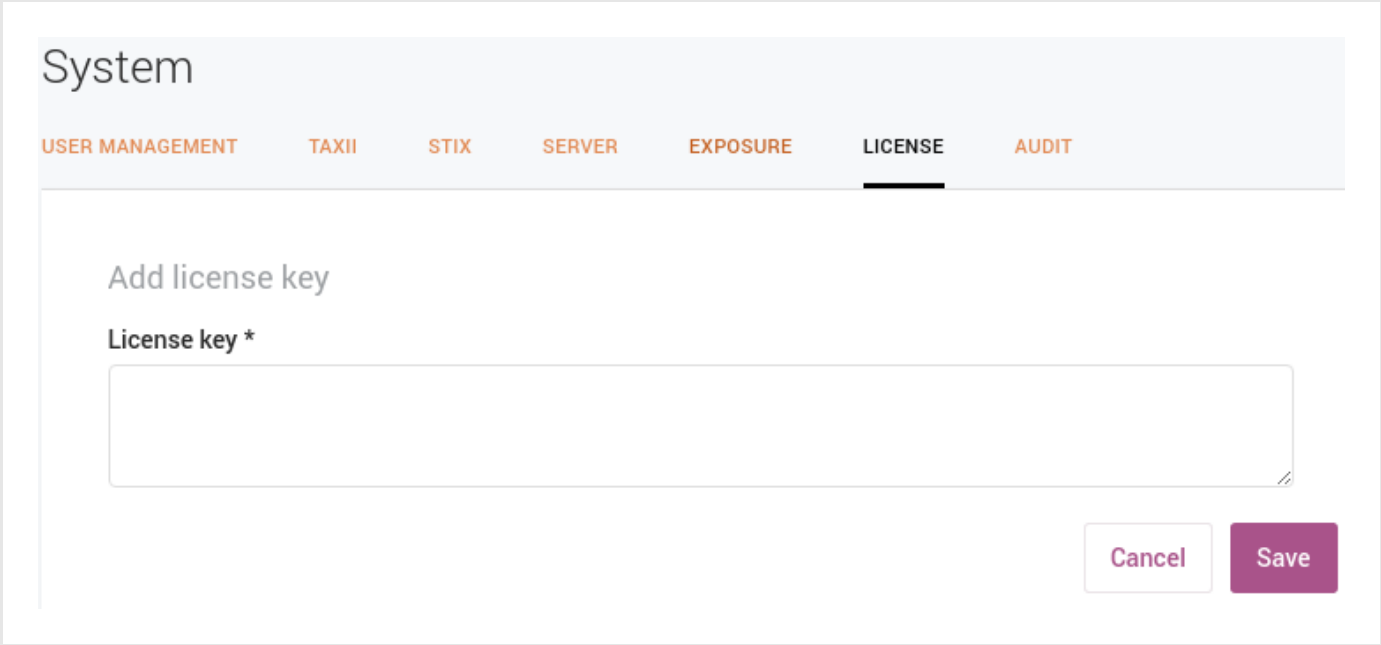

Valid license information populates the license information section:

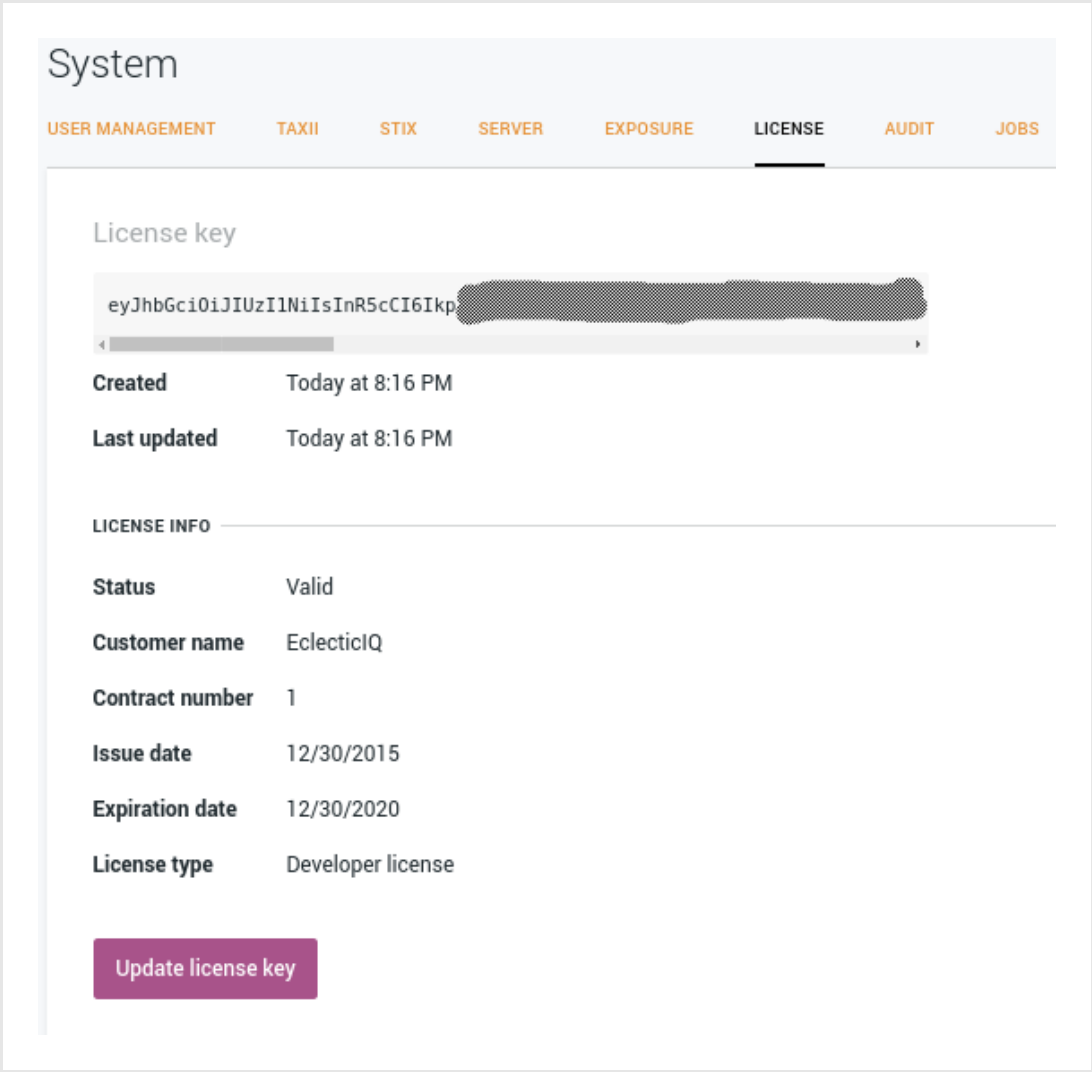

The license type is displayed on the status bar:

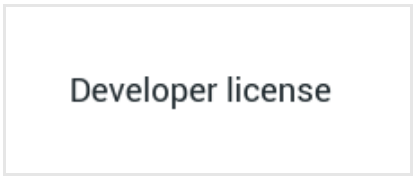

If the platform is unlicensed, a notification message is displayed on the status bar instead of the license type.

Monitor the system

You can monitor the system through the web-based GUI by accessing the following sections:

- Exposure gives you insight into ingested intelligence leveraging
- Audit offers an overview of system audit logs
- Jobs offers an overview of all platform tasks

Exposure

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

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 **Rules > Observable > ✚ Rule > Action > Mark as malicious > Confidence**.
- **Entity age**: it defines a time interval ranging from now, that it, 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.

After configuring it, you can start leveraging exposure.

Audit

The **System > Audit** tab provides a clear and searchable overview of system audit logs. To view audit logs in the platform web interface, do the following:

- On the left-hand navigation sidebar click ⚙ **> System settings > Audit**.
- If audit logging is enabled, and if the audit log file is populated, the matching audit log records are returned.
  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 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 specified user name(s).
Level	Displays only the search result items with the specified message level flag(s): **Info**, **Warning**, **Error**.
Method	Displays only the search result items with the specified HTTP method(s): **Delete**, **Post**, **Put**.
Response	Displays only the search result items with the specified HTTP response status code(s): *2xx*, *4xx*, or *5xx*.

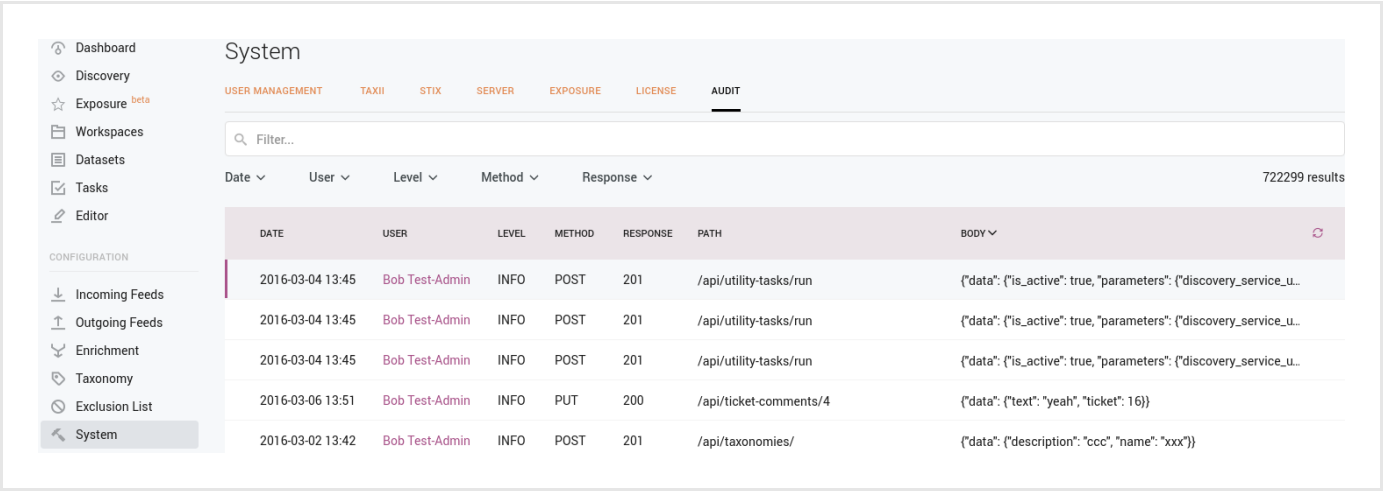

Jobs

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.

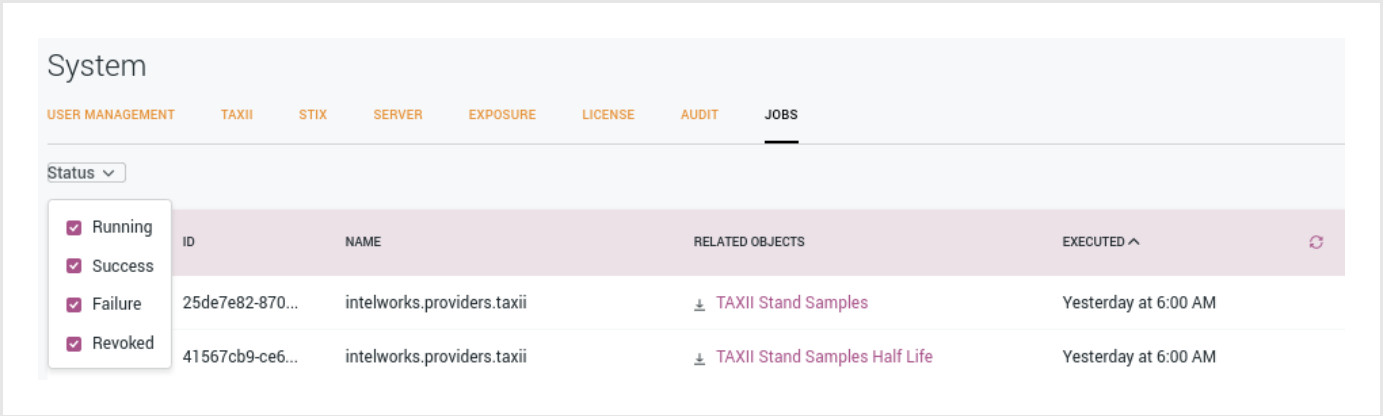

- **Related objects**: shows the platform objects the system job acts on. In this context, the objects are the channels the platform uses to ingest and to publish information: incoming and outgoing feeds, discovery, and entity or observable rules.
  On the job detail pane you can click a related object 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.

Configure users and roles

Configure users, their roles and permissions, and create user groups.

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.

Configure 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)) 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.

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).

Configure roles

To add a new user, do the following:

- On the left-hand navigation sidebar, click **System**.
- Under **User management**, click **Roles**.
- Under **Roles**, click the **+ Role** button.

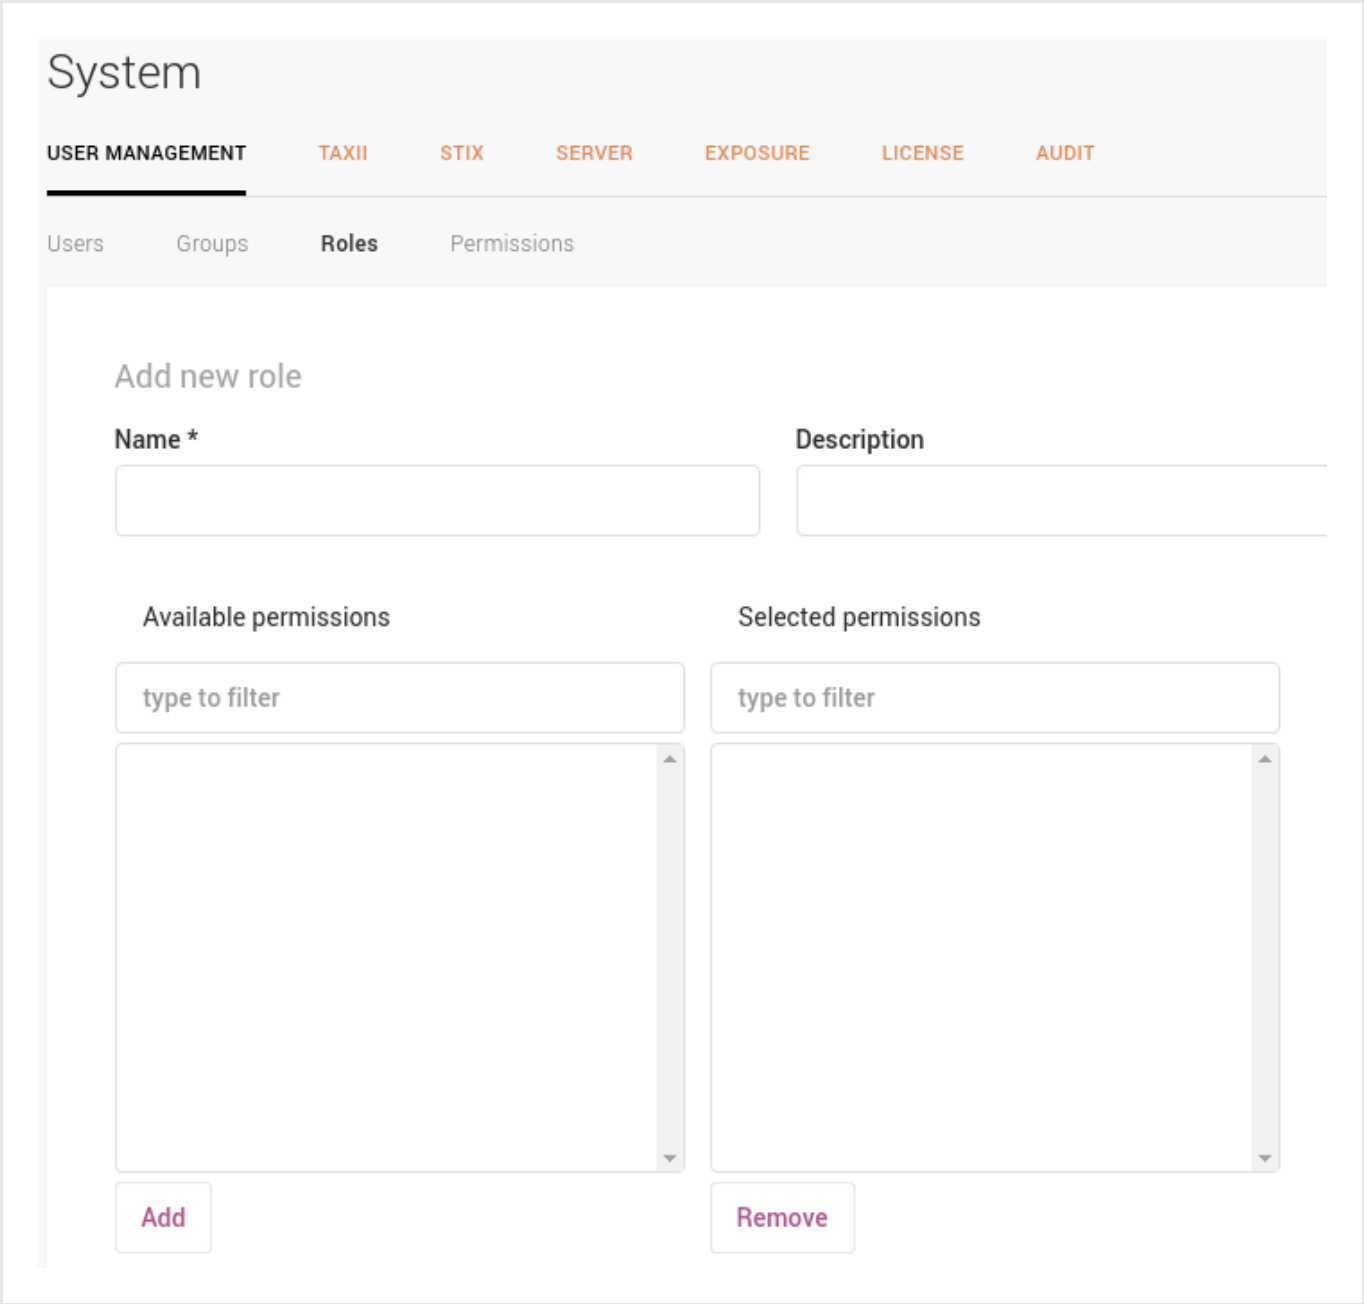

- Under **Add new role**, define the following configuration settings:
  - **Name**: a descriptive name for the role.
  - **Description**: a short description of the role and its purpose.
  - **Available permissions**: this pane lists all available permissions the new role can be granted.
    - Select one or more permissions from the list.
    - Click **Add** to grant the role the permission(s) listed in the **Selected Permissions** pane.
    - Alternatively, start typing a permission name in the autocomplete text input field above the pane.
    - Select one or more filtered permissions from the list.
    - Click **Add** to grant the role the permission(s) listed in the **Selected Permissions** pane.
    - To revoke one or more permissions for the role, select the relevant entries under **Selected permissions**, and then click **Remove**.
  - Click **Save** to store your changes, or **Cancel** to discard them.

Configure permissions

- Permissions are associated to 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.

Configure 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.

> ✔ On the forms, 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.

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 goup, 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).

Incoming feeds

Configure incoming feeds to ingest data from selected sources in many different formats.

When you launch the platform for the first time, the dashboard may look empty and uninformative. Therefore, one of the first things you want to do is ingest data. The platform can acquire data in several ways, one of them being through incoming feeds.

You can populate the platform with data by defining one or more incoming feed sources.
Once it is set up and it is running, an incoming feed provides a data stream that the platform ingests and processes automatically.

A minimal incoming feed configuration includes:

- A *data source*: the intel origin the incoming fed fetches data from.
  For example, a URI, a path pointing to a network location, or an IP address to an API endpoint.
- A *transport type*: the vehicle carrying the data.
  Typically, this is a communications protocol like TAXII, HTTP, FTP, or IMAP.
- A *content type*: the incoming data format the platform should expect from the incoming feed.
  For example, STIX, JSON, CSV, or PDF.

Content types

Content type	Feed type	Description
Anubis Cyberfeed JSON	in	JSON format representing entity data as JSON objects.
ArcSight CEF	out	The Common Event Format is a text-based standard for log records proposed by ArcSight. It allows sharing, consuming, and parsing event information across devices such as SIEM platforms and Syslog servers.
Cisco Threat Grid Samples JSON	in	JSON format representing entity data as JSON objects.
EclecticIQ Entities CSV	out	Comma separated CSV format for tabular data representation of entities.
EclecticIQ JSON	in, out	JSON format representing entity data as JSON objects.
EclecticIQ Observables CSV	out	Comma separated CSV format for tabular data representation of observables.
Group-IB accounts, Group-IB cards, Group-IB IMEIs	in	Group-IB proprietary data format to exchange information on compromised accounts, payment cards, and mobile devices.
Intel 471	in	Intel 471 proprietary data format.

Content type	Feed type	Description
PDF	in, out	Standard PDF format, preferably native (not scanned).
STIX 1.0	in, out	STIX data model **v. 1.0** (http://stixproject.github.io/data-model/1.0/).
STIX 1.1	in, out	STIX data model **v. 1.1** (http://stixproject.github.io/data-model/1.1/).
STIX 1.1.1	in, out	STIX data model **v. 1.1.1** (http://stixproject.github.io/data-model/1.1.1/).
STIX 1.2	in, out	STIX data model **v. 1.2** (http://stixproject.github.io/data-model/1.2/).
Text/Plain text value	in, out	Plain text format. This content type allows you to enter free text and literals, wildcards (where supported), as well as JSON paths to point to specific entity property fields, and regex patterns to filter data.
Threat Recon	in	Threat Recon JSON output returned by the **Threat Recon API** (https://threatrecon.co/api). Threat Recon focuses on providiung information about indicators.
FireEye iSIGHT STIX 1.1.1	in	FireEye iSIGHT Intelligence Report API outputs reports in STIX 1.1.1 format. Reports concern threat topics such as vulnerabilities, malware, threat actors, stategies, tactics, and techniques.

Transport types

Transport type	Feed type	Description
Anubis Cyberfeed	in	Provides data on bank Trojans, compromised DNS servers, malware-infected web site and malware files.
Cisco Threat Grid Curated Feed	in	Provides data on compromised IP addresses, domains, hashes, registry keys, and network streams.
Cisco Threat Grid Samples API	in	Allows submitting malware samples for analysis, as well as investigating a domain, an IP, or a URL to obtain information about potential threats.
FTP download	in	Custom feed ingesting data through FTP.
Group-IB JSON API	in	*Accounts*: provides information on compromised logins, passwords, corporate email accounts, and so on. *Cards*: provides information on compromised bank card numbers and online banking keys. *IMEI*: provides information on compromised mobile devices like IMEI/IMSI, and the ICCID of compromised SIM cards.
HTTP download	in	Custom feed ingesting data through HTTP.

Transport type	Feed type	Description
IMAP email fetcher	in	Custom feed using the IMAP email protocol to ingest data included in emails as attachments.
Intel 471 API	in	Provides data on compromised IP addresses, domains, URLs emails, with a focus on threat actors.
Mount point download	in	Custom feed using a local or a network drive as a data source.
TAXII inbox	in	Custom feed ingesting data through the TAXII inbox service.
TAXII poll	in	Custom feed ingesting data through the TAXII poll service.
Threat Recon JSON API	in	Provides data on compromised IP addresses, domains, as well as whois information.
FireEye iSIGHT Intelligence Report API	in	Provides intelligence reports on vulnerabilities, malware, and other threats such as threat actors, stategies, tactics, and techniques.
IMAP email attachment fetcher	in	Ingests email attachments that are then transformed into reports. The reports can be used for analysis (human consumption) or instrumentation (machine consumption/automation).
Crowdstrike Falcon Intelligence Indicator Feed	in	Ingests information about indicators by polling the Falcon Intel Indicator API. It is possible to set a start date to poll data from.

Configure incoming feeds

You can configure incoming feeds to acquire and ingest a stream of cyber threat data from one or more source producers.

Set up incoming feeds

To set up an incoming feed to populate the platform with entities, do the following:

- On the top navigation bar, select **Configuration**, and then **Incoming feeds**.
- On the top-left corner of the page click the ⊕ icon to open the incoming feed editor.

The **Incoming feeds** page displays an overview of the configured incoming feeds ingesting data from the specified intel providers and data sources.

- On the top-right corner of the screen, click the **+ Incoming feed** button.
- On the **✚ > Data management > Incoming feeds > Create** form you can populate the input fields to define the intel provider/data source for the feed, and the feed behavior.

> ✔ On the forms, input fields marked with an asterisk are required.

Field	Type	Description	Example
Name	String, alphanum. [A-Z a-z][0-9]	*Required* — The name you assign to the incoming feed. On the forms, input fields marked with an asterisk are required.	*EclecticIQ threat feed*
Organization	String, alphanum. [A-Z a-z][0-9]	*Required* — The name of the source organization that serves as the source for the incoming feed.	*EclecticIQ*
Override TLP color	Radio button(s)	You can override any existing TLP value and assign a custom TLP color code to all the entities ingested through the incoming feed.	*Amber*
Source reliability	Single choice drop-down menu	You can choose a value from the drop-down menu to flag the level of reliability of the source.	*B - Usually reliable*
Content type	Single choice drop-down menu	*Required* — It defines the data format of the incoming feed. Its value needs to match the actual feed file format.	*STIX 1.1*
Transport type	Single choice drop-down menu	*Required* — It defines the protocol used to carry the data.	*TAXII poll*

Field	Type	Description	Example
Authorized groups	Single choice drop-down menu	You can restrict access to this feed to a single user group.	*Analysts*

- Depending on the selected transport type, you may need to specify additional settings under **Transport configuration**. For example:
  - A URL endpoint corresponding to the API service exposing the data source for the incoming feed.
  - A valid API key to grant you access to the feed data source.
  - Any required login credentials to obtain access to the feed data source.
- After populating the fields with the necessary details, click the **Save** button to save the newly created feed and to make it available.
- The new feed is now included in the overview table on the **Incoming feeds** page.

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.

Run feeds

You can schedule feed tasks to run at specific times, as well as manually trigger feed task execution.

If you set an execution schedule for a feed, you don’t need to do much, unless you want to poll the feed origin right away to retrieve the corresponding feed data.

Manually run a feed task

To manually run a feed task, do the following:

- Sign in to the platform.
- - On the top navigation bar, select **Configuration**, and then **Incoming feeds**.
  - On the top-left corner of the page click the ⊕ icon to open the incoming feed editor.

- The **Incoming feeds** page displays an overview of the configured incoming feeds ingesting data from the specified intel providers and data sources.
- Click any area on the row corresponding to the feed you want to run.
- On the feed detail pane, select the **Logs** tab.
- On the **Logs** tab, click the **Run now** button.
- The task is executed.

Check task results

- On the **Logs** tab you can check the task status under **Status**.
- On the **Content** tab you can examine all the entities retrieved so far with the feed.
- To see how many entities have been ingested in total in the platform, go to the dashboard.

Configure outgoing feeds

You can configure outgoing feeds to relay a stream of cyber threat data that you can share cross-teams within your organization, or make available to third-parties.

Set up outgoing feeds

To set up an outgoing feed and make entities available for retrieval, do the following:

> ✔ On the forms, input fields marked with an asterisk are required.

Under **Transport and content** you can define *what* you want to publish and *how*, that is, the data content type and the data transport type.

- Under **Feed name**, enter a name for the feed you are creating. It should be descriptive and easy to remember.
- **Transport type**: from the drop-down menu select the appropriate transport type to publish data through the outgoing feed. This can vary, based on the carrier used to distribute the data.
- Depending on the selected transport type, you may need to specify additional settings under **Transport configuration**. For example:
  - A URL endpoint corresponding to the API service exposing the data source for the incoming feed.
  - A valid API key to grant you access to the feed data source.
  - Any required login credentials to obtain access to the feed data source.
- **Content type**: from the drop-down menu select a content type that matches the data format for the feed and configure the appropriate parameters under **Content configuration**, when applicable.
- **Dataset**: from the drop-down menu select one or more datasets as data sources for the outgoing feed.
- **Update strategy**: from the drop-down menu select the preferred method to update the data:
  - **Append**: every time the outgoing feed task runs, only new data from the latest task run, that is, only new entities, is appended to the existing data.
    When the outgoing feed task runs, it includes only new entities.
  - **Replace** every time the outgoing feed task runs, it publishes only new data.
    When the outgoing feed task runs, it produces new content that can include new, as well as existing entities.

Set a schedule

Under **Execution schedule** you can define how often you want to run the feed task:

- **None**: no schedule is defined. You need to manually trigger the task to ingest or to publish data through an incoming or an outgoing feed, respectively.
- **Minute**: the feed task runs automatically every *N* minutes, where *N* defines the selected time interval in minutes.
  You define the execution interval in 5-minute increments from the corresponding drop-down menu.

- **Hour**: the feed task runs automatically every hour.
  You define how long in minutes after the beginning of an hour the task should run from the corresponding drop-down menu.
- **Day**: the feed task runs automatically once a day.
  You define the time of the day when the task should run from the corresponding drop-down menu.
- **Week**: the feed task runs automatically once a week.
  You define the day of the week and time of the day when the task should run from the corresponding drop-down menu.
- **Month**: the feed task runs automatically once a month.
  You define the day of the calendar month and time of the day when the task should run from the corresponding drop-down menu.
  Keep in mind that not all months of the year have 31 days.

Set a TLP override

- **Override TLP** overwrites the **TLP** `(https://www.us-cert.gov/tlp)` color code associated to the feed entities with the one you set here. The selected TLP value is assigned to all the entities in the feed.

  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.
- The **Filter TLP color** options allow including in the feed data only an entity subset, based on the selected **TLP** `(https://www.us-cert.gov/tlp)` value.
  If you set a TLP color as a filter, the feed includes all the entities flagged with the selected TLP color code, as well as the entities whose TLP color indicates that they are suitable for progressively broader audiences. For example, if you select green, the feed includes entities with a TLP color set to green and entities with a TLP color set to white.

Set reliability and relevancy

- **Source reliability**: from the drop-down menu select an option to flag the content of the outgoing feed with a predefined reliability value to help other users assess how trustworthy the feed source is.
  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)`.
  Example: *B - Usually reliable*
- **Relevancy threshold (%)** allows you to set a filter to include in the feed only entities whose relevancy is higher than the value defined here.

Set observable filters

- **Allowed observable states**: from the drop-down menu select one or more observable states to include in the feed data only entities whose observable states match at least one of the selections defined here.
- **Observable types**: from the drop-down menu select one or more observable types to include in the outgoing feed only entities whose observable types match at least one of the selections defined here.
- **Enrichment observable types**: from the drop-down menu select one or more enrichment observable types to include in the outgoing feed only entities whose enrichment observable types match at least one of the selections defined here.
- Click **Save** to store your changes, or **Cancel** to discard them.

The filters work independently of each other: there are no Boolean `AND` or `OR` to join multiple filters into a serial pipeline.

The new outgoing feed is now included in the overview table on the **Outgoing feeds** page.

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.

Depending on the selected transport type, you may need to specify a URI:

- A URL endpoint corresponding to the source of the outgoing feed.
- You may need to provide also any relevant login credentials to be granted access to the source of the outgoing feed.

Enrichment

Enrichment improves the quality of the intelligence you obtain from cyber data analysis. Enrich entities and integrate entity observables with additional raw data to access a broader context and gain deeper insight.

Enriching entities via observables

The platform can ingest cyber threat intelligence through incoming feeds, by manually uploading one or more files, or by creating an entity in the entity editor.

After ingesting and saving entities to the database, you can integrate the existing information with additional details. The extra information is raw data that augments the entity intelligence value by adding more context and meaning to it. The data is extracted from different sources such as feeds, reports, database searches, curated intel distribution lists, and so on.

The platform uses enrichers to fetch and extract the data. Enricher rules sift through the data to link it to relevant entities as enrichment observables.
This process does not alter core entity data: each bit of enriching information is saved to observables, which are related to entities.

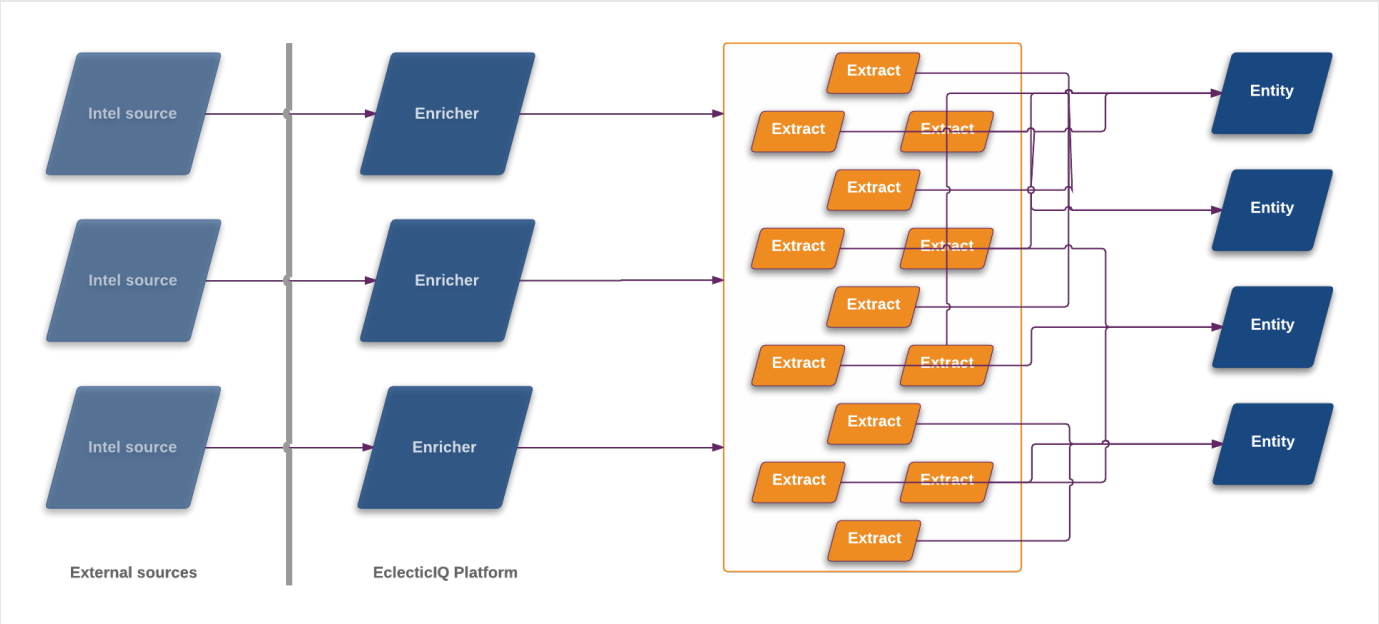

For example, let's assume a scenario where an analyst is investigating a threat actor entity. The entity includes some observables, and one of them is a DNS domain name.
The analyst looks up the domain name by running it through a whois service. The lookup results include an email address. During the investigation, the analyst retrieves also a file hash related to the domain name. An examination reveals that the hash is related to an incident. Information about the incident includes the same email address detail the DNS domain name returned.
There is an indirect relationship between the threat actor and the incident that would not have been noticeable without extra context, which in this example is provided by the hash.

Enrichments help get a broader and sharper picture: by adding meaningful context, they help discover broader, indirect relationships that are not immediately visible.

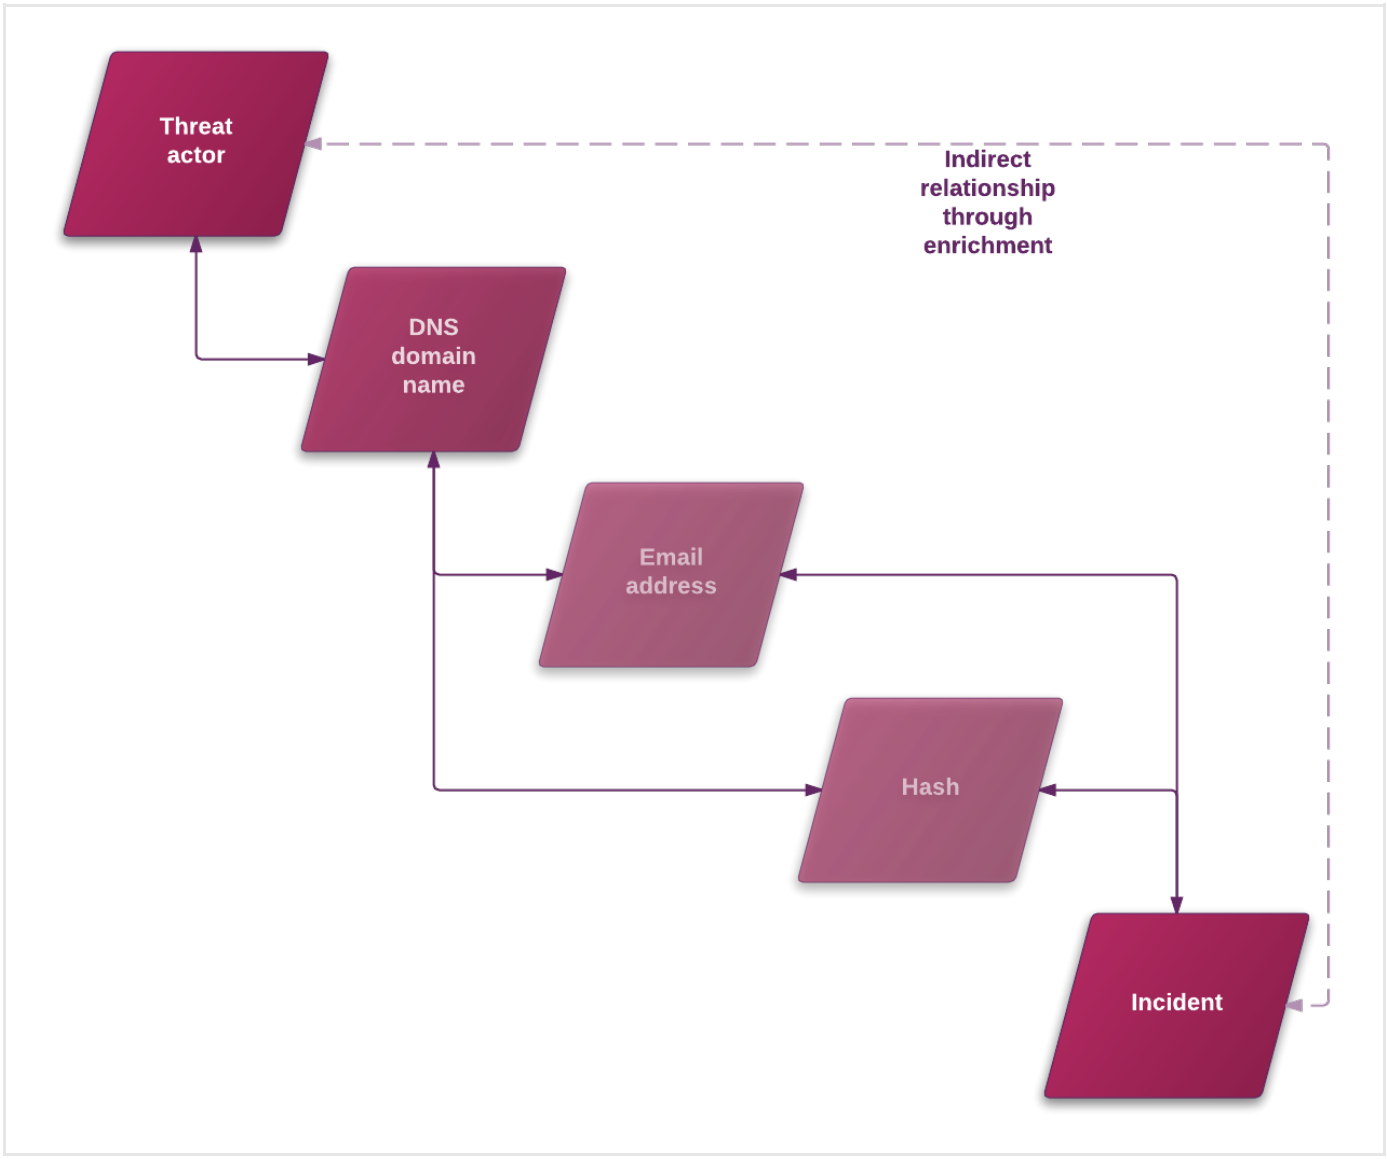

Enrichments augment observables with raw data information related to entities:

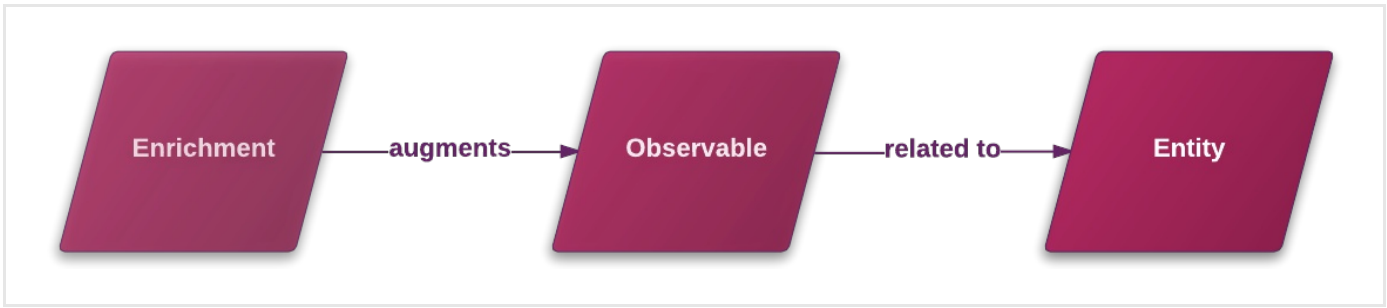

Enrich entities

You can enrich entities in the following ways:

- Automatically, or
- Manually.

Enrichment rules and enrichment tasks drive the enrichment process to:

- Poll selected and trustworthy intelligence data sources;
- Retrieve relevant, accurate, and reliable data to augment platform entities with additional bits of information that provide additional context.

Rules

Enrichment rules define what to do with the retrieved enrichment data.
Rules act like filters, and they set the logical constraints defining:

- The platform data sources to augment with the enrichment information. Data sources can be incoming feeds, as well as other enrichers.
- Within the selected platform data sources, the entity type(s) to augment with the enrichment information.
- The enrichers to use to fetch the enrichment data.

Tasks

Enrichment tasks define process execution by setting the following options:

- The data fetching mechanism; for example, an API endpoint exposing the enrichment data service.
- Specific data sources; for example, datasets targeting threat actors like hackers and terrorist groups.
- Data rate limit and monthly execution cap values to control the amount of polled data.
- A source reliability flag for the incoming enrichment data to simplify assessing the quality of the retrieved data.

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.

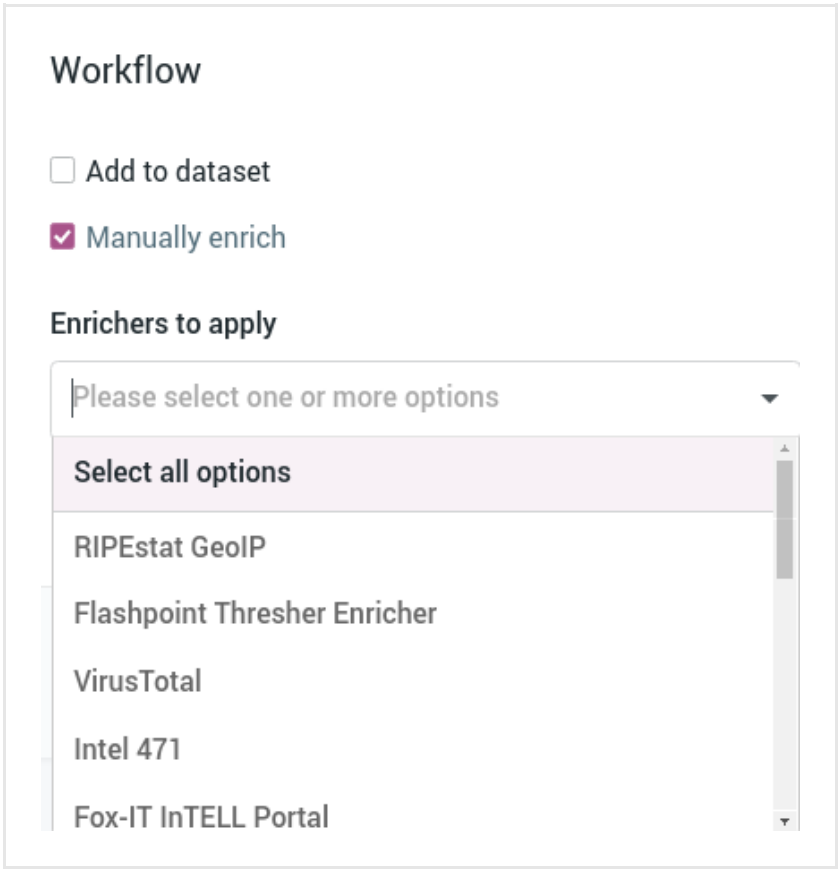

- Click **Save draft** to store your changes without publishing the entity, **Publish** to release the new version of the entity including your changes, or **Cancel** to discard the changes.

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 **Enrich** drop-down menu, select **Enrich all observables**.
- The platform polls all applicable enrichers for the entity, and it enriches all the entity observables with the retrieved data.

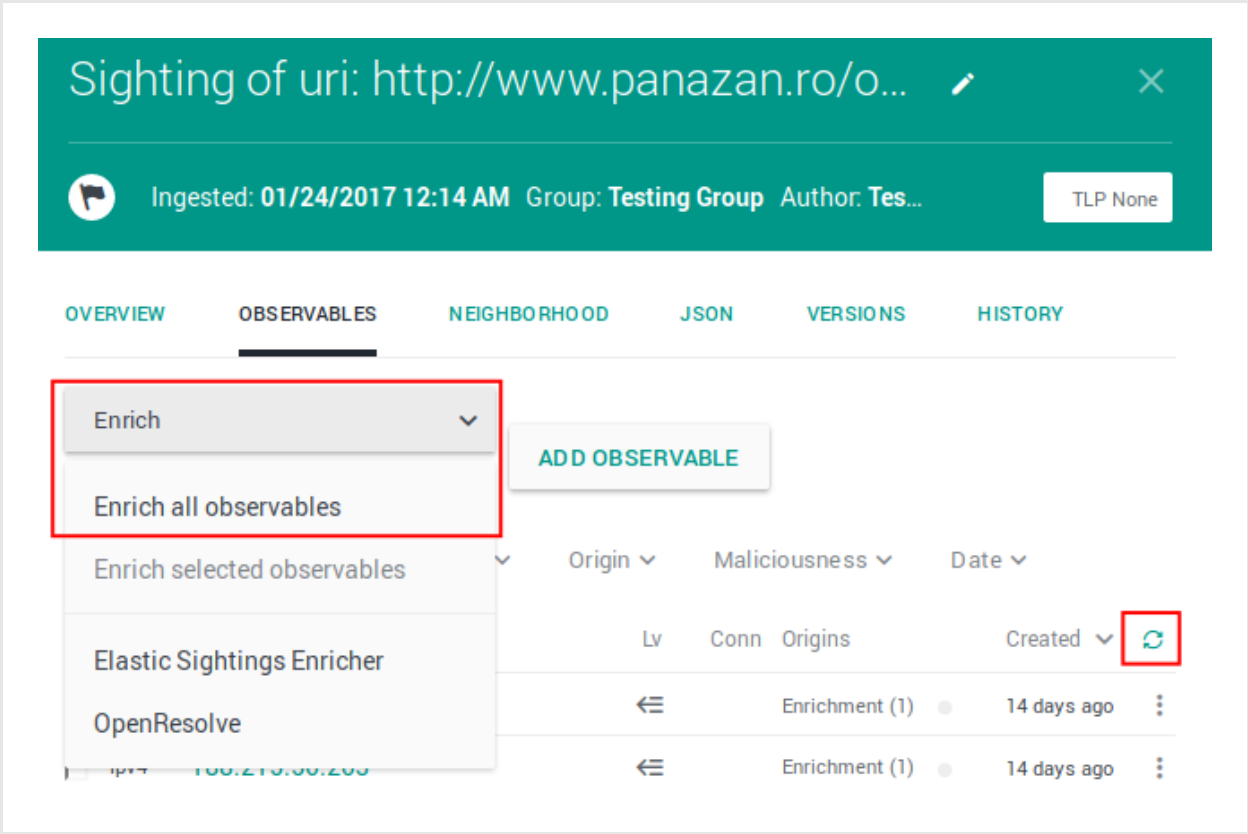

To poll a specific enricher:

- Select it from the **Enrich** drop-down menu, and then click it.
- The platform polls the specified enricher for the entity, and it enriches all the entity observables with the retrieved data.

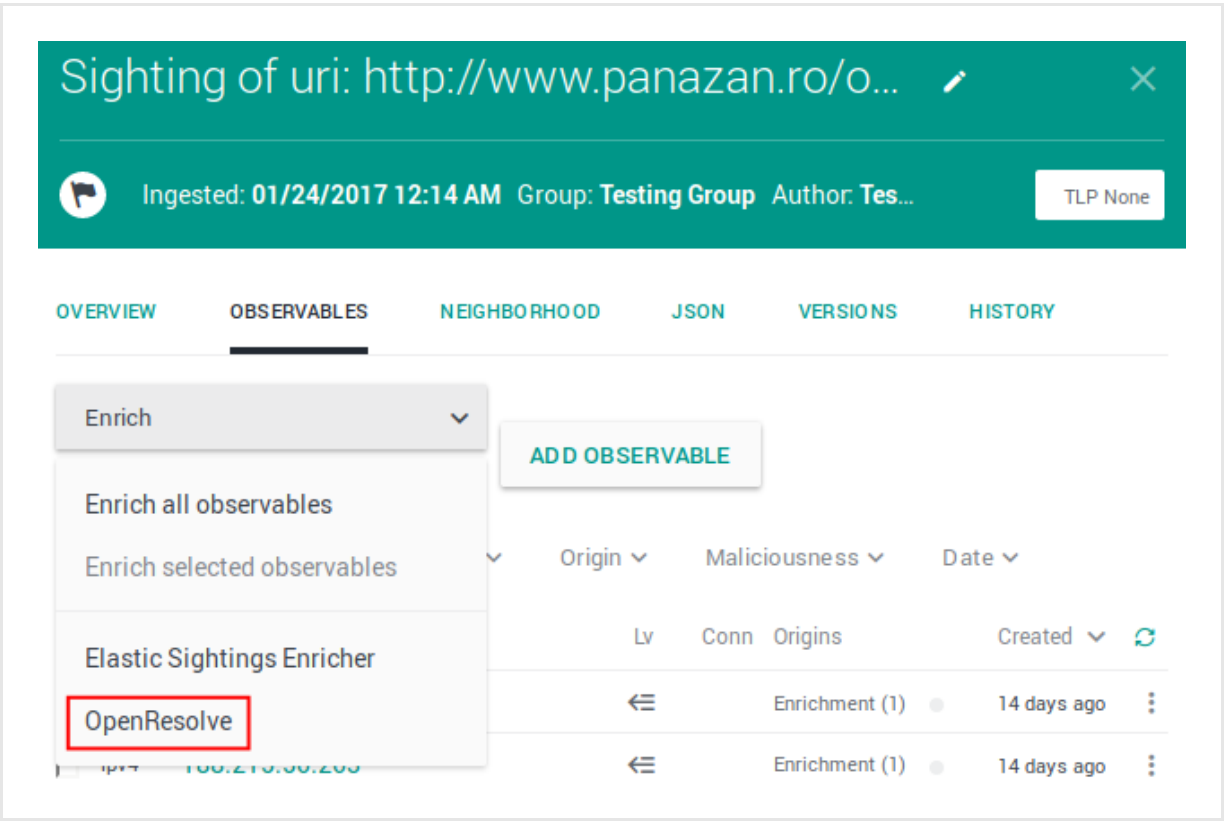

To enrich only specific observables:

- On the **Observables** tab, select the checkboxes corresponding to the observables you want to enrich.

- From the **Enrich** drop-down menu, select **Enrich selected observables**.
- The platform polls all applicable enrichers for the entity, and it enriches the selected entity observables with the retrieved data.

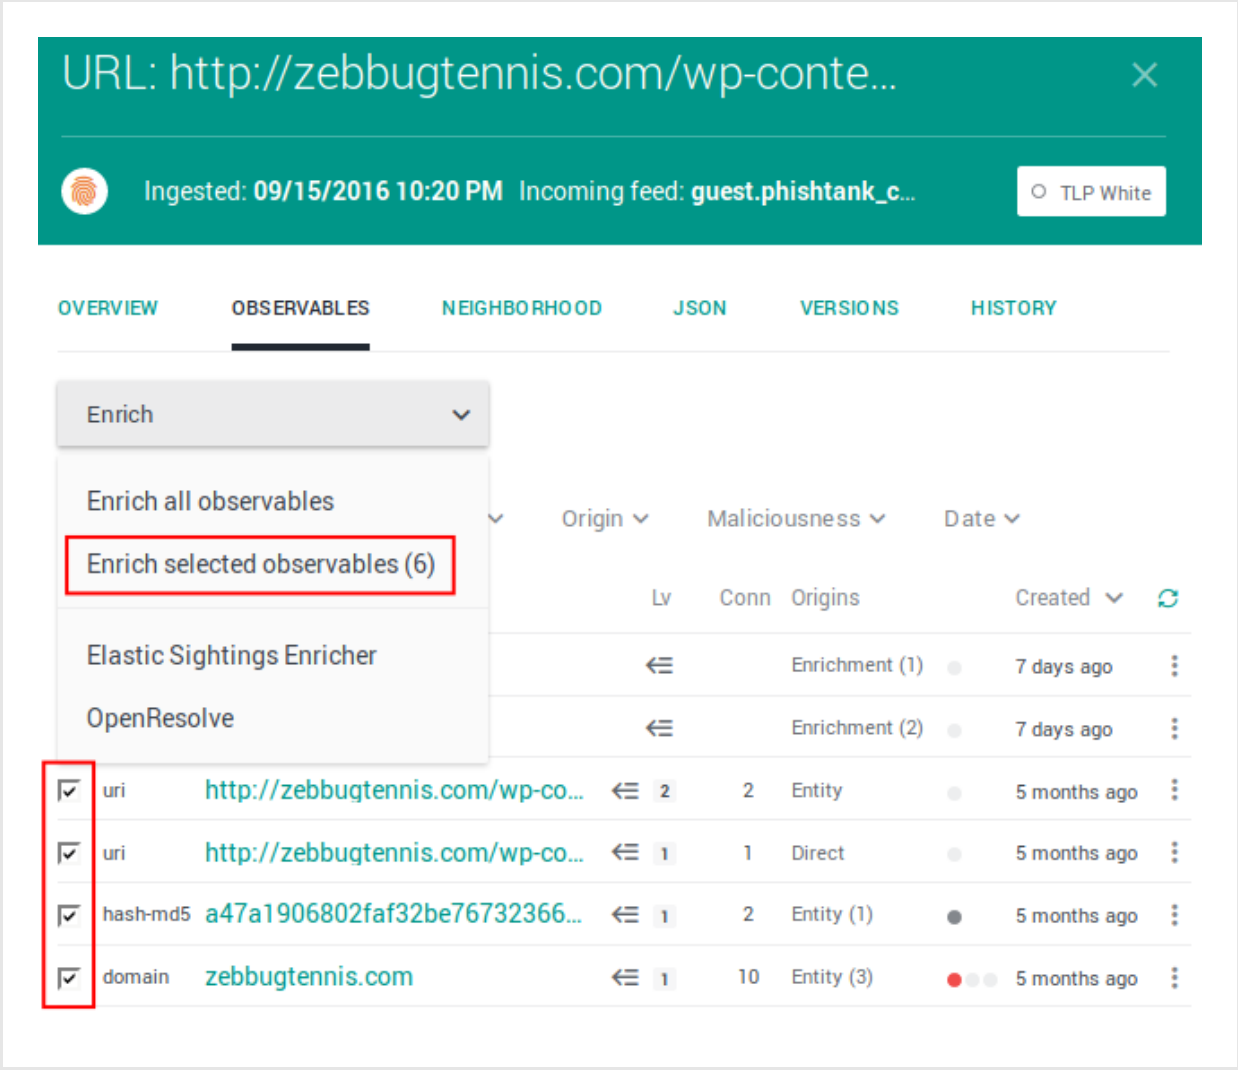

The available enricher tasks in the drop-down menu are automatically filtered to show only the applicable enrichers for the entity.
Enrichers automatically augment all the entities that accept the enricher's content type as an observable. In other words, the observable types an entity supports define the applicable enrichers an entity can use.

Enricher rules

View enricher rules

To view enricher rules, do the following:

- On the top navigation bar, click the ⚙ icon next to the user avatar image.
- 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 enricher rules

To add a new enricher rule, do the following:

- On the top navigation bar click **✚ > Rules > Enrichment**.

Alternatively:

- On the top navigation bar, click the ⚙ icon next to the user avatar image.
- 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.
- Click the **+ Rule** button.

> ✔ On the forms, input fields marked with an asterisk are required.

On the **Rules > Enrichment > Create** 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 or the enricher whose observables you want to augment with additional information.
- **Entity types**: from the drop-down menu select the entity type whose observables 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 enricher rules

To edit enricher rules, do the following:

- On the top navigation bar, click the ⚙ icon next to the user avatar image.
- 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 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 detail pane, click **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**.

> ✔ On the forms, 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 or the enricher whose observables you want to augment with additional information.
- **Entity types**: from the drop-down menu select the entity type whose observables 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 enricher rules

To delete an enricher rule, do the following:

- On the top navigation bar, click the ⚙ icon next to the user avatar image.
- 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.
- 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.
- Click **Delete** on the rule detail pane.

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.

Enricher tasks

View enricher tasks

To view enricher tasks, do the following:

- On the top navigation bar click ✚ **> Data management > Dataset > Enrichment** .

Alternatively:

- On the top navigation bar, click the ⚙ icon next to the user avatar image.

- From the drop-down menu select **Data management**.
- On the left-hand navigation sidebar click **Enrichment**.
- Click the enricher you want to examine.
- On the enricher detail page, you can view all the details about the selected enricher, including the rules driving the enricher behavior, recently executed enriching tasks, and the state.
- You can click the state value or an enrichment rule to display additional information.

> When the state value returns **FAILURE**, click the link to view the task execution traceback and to begin troubleshooting.

The **Data management > Enrichment** view shows all configured enrichers polling third-party and/or external services to acquire additional information to integrate observables with, so that they can provide more context to the cyber threat entities they belong to.

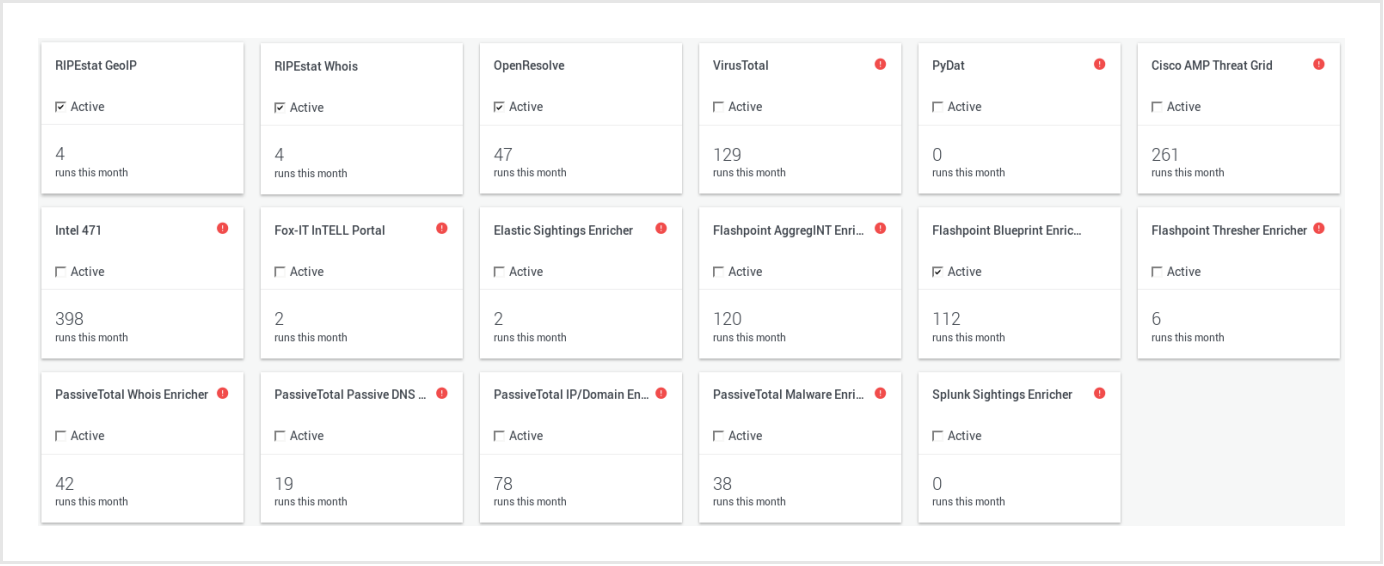

Edit enricher tasks

To configure or to edit an enricher task, do the following:

- On the top navigation bar click **+ > Data management > Dataset > Enrichment** .

Alternatively:

- On the top navigation bar, click the ⚙ icon next to the user avatar image.
- From the drop-down menu select **Data management**.
- On the left-hand navigation sidebar click **Enrichment**.
- Click the enricher you want to configure or modify.
- On the enricher detail page, click the **Edit** button.

> ✔ On the forms, input fields marked with an asterisk are required.

Warning:
Some enricher tasks include an additional API key field where you specify the API key issued by the source of the enricher, along with the necessary authentication and authorization credentials.
Contact the intel service provider whose data you want to use as a source for the enricher to request an API key and any other required credentials.

You need to install and set up PyDat locally. The product does not work outside a local network.
You need to configure the host before you can access PyDat features through the API endpoint.
See also:

- **Mitre blog on PyDat** `(http://www.mitre.org/capabilities/cybersecurity/overview/cybersecurity-blog/using-whois-and-passive-dns-for-intelligence)`
- **PyDat GitHub repo** `(https://github.com/mitrecnd/whodat)`

Taxonomy

The Taxonomy page offers an overview of the tags used to label entities in the platform. Besides using tags to organize entities, you can design taxonomies to structure the tags, and to create a controlled tag corpus to improve information retrieval.

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.

The Taxonomy feature

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.

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
Fairly reliable	Possibly True
Not usually reliable	Doubtful
Reliability cannot be judged	Improbable
Unreliable	Truth cannot be judged

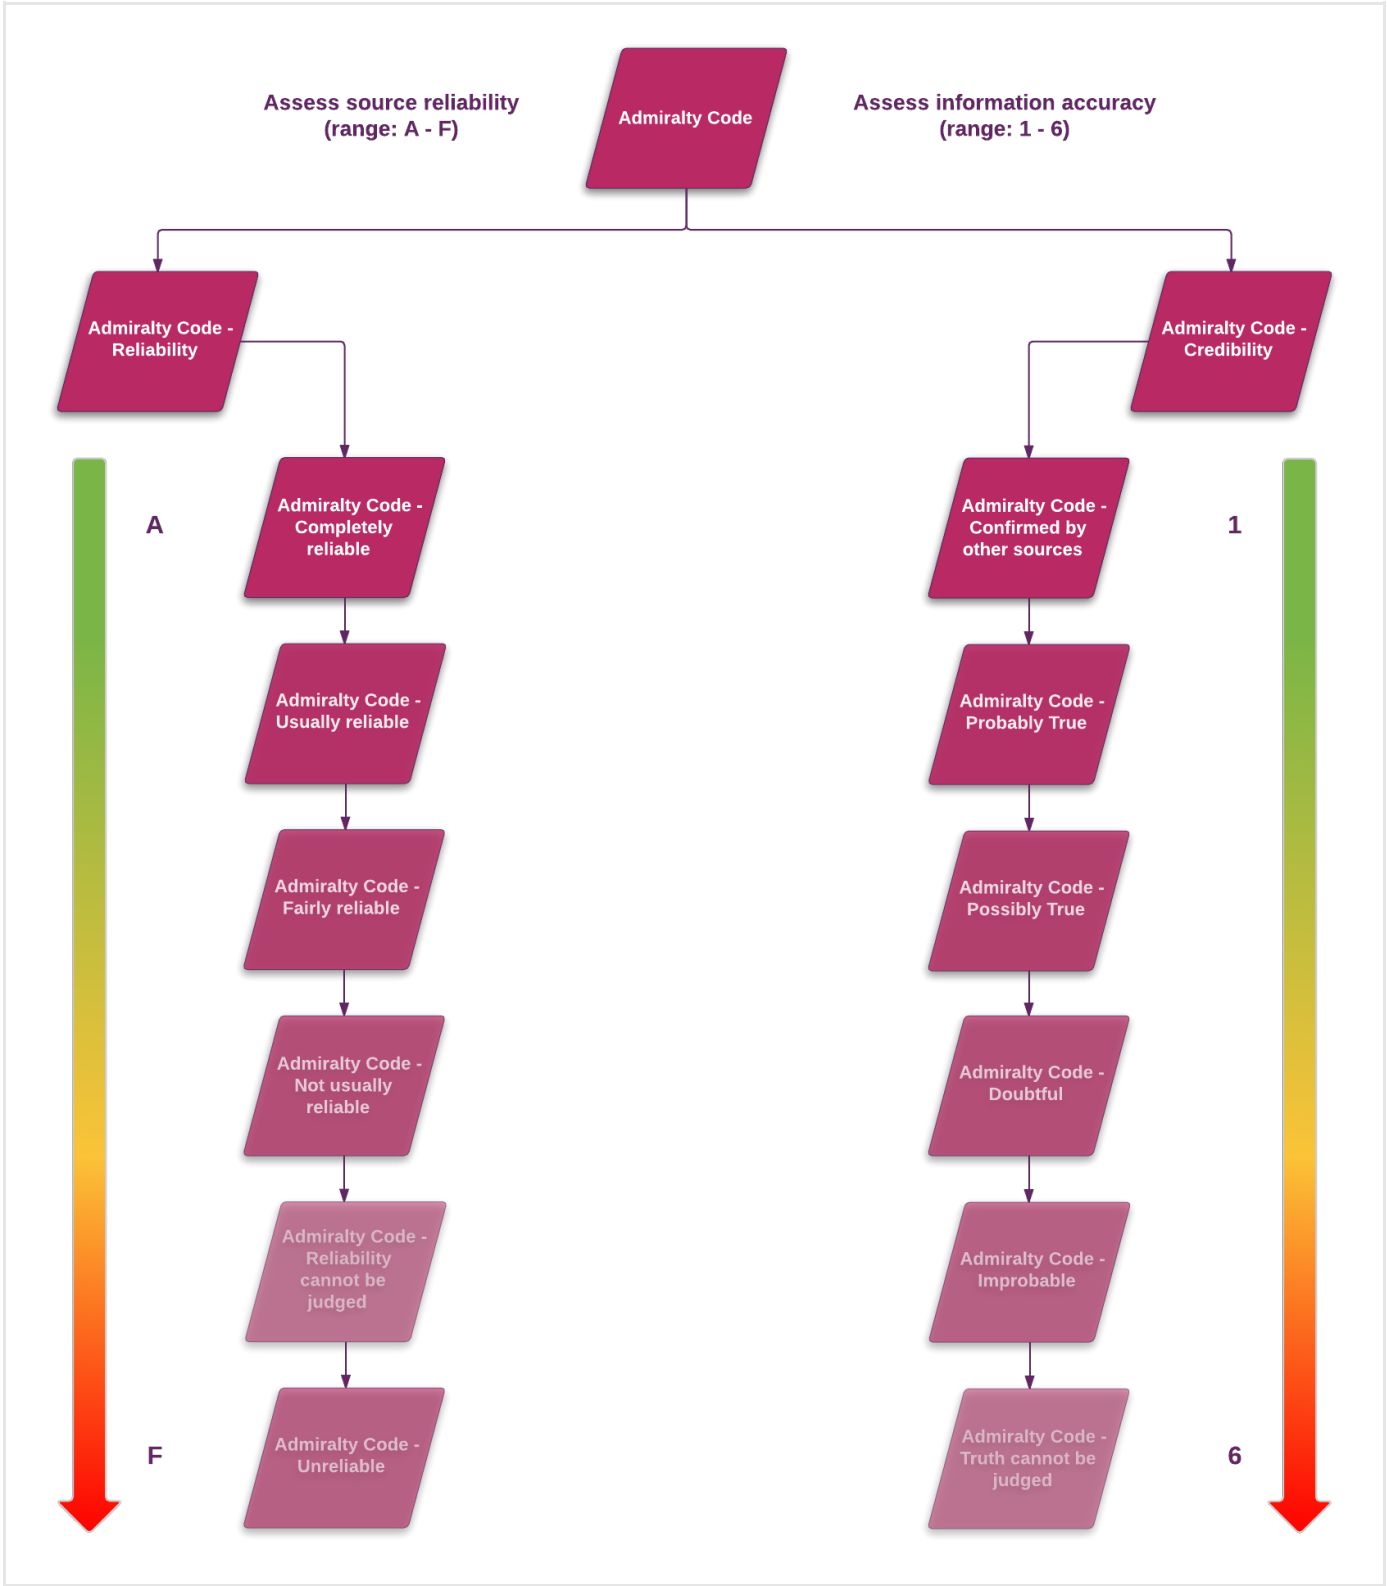

Kill chain

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.

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

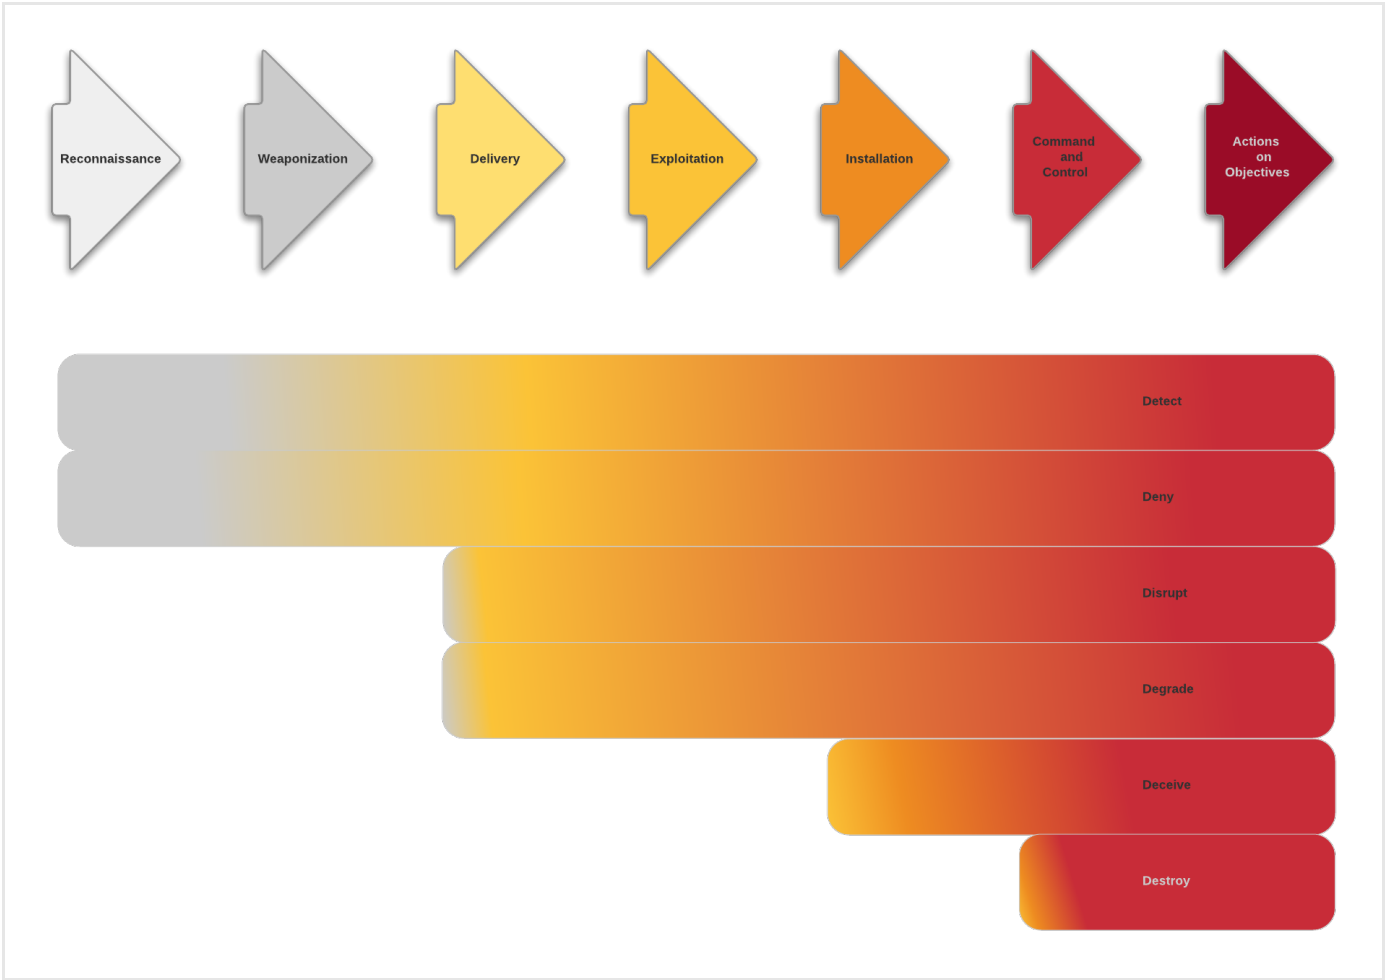

Create a taxonomy entry

✔ On the forms, 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 management > Taxonomy** .

Alternatively:

- On the top navigation bar click **⚙ > Data management > Taxonomies > ✚ Taxonomy**.

- On the **Data management > Taxonomy > Create** 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 beaneath.
- 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 management > Taxonomies** .
  The **Data management > Taxonomy** 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**.

- On the **Data management > Taxonomy > Edit** 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 management > Taxonomies** .
  The **Data management > Taxonomy** 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**.

- 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.

Rules

Rules give you granular control over entity observables by automatically flagging them as malicious, safe, or irrelevant. Entity rules enable you to automatically assign taxonomy tags, or to add entities to a dataset at the end of the ingestion process.

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.

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.

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 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.

Rule types

- Entity rules help you automate entity tagging and entity adding to one or more datasets, based on a predefined set of criteria.
- Observable rules act like filters: they filter entity observables as malicious, safe, or ignorable, based on a predefined set of criteria.

Entity rules

Add an entity rule

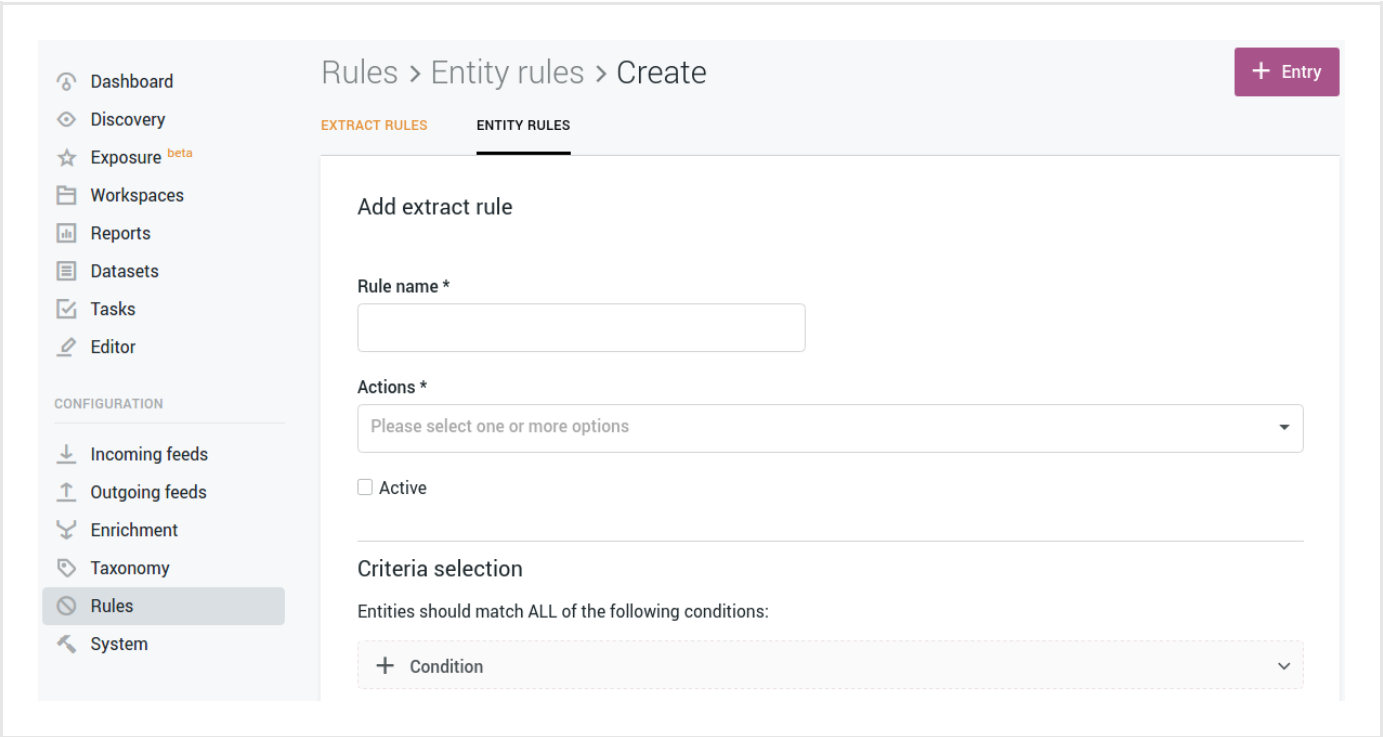

✔ On the forms, input fields marked with an asterisk are required.

To create a new entity rule, do the following:

- On the top navigation bar click **+ > 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 of the following 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:

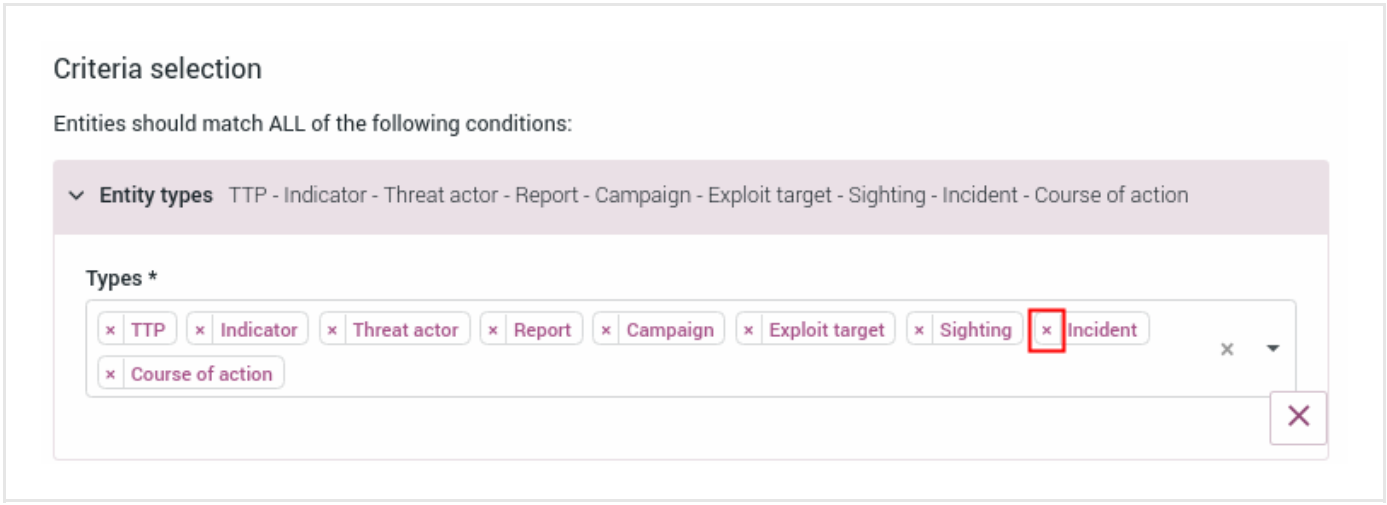

- **Content criteria**: key/value pairs define the content criteria the rule should apply.
  The input format for the *key* field is a *JSON* path. It points to an entity field/entity location in the entity structure.
  The input format for the *value* field is a *regex*. It specifies the content pattern.
  By default, **Content criteria** JSON path expressions are relative to the `data` field. `data` is the default root of any JSON path expression defined here.
  The `data` root is implied. To point to the title or to the description fields of an entity, you only need to enter `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": "robohelptesting.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.

  Value 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.

Observable rules

Add an observable rule

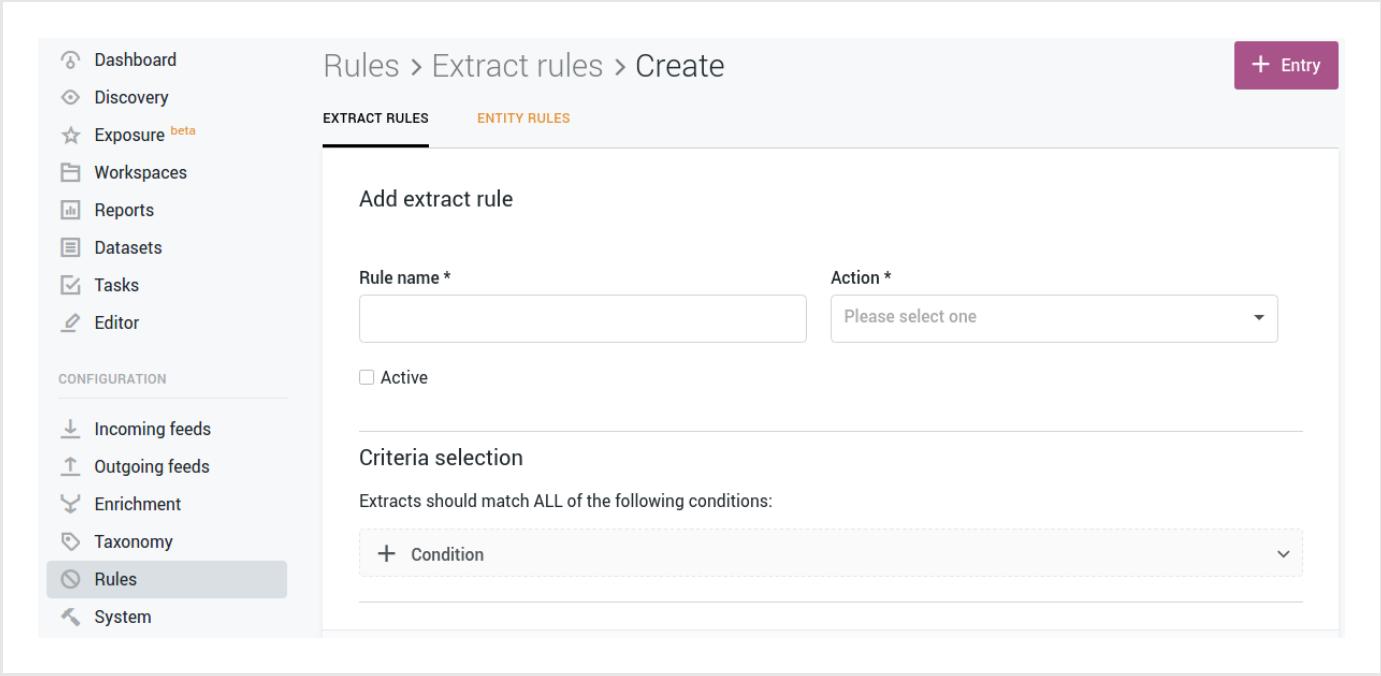

✔ On the forms, input fields marked with an asterisk are required.

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 of the following 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:

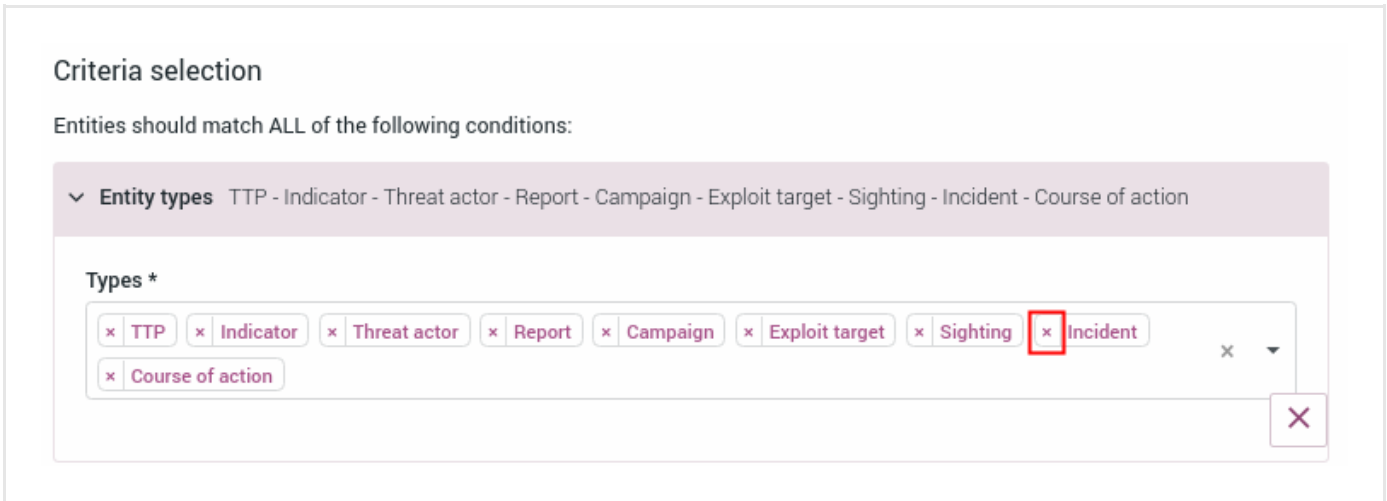

- **Observable types**: from the drop-down menu select one or more observable types, that is, the data types describing the corresponding entity observable data.
  For example, you may want to include in the rule processing a specific city name, an actor, a range of IP addresses, or a telephone number.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

  Value 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.

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 ⚙ **> 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 ⚙ **> 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**: *data.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.

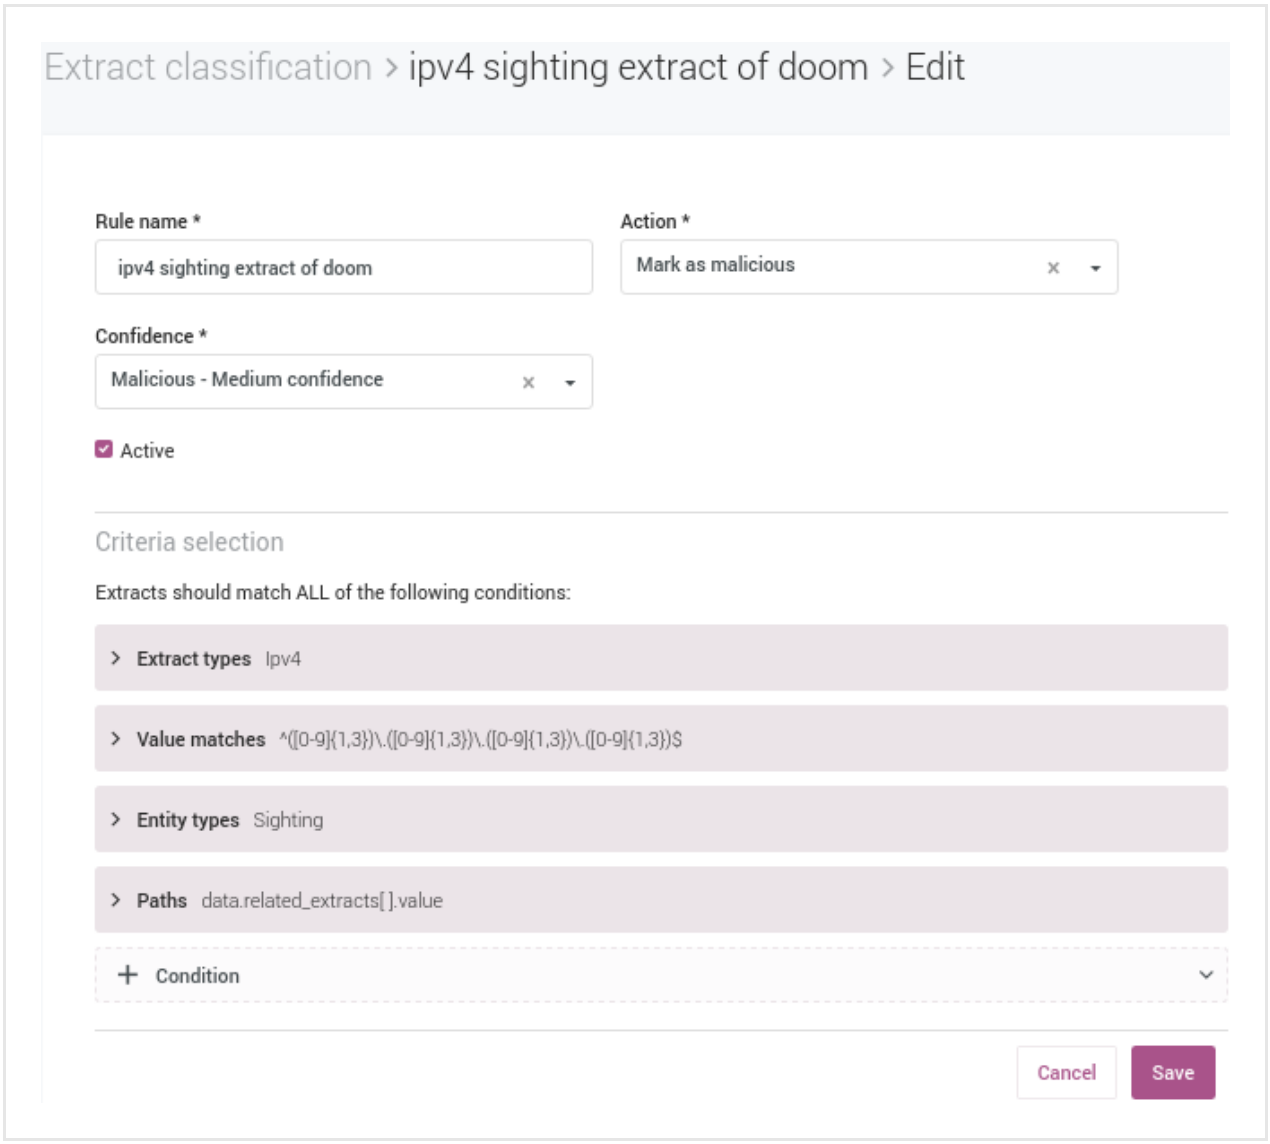

The path matching the specified pattern points to the ipv4 values in the second and third members of the following array:

```
{
  "data": {

    "related_extracts": [

      {
        "kind": "domain",
        "value": "robohelptesting.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%2Cdcd5701fc75f672e%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
```

```
    "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": []
}
```


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.

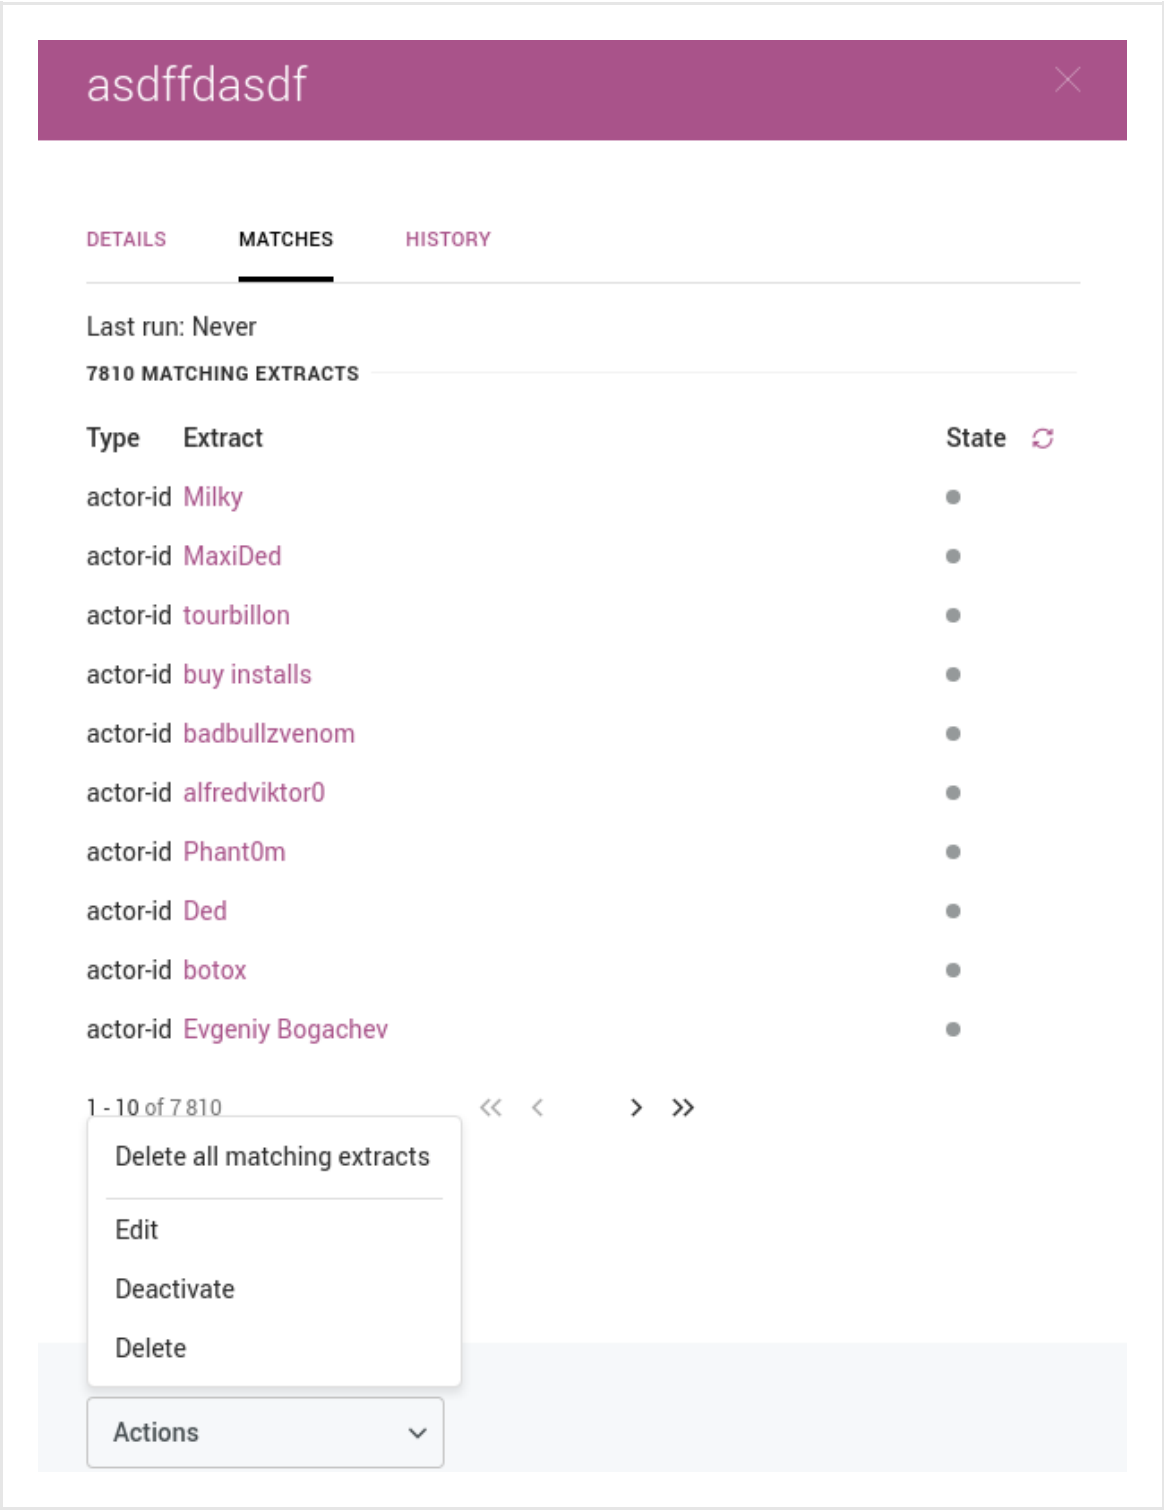

The **Matches** tab shows the matching observables the rule retrieved:

- **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.

Get observable types via API

You can retrieve a JSON response with all supported observable types by making an API call:

- Authenticate to obtain the token you pass with each API call.
- Send a request to the `/api/extracts/kinds/` (trailing slash included) API endpoint:

```
$ curl -X GET
       -v
       --insecure
       -i
       -H "Content-Type: application/json"
       -H "Accept: application/json"
       -H "Authorization: Bearer <token>"
       https://platform.host/api/extracts/kinds/

# 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/extracts/kinds/
```

- The response is a JSON array listing all supported extract types:

```
{
  "data": [

    {
      "kind": "uri"
```

```
},

{
  "kind": "product"
},

{
  "kind": "mac-48"
},

{
  "kind": "card-owner"
},

{
  "kind": "eui-64"
},

{
  "kind": "ipv4"
},

{
  "kind": "hash-sha256"
},

{
  "kind": "ipv6"
},

{
  "kind": "mnt-domains"
},

{
  "kind": "geo"
},

{
  "kind": "mutex"
},

{
  "kind": "hash-sha512"
},

{
  "kind": "fox-it-portal-uri"
},

{
  "kind": "email"
},

{
  "kind": "industry"
},

{
  "kind": "port"
},

{
  "kind": "email-subject"
},
```

```
{
  "kind": "company"
},

{
  "kind": "process"
},

{
  "kind": "telephone"
},

{
  "kind": "organization"
},

{
  "kind": "bank-account"
},

{
  "kind": "street"
},

{
  "kind": "nationality"
},

{
  "kind": "uri-hash-sha256"
},

{
  "kind": "handle"
},

{
  "kind": "hash-md5"
},

{
  "kind": "raw-artifact"
},

{
  "kind": "card"
},

{
  "kind": "person"
},

{
  "kind": "country"
},

{
  "kind": "descr"
},

{
  "kind": "name"
},

{
```

```
      "kind": "mnt-by"
    },

    {
      "kind": "netname"
    },

    {
      "kind": "hash-sha1"
    },

    {
      "kind": "postcode"
    },

    {
      "kind": "actor-id"
    },

    {
      "kind": "malware"
    },

    {
      "kind": "city"
    },

    {
      "kind": "host"
    },

    {
      "kind": "winregistry"
    },

    {
      "kind": "file"
    },

    {
      "kind": "domain"
    },

    {
      "kind": "asn"
    },

    {
      "kind": "mnt-routes"
    },

    {
      "kind": "cve"
    },

    {
      "kind": "inetnum"
    }

  ]
}
```

Observable types

The available observable types are:

actor-id
asn
bank-account
card
card-owner
company
cve
domain
email
email-subject
eui-64
file
forum-name
forum-thread
forum-room
fox-it-portal-uri
geo
geo-lat
geo-long
city
country
country-code
address
street
postcode
hash-md5
hash-sha1

hash-sha256
hash-sha512
handle
host
industry
inetnum
ipv4
ipv6
mac-48
malware
mutex
name
nationality
netname
organization
person
port
process
product
raw-artifact
registrar
telephone
uri
uri-hash-sha256
winregistry

Dashboard

The dashboard is the main entry point to the platform. 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.

Customize the dashboard

You can customize the dashboard by adding and removing gauges: click the **Customize** button at the top of the page to view a list of the available gauges.

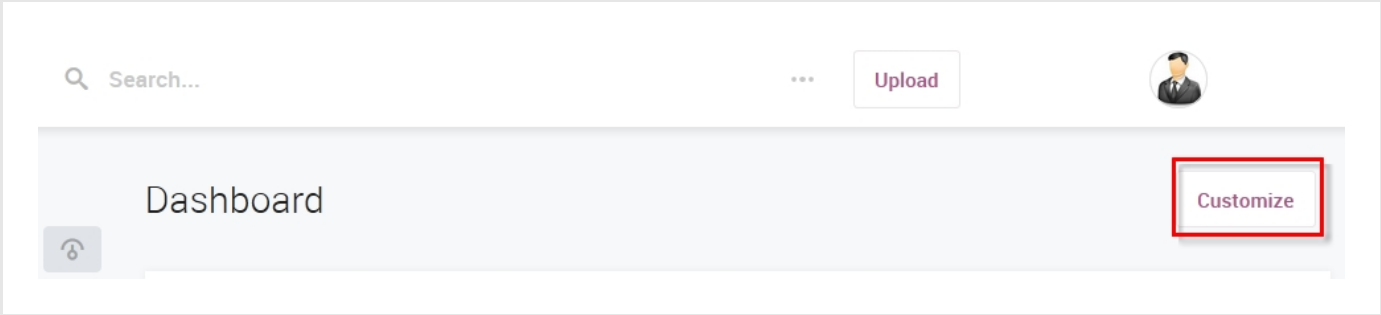

- To add a gauge from the list to the dashboard view, select the corresponding checkbox.
- To remove a gauge from the dashboard view, deselect the corresponding checkbox.
- When you are done, click **Return to the dashboard** to apply the changes and go back to the dashboard.

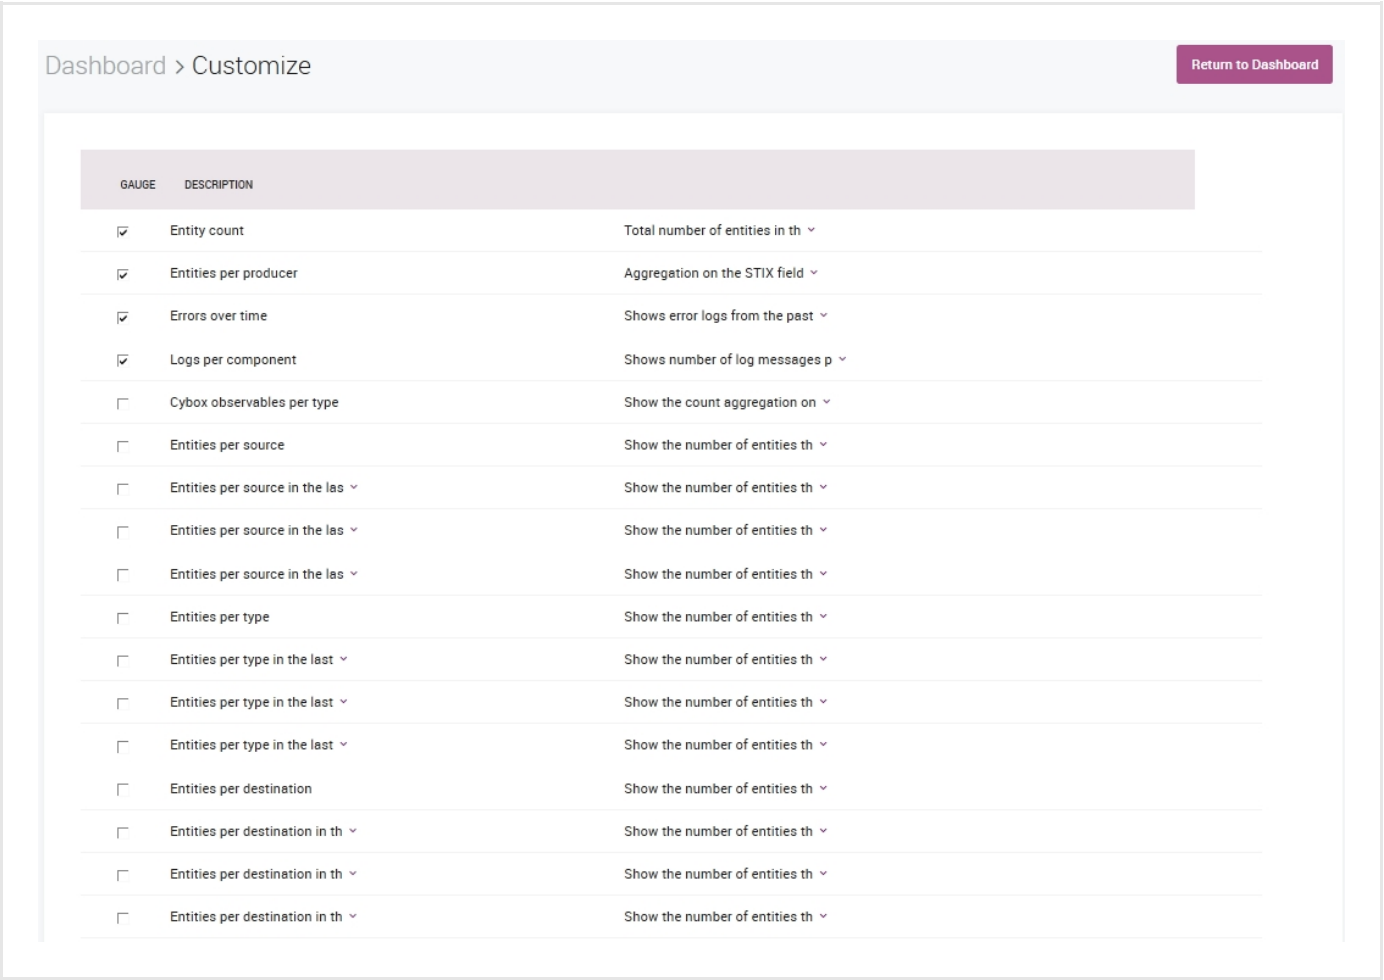

Warning: Try to limit the number of active gauges on the dashboard. Each gauge polls data from the database. Therefore, a dashboard with many gauges may become resource-intensive, and it may take longer to load.

Search the platform

Use the search field to look for entities and indicators in the platform.

Search

You can find the search box on the top bar:

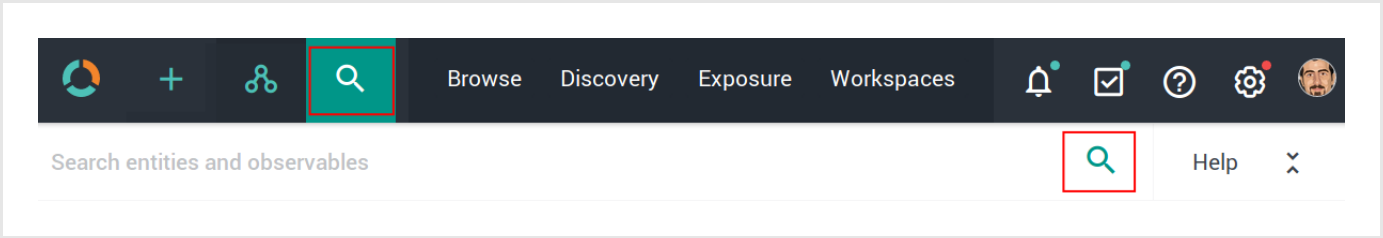

Enter search terms and search queries, and then press **ENTER** or click the search icon to run the search. Searches you run through this search box are executed platform-wide.

> 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.

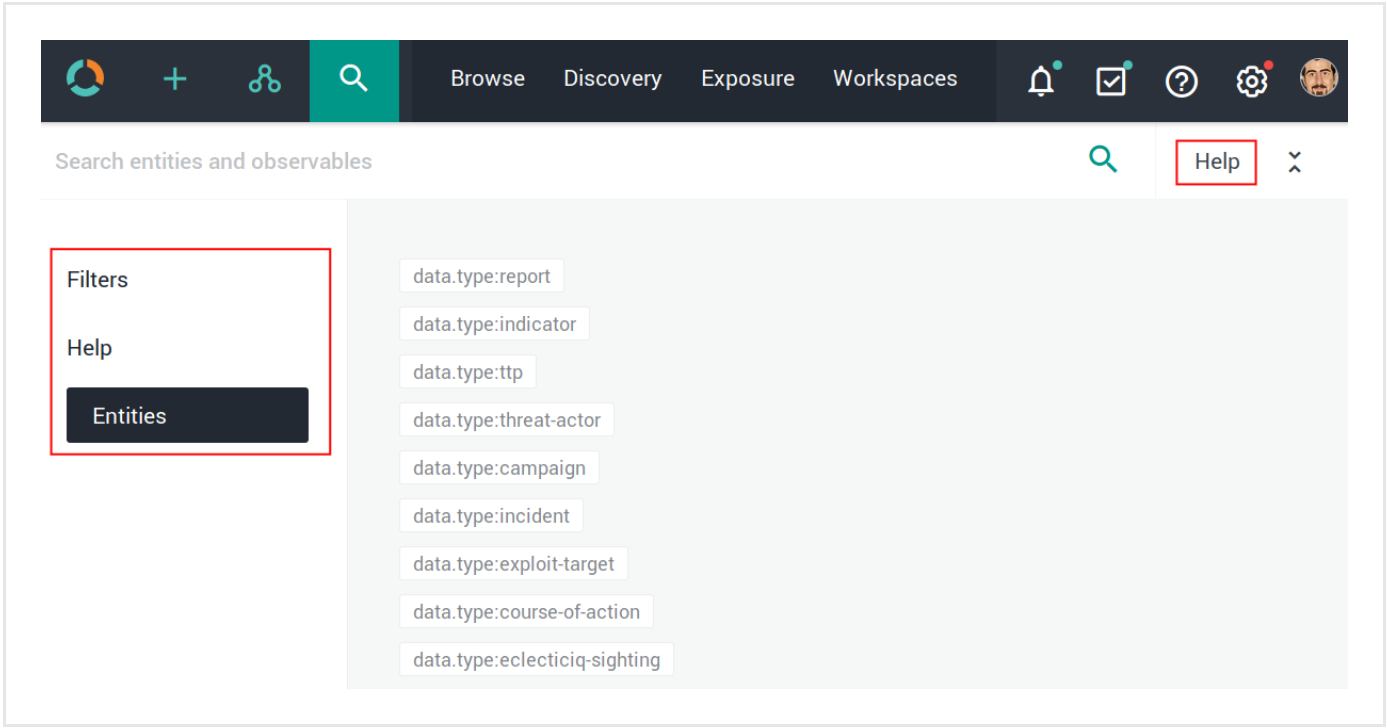

Besides full text search, you can use Boolean operators, wildcards, regex, and you can combine these filtering options to create more refined searches.

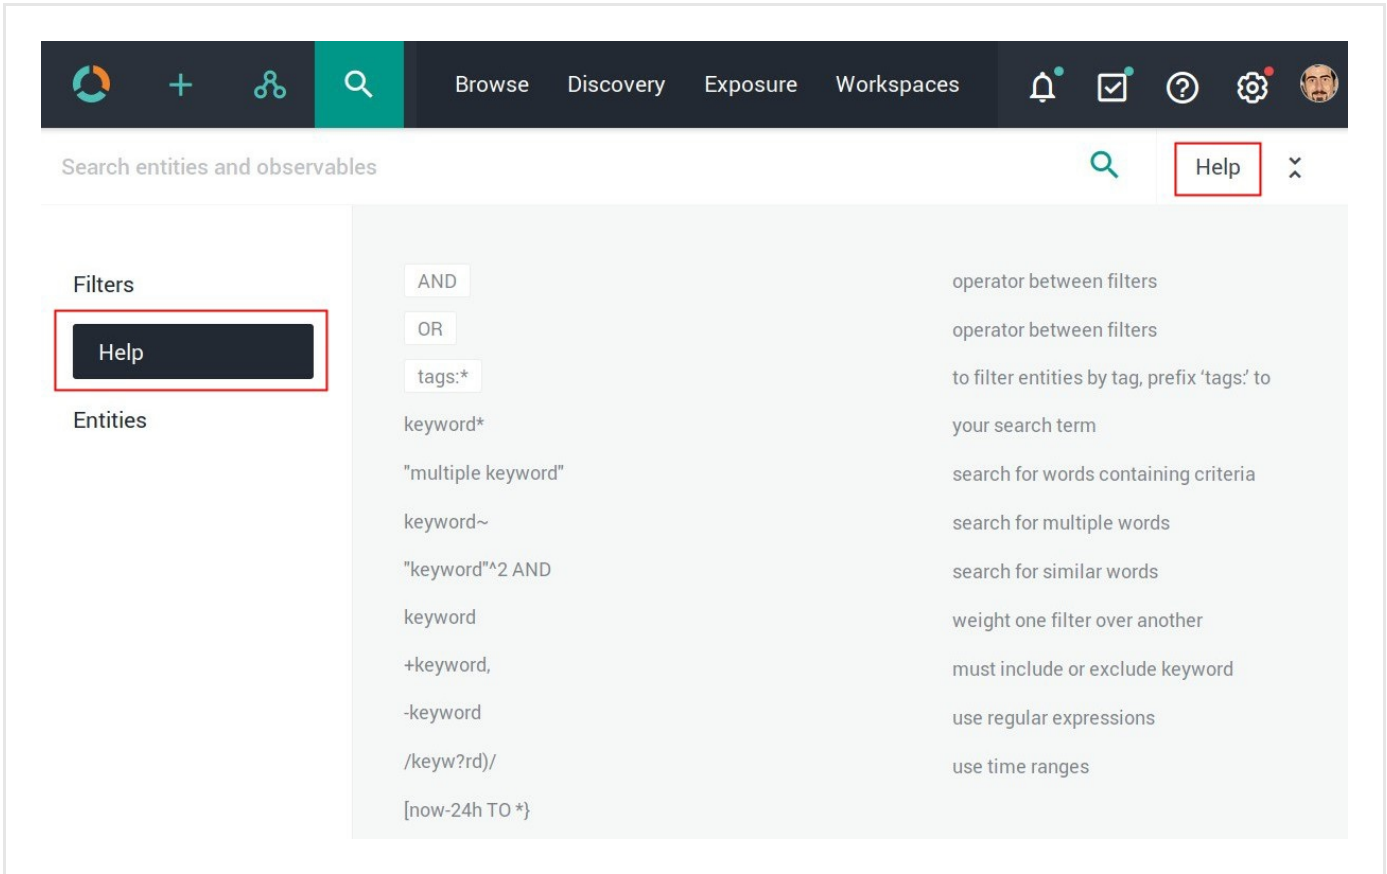

Search query fields

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

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

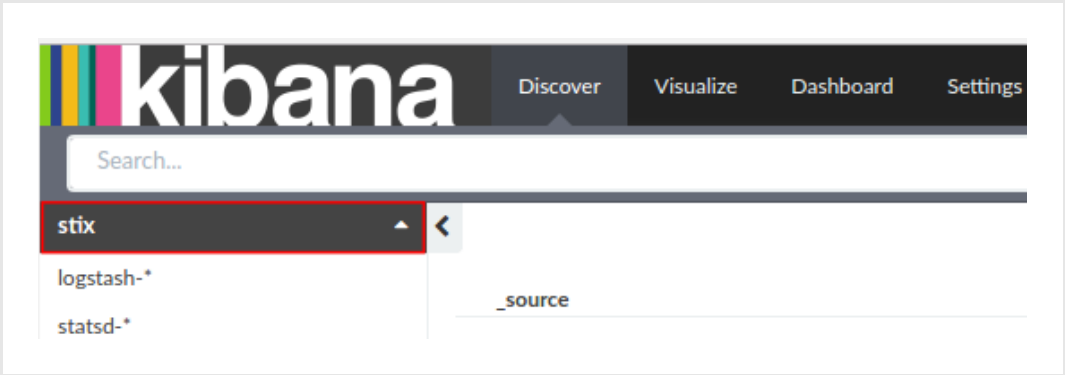

- On the main menu bar, select **Settings**:

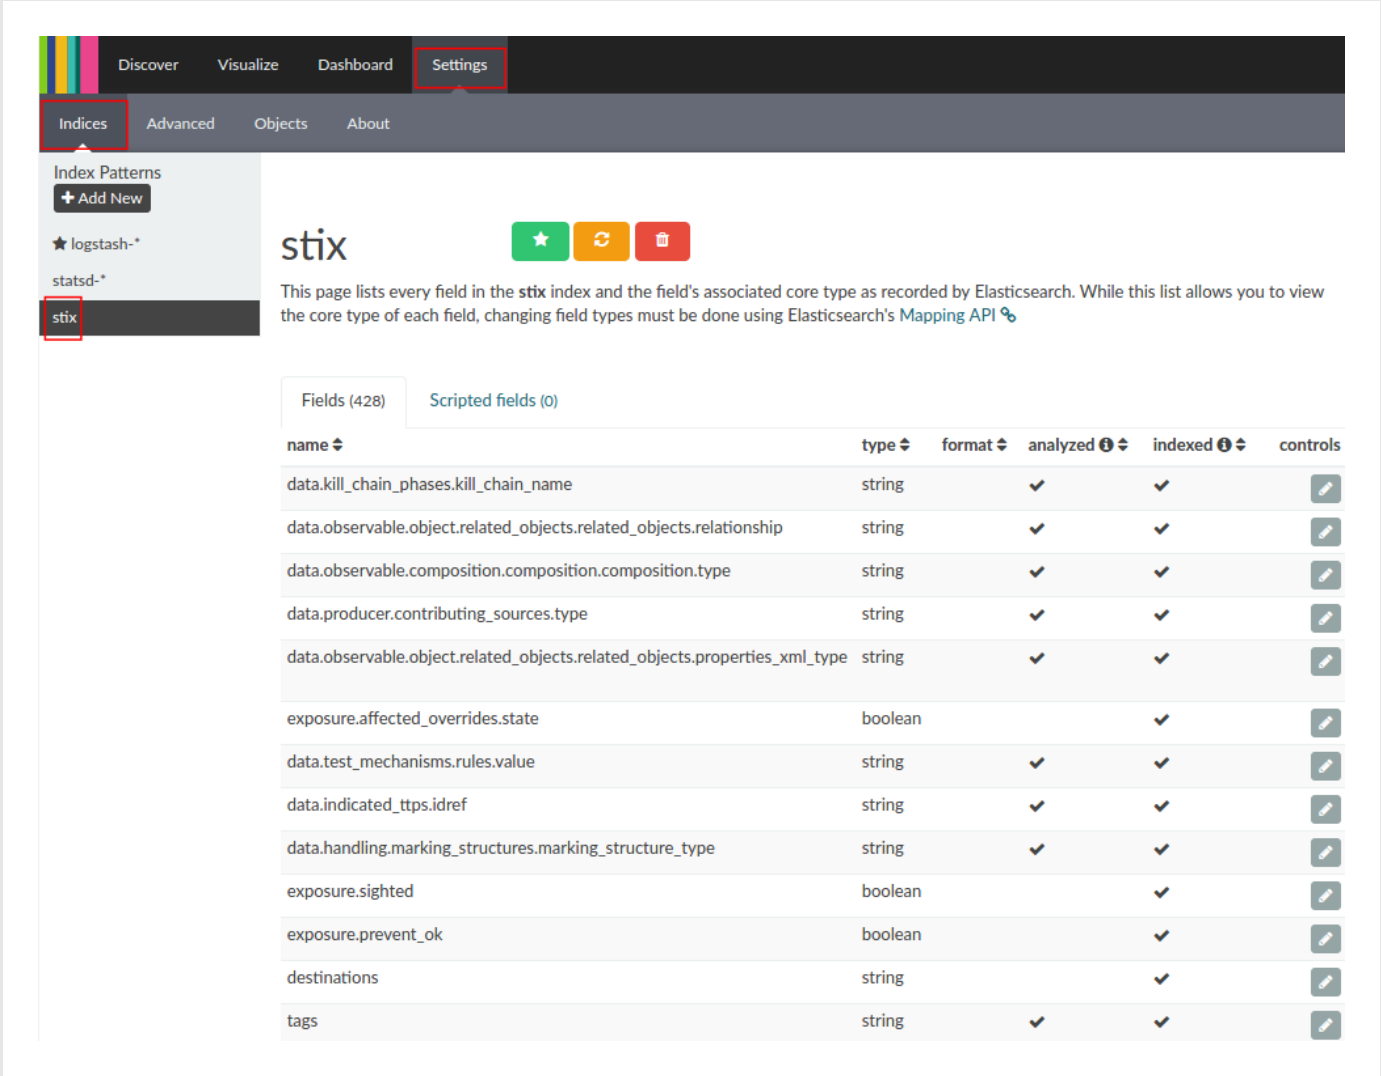

Search timeout

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, and it needs to be an integer.
- Save the configuration file.

Upload data files

Use the upload option to add data files and compressed archives to the platform.

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

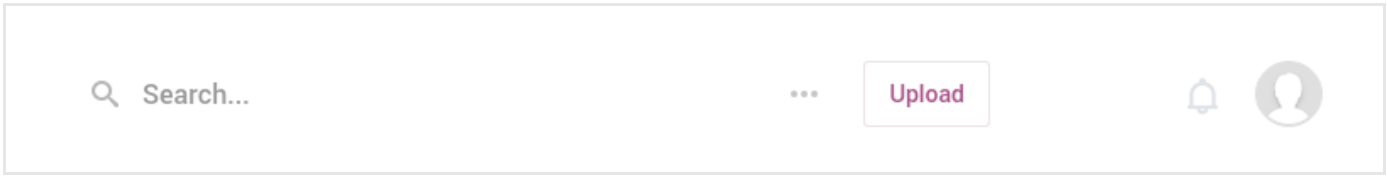

Upload data

- **Content type**: from the drop-down menu select a content type corresponding to the file format you are about to upload. The available options on the list correspond to the allowed content type formats that the platform can ingest and process.
- **Extraction ignore levels**: from the drop-down menu select at least one option if you want to *exclude extracted content* from the ingested data.
  If you select at least one value, the filter excludes extracted data from ingestion, based on the direct or indirect relationship the data has with the entity it refers to.
  In other words, the filter ignores specific data, based on the data location in the entity data structure:
  - **Extraction ignore levels — 1**: the extracted data is inside a CybOX object. The ingestion process ignores and it does not include extracted data found inside observable CybOX objects embedded in STIX indicators.
  - **Extraction ignore levels — 2**: the extracted data is outside a CybOX object. The ingestion process ignores and it does not include extracted data found inside STIX fields. For example, STIX headers, titles or references.
- **Archive**: select this checkbox if you are uploading compressed *zip* archives.
  If the archive(s) you are manually uploading are password-protected, select the **Password protected archive** checkbox, and then enter the password in the **Archive password** field. The specified password acts as a master password, and it is used to unlock all the archives included in the same upload operation.
- Drag and drop files onto the upload area, or click it to open a system file manager window, and browse to the desired file location.
- After selecting the appropriate file type and the file(s) you want to upload to the platform, click **Upload** to complete the action.
- Under **Current uploads** you can see the upload queue.
  File upload progress for each file is expressed as a percentage.
  A confirmation message notifies a successful file submission.
- To submit a new file for upload, click **New upload**.

About archives

- The archive you want to upload should be in *zip* format.
- When you prepare an archive for upload to the platform, you should not mix file types: to be correctly processed, all the files inclided in the archive need to share the same content type.
- You can upload report documents in plain text *(txt)*, STIX or PDF format by including them in an archive, which you subsequently upload to the platform. Make sure all reports in the zipped archive share the same content type.
- The platform automatically extracts and produces entities from successfully uploaded archive files.
- The maximum file size you can upload is 25 MB.

Discovery

Use Discovery to run automatic searches returning specific cyber threat information.

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

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 to 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.

> ✔ On the forms, 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 to 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.

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.

Run rules manually

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

Workspaces help you organize your threat analysis tasks and manage collaboration efforts with other colleagues. You can create private and public workspaces.

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 easily 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 adress 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 to 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.

If there are no workspaces yet, you can easily configure them.

Create a workspace

To create a workspace, do the following:

- On the top navigation bar click **✚ > Workflow > Workspace**.

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.

> ✔ On the forms, 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.

Workspace types

Workspaces provide user-friendly thematic environments to help you efficiently organize your tasks and your data.

The **Workspaces** page displays all available workspaces as tiles.

Workspace tiles

Have a quick look at a workspace tile to quickly get high-level information about it. For example when it was last edited, by whom, how many collaborators are participating in the workspace, and how many entities it holds.

Hover the mouse cursor on the top half of the tile to display a free-text description, if available, that provides further information about the workspace.

When a tile shows a closed lock icon, it means that the corresponding workspace is private, and therefore only its members can access it.

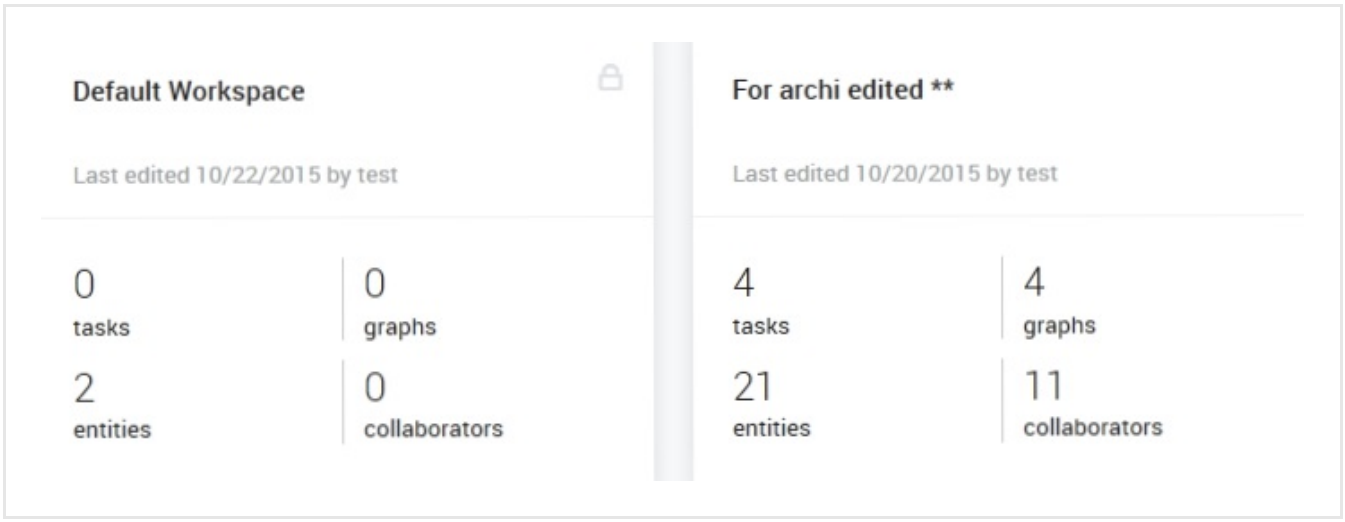

A locked, private workspace and a public one, respectively

Manage workspaces

On this page you can carry out basic workspace management operations:

- Sort the workspace display order:
  - Either alphabetically;
  - Or in reverse chronological order, based on the modification date (**Last Modified**)

- Filter workspaces to view only specific ones:
  - Click **Show**.
  - From the drop-down menu select one or more checkboxes to display only the workspace types you want to view:

Workspace type	Descritpion
Archived	An archived workspace is no longer updated, but its content can be useful for reference.
Generic	A generic repository, it can be the initial step towards a more focused workspace with limited scope.
Team	Helps you organize tasks and information shared across collaborators belonging to the same team.
Topic	Helps you organize actions and information concerning a specific topic, for example Chinese malware, or cyber threats affecting the banking sector.
Case	Helps you organize actions and information concerning a case, for example an ad-hoc cyber-attack that took place on a specific date to hit a specific target.

- Add a new workspace.

Workspace Overview tab

The Overview tab sums up relevant workspace details.

The workspace **Overview** tab gives you detailed information about the selected workspace content and its purpose.

For example you can read a short description for the workspace, if available, view thumbnails of any saved graphs, check if there are any scheduled running or pending tasks, view any file attachments, and examine any entities belonging to the workspace.

Add and remove collaborators

You can also see how many collaborators participate in the workspace, and an email address to contact the group, if provided.

If you belong to the workspace, and if you have the appropriate user rights, you can add new collaborators by clicking the pencil-shaped icon on the top-right corner of the tab. In the popup dialog window you can see a list with the current collaborators.

To add a new collaborator:

- In the popup dialog window, click **Add Collaborator**.
- From the **User** drop-down list, select the name of the person you want to add to the workspace.
- Click **Save Collaborator**.

To remove a collaborator:

- In the popup dialog window, click the delete icon on the right, corresponding to the collaborator you want to remove from the workspace.

Workspace Tasks tab

The Tasks tab displays the ongoing activities in the workspace, and the collaborators who are assigned to carry them out.

The **Tasks** tab is the *what's going on agenda* for the workspace: here you can see who is doing what within the workspace.

The workspace task overview tab shows you task information in table format:

Column	Description
Assignee	The owner of the task.
Status	The task status in the task flow.
Title	A short title to identify the task. It should be descriptive and easy to remember.
ID	Num., integer. An automatically generated task reference identifier.
Due date	Deadline: the scheduled completion date for the task.

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 to 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 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.

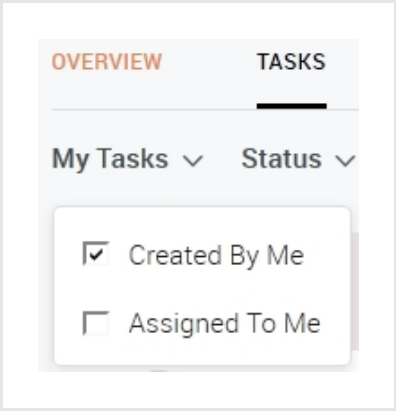

View tasks by status

- 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.

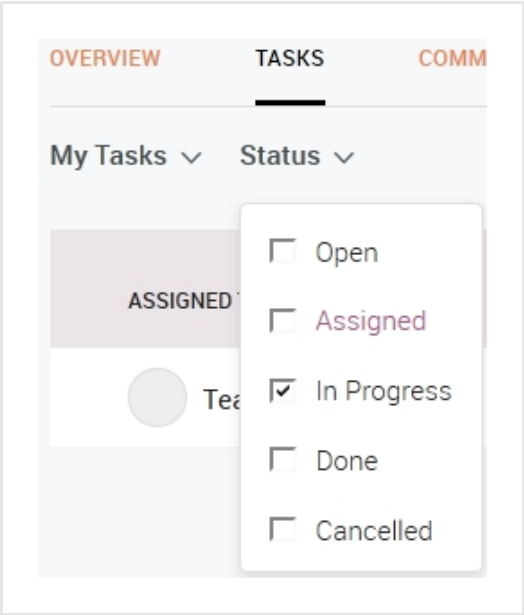

View task details

- On the 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**.

View 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.

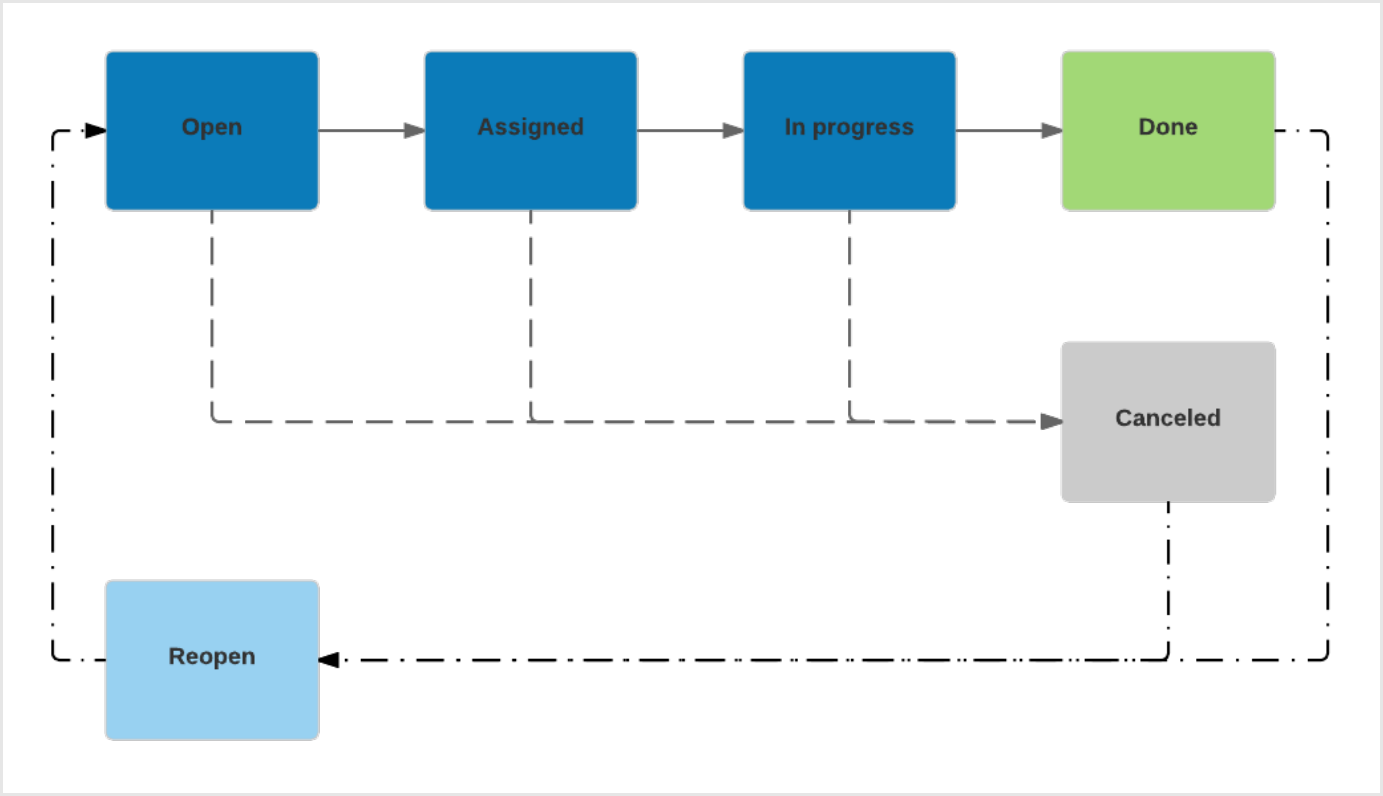

A standard task status flow

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.

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.

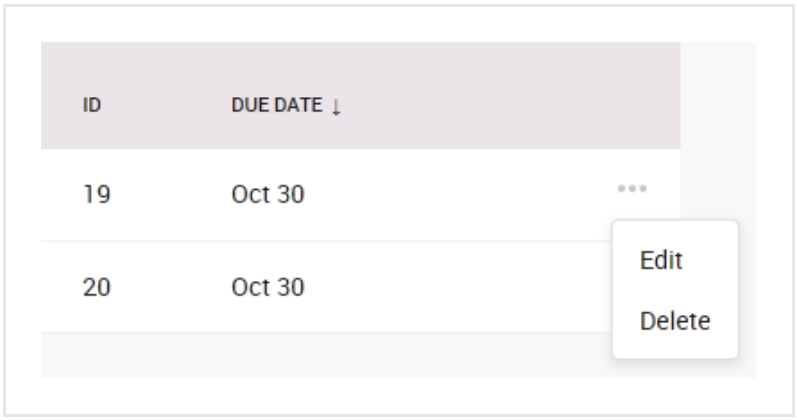

Grayed-out options in the menu are disabled for the selected item.

You can access the same action options by clicking the **Actions** the pop-up menu on the task detail pane.

Workspace Comments tab

The Comments tab displays shared comments and information the collaborators exchange in the workspace.

In the **Comments** tab you can review any comments the workspace collaborators added to provide additional context or to explain unclear items.

Besides providing valuable information, the comment thread acts as a history of the workspace, since it is here that collaborative exchanges, questions and answers, tips, and so on are recorded for reference.

Workspace Saved graphs tab

On the Saved Graphs tab you can find stored graph visualizations for reference or further analysis.

The **Saved Graphs** tab gives you access to any saved graphs the workspace collaborators are keeping for further investigation, analysis or to look them up at a later moment.

The workspace **Overview** tab offers a quick access to a selection of saved graphs. The **Saved Graphs** tab displays all saved graphs for the workspace.

To load or delete a graph, click the ⋮ icon on the relevant graph tile:

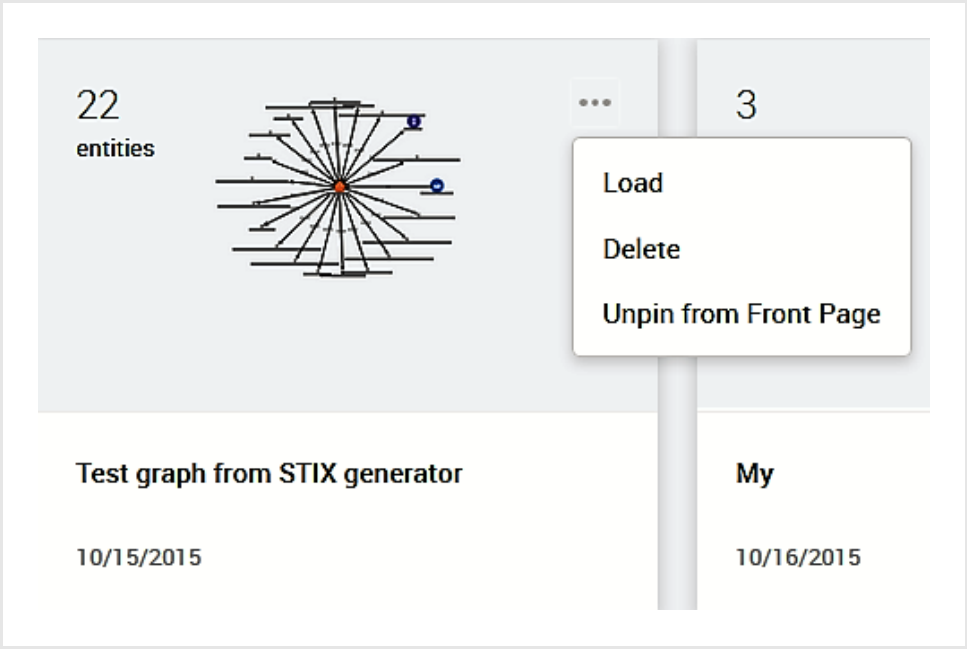

- **Load** allows you to load the selected graph to analyze entity relationships and references in an intuitively visual way.
- **Delete** purges and discards the graph and its content.
- **Unpin from Front Page** removes the graph quick access reference from the **Overview** tab, but keeps the graph available on the **Saved Graphs** tab.

To use graphs, you first need to populate them with entities and/or datasets.

Workspace Entities tab

The Entities tab displays all the entities in the current workspace.

> In the platform context, an entity is an **IOC** (https://en.wikipedia.org/wiki/indicator_of_compromise).

This area allows you to access, examine and review all the entities included in the workspace.

The **Entities** tab is related to the **Saved Graphs** one: use the latter to choose and display visual representations whose data you define and select in the **Entities** tab.

It shows an overview where each entity is assigned a row. You can search entities using the search bar on the upper half of the page; moreover, you can filter and manipulate entities by executing actions and by editing them.

Actions

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 onto 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.

Alternatively, in the workspace **Entity** tab or in the entity editor you can click the ⋮ icon on the row corresponding to the entity you want to modify, and then choose the desired action from the context menu. Grayed-out options in the menu are disabled for the selected item.

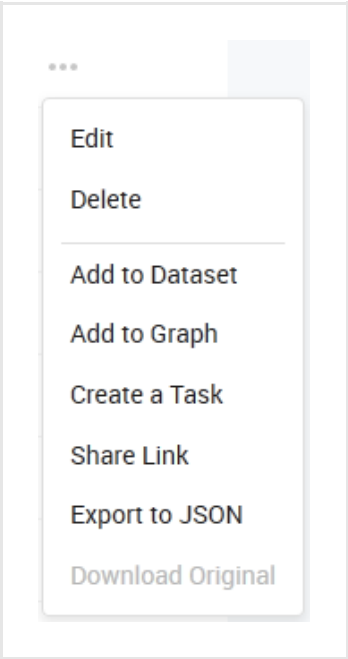

View entity details

To view detailed information about an entity, do the following:

- Go to the entity editor published entity overview.
- Click anywhere on the row corresponding to the entity you want to inspect.
- An overlay slides in from the side of the screen.

The entity name is always visible at the top of the entity detail pane. Click the edit icon next to the entity name to edit it. Below the entity name, you can see summary metadata details:

- Ingestion time
- Entity author user name
- User group the entity author belongs to
- **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).

The TLP value is displayed on the bottom-right corner of the header section. Click it to change its value, if necessary.

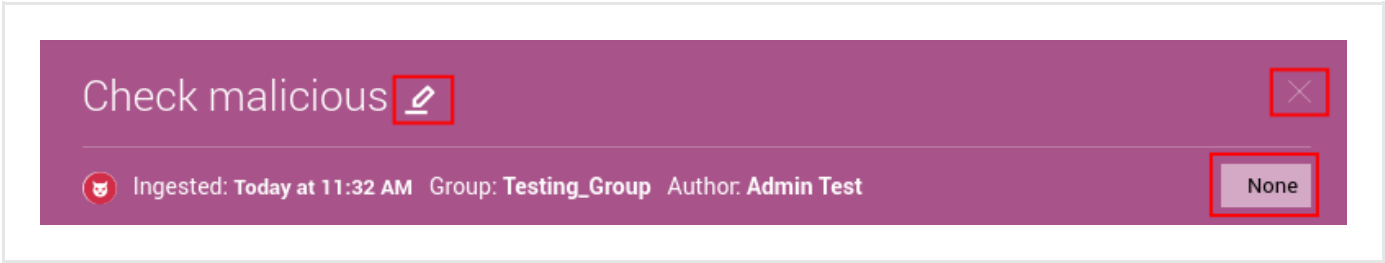

Create a new entity

To create a new entity, in the **Entities** tab select the **Draft entities** sub-tab. **Draft entities** retains the same look and feel as its parent tab, and it allows you to create new entities that you can add to the workspace.

- Click **Create New Entity**.
- From the drop-down menu select the type of entity you want to create.

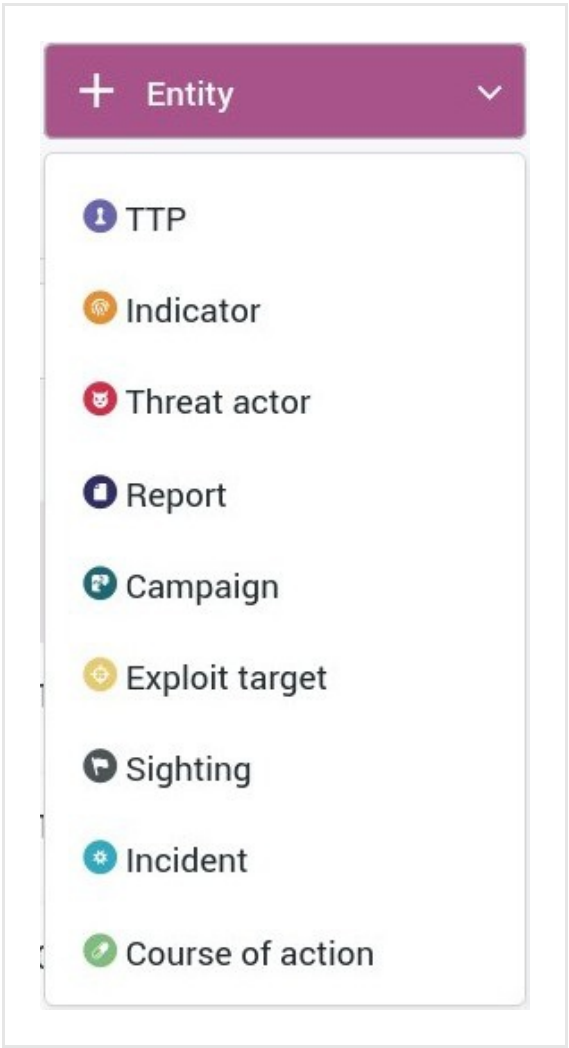

The drop-down menu discloses the available entity types you can create

Entity type	Description
Campaign (https://stixproject.github.io/data-model/1.2/campaign/campaigntype/)	A campaign is a series of planned actions aiming at achieving a specific goal. It groups a set of related threat actors, TTPs, and incidents sharing a common intent or goal.
Course of action (https://stixproject.github.io/data-model/1.2/coa/courseofactiontype/)	A course of action details a set of clear, specific recommendations and measures to mitigate an incident, address affected exploit targets, and effectively respond to a cyber threat.
Exploit target (https://stixproject.github.io/data-model/1.2/et/exploittargettype/)	An exploit target is a vulnerability or a weakness in software, hardware, systems, or networks that a threat actor can leverage and take advantage of to intrude or carry out an attack.
Incident (https://stixproject.github.io/data-model/1.2/incident/incidenttype/)	An incident describes a specific occurrence of one or more indicators affecting an organization. It includes information on threat actors, tools or skills, timeframes, techniques, as well as impact assessment and the recommended response course of action.
Indicator (https://stixproject.github.io/data-model/1.2/indicator/indicatortype/)	An occurrence or a sign that an incident may have occurred or may be in progress. See also the definition provided in the **Cybersecurity Information Sharing Act of 2015 (CISA)** (https://www.congress.gov/bill/114th-congress/senate-bill/754/text).
Report (https://stixproject.github.io/data-model/1.2/report/reporttype/)	A detailed account of an indicator of compromise (IOC), a threat, a campaign or other threat activity as a result of an investigation or an analysis. A report tells a story about a piece of threat intelligence by providing background, context, and by pulling threads together to weave a clear and meaningful description of a security breach, a cyber attack, or a series of attacks.
Sighting ()	A sighting records a specific observation of a malicious indicator by matching fingerprints. For example, it can record the occurrence of a malicious IP address at a specific date and time,
Threat actor (https://stixproject.github.io/data-model/1.2/ta/threatactortype/)	An individual or a group carrying out or planning to execute malicious activities. Threat actors include information on their identity, suspected motivation, and suspected intended effect.
TTP (https://stixproject.github.io/data-model/1.2/ttp/ttptype/)	Tactics, Techniques and Procedures. Sometimes referred to also as Tools, Techniques, Procedures. TTPs describe the behavior of cyber adversaries. Tactics describe *"the employment and ordered arrangement of forces in relation to each other"*. Techniques are *"non-prescriptive ways or methods used to perform missions, functions, or tasks."* Procedures are *"standard, detailed steps that prescribe how to perform specific tasks."* (definitions from *"Joint Publication 1-02, Department of Defense Dictionary of Military and Associated Terms, 8 November 2010 (as amended through 15 February 2016)"*)
Package (https://stixproject.github.io/data-model/1.2/stix/stixtype/)	A package is a wrapper containing one or more STIX objects such as indicators, threat actors, TTPs, and so on. When the platform ingests packages, it extracts the STIX objects and it converts them to its internal JSON data model.

When you select the desired entity type you want to create, the entity editor is displayed. In the editor you can create STIX-compliant entities in a fillable form page.

For further information on building and structuring entities, see the **STIX data model** (http://stixproject.github.io/data-model/) and the recommendations about using a **controlled vocabulary** (http://stixproject.github.io/documentation/concepts/controlled-vocabularies/).

After filling out the necessary input fields to record the new entity, you can save it as a draft or save it and publish it immediately:

- **Save draft**: Click **Save draft** to store your changes, or **Cancel** to discard them.
  The new entry is saved as a draft, but it is not published, i.e. it is inactive. It is available in the entity editor under **Draft entities**.
  If you are creating a new entity and you have not yet saved it for the first time, you can also click the drop-down arrow and select **Save draft and new** to:
  - Save the current populated form as a draft without publishing it to the platform;
  - Create and open a new draft form in the editor.
- **Publish**: Click **Publish** to store your changes, or **Cancel** to discard them.
  The new entry is saved and published to the platform. As soon as it is indexed, it is available in the platform in the entity editor under **Published**.
  Published entities associated to a workspace or included in a dataset are available also through the corresponding workspace and dataset.
  If you are creating a new entity and you have not yet saved it for the first time, you can also click the drop-down arrow and select **Publish and new** to:
  - Save the current populated form and publish it to the platform;
  - Create and open a new form in the editor.

Filter entities

You can filter entities to zero in on specific subsets:

- Use the drop-down menus above the entity list to select the criteria you want to apply to filter the entities.
- You can make multiple selections per menu and across the menus to further refine your results.
- The available menu options may vary, since they are based on the metadata of the entities in the workspace.

The available filter menus are:

Filter menu	Description	

Filter menu	Description	
Entity Types	Filter entities by type.	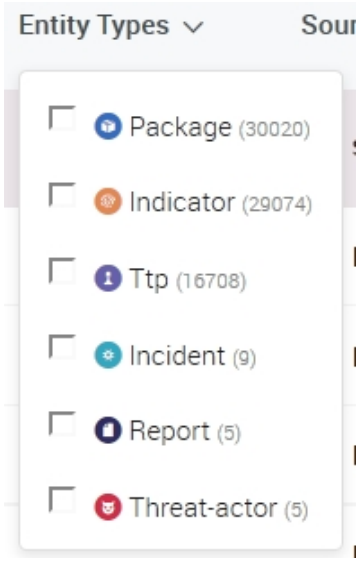
Source Types	Filter entities by source/origin.	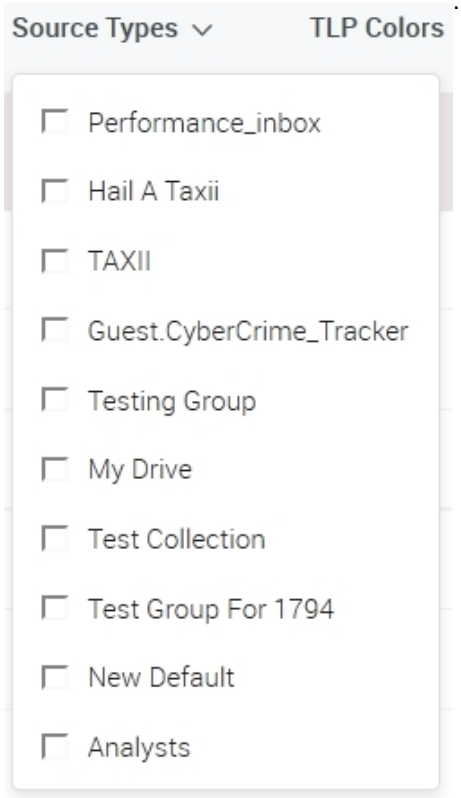
TLP Colors	Filter entities by Traffic Light Protocol color code.	

Filter menu	Description	
Date	Filter entities included in a date range.	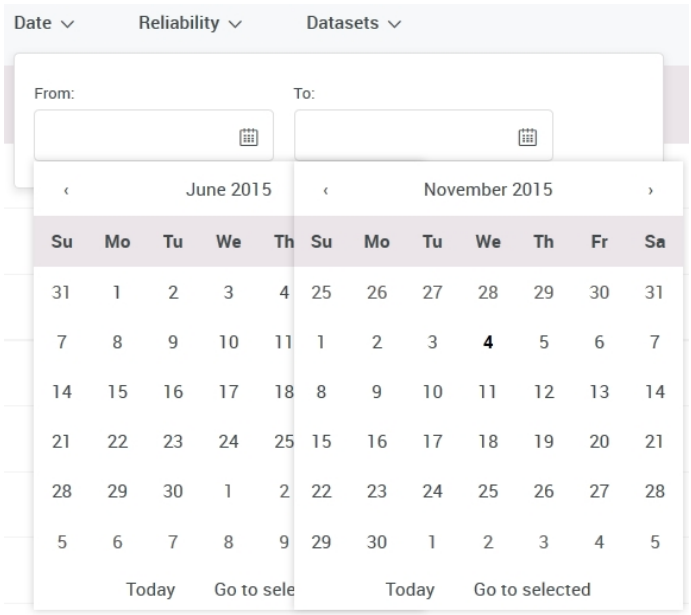
Reliability	Filter entities based on their reliability index.	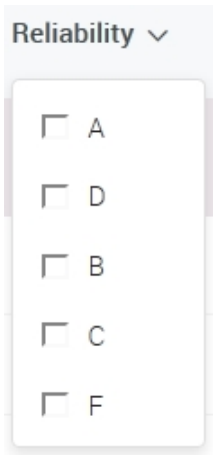
Datasets	Filter entities based on the dataset they belong to.	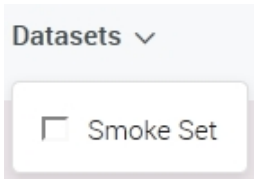

Entity reliability

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)`.
The allowed values for this field are:

Reliability value	Source accuracy
A	Completely reliable
B	Usually reliable
C	Fairly reliable
D	Not usually reliable
E	Unreliable
F	Reliability cannot be judged

Workspace Files tab

Use the Files tab to upload files to the workspace.

The **Files** tab is your upload point for the workspace. Use it to manually upload files as attachments to the current workspace. For example, besides entities you may want to access other reference and background information from within a workspace, so that it is readily available when needed.

The files you upload through the **Files** are not processed, they are simply stored. Besides the content types the platform can ingest, you can upload other formats like image files or presentations.

Upload a file

To upload a file through the **Files** tab:

- Browse to the location where the files you want to upload are located.
- Select one or more files, then drag and drop them onto the upload area, which is marked with an upload icon:

- Alternatively, click the upload icon to open your machine file manager.
- In the file manager, select the files you want to upload to the workspace, and then confirm your selection to populate the upload area in the **Files** tab.
- The file upload starts automatically.

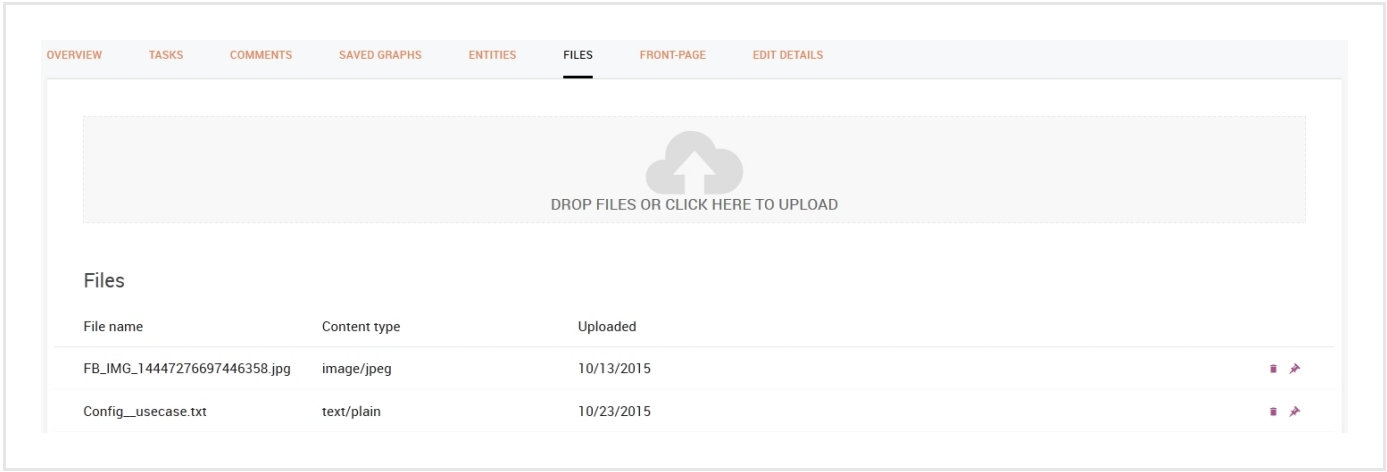

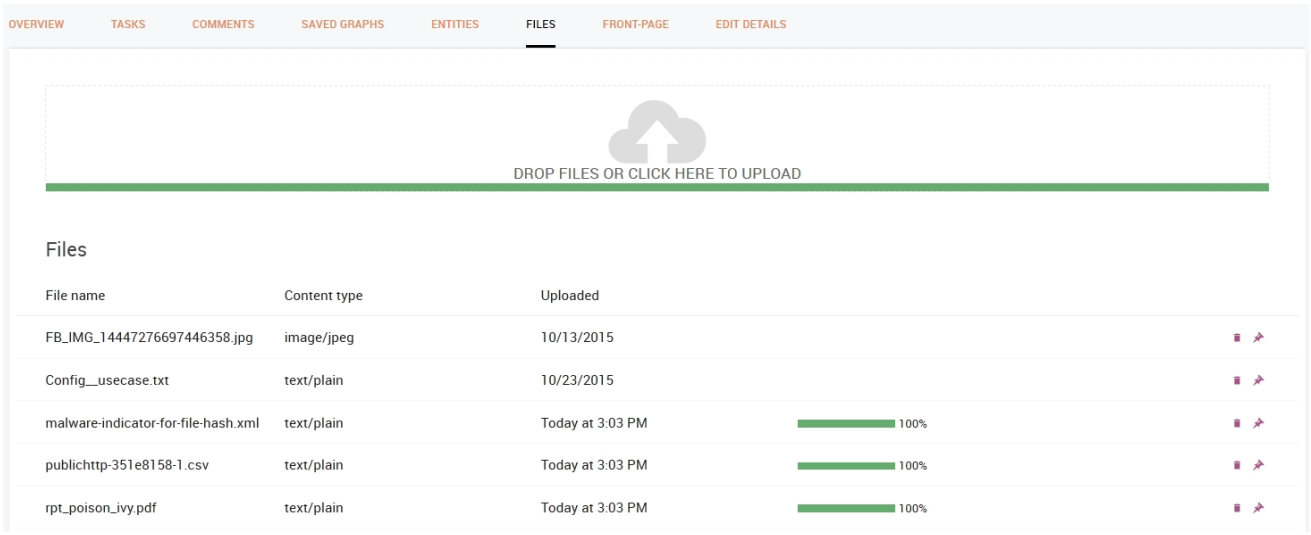

Workspace Edit details tab

Use the Edit details tab to change the basic structure and information of a workspace.

A workspace can evolve in time to adapt to collaborators joining or leaving the workspace, or to changes in scope and purpose.

You can update workspace information and carry out basic workspace maintenance in the **Edit details** tab.

Except for the **Is Public** flag, the details you edit here correspond to the ones you fill out when you create a new workspace.

Archive and delete a workspace

Besides updating a workspace information, you can also archive it and delete it.

UI option	Description
Archive Workspace	Select this option to make the workspace read-only. The workspace cannot be updated any longer, and it is available for reference only. This is a non-permament change: you can restore archived workspaces anytime to make them available in write and read mode again.
Delete Workspace	Select this option to completely remove a workspace. This action cannot be undone: upon deletion, a workspace and its data are lost.

Use the graph

Load entities onto the graph to analyze them, explore relationships, and manipulate the 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 **Browse** or **Exposure** on the top navigation bar, or by clicking ⚙ **> Data management > Datasets > <dataset_name>** or ⚙ **> Data management > Incoming feeds > <incoming_feed_name>**.

- 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 onto 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:

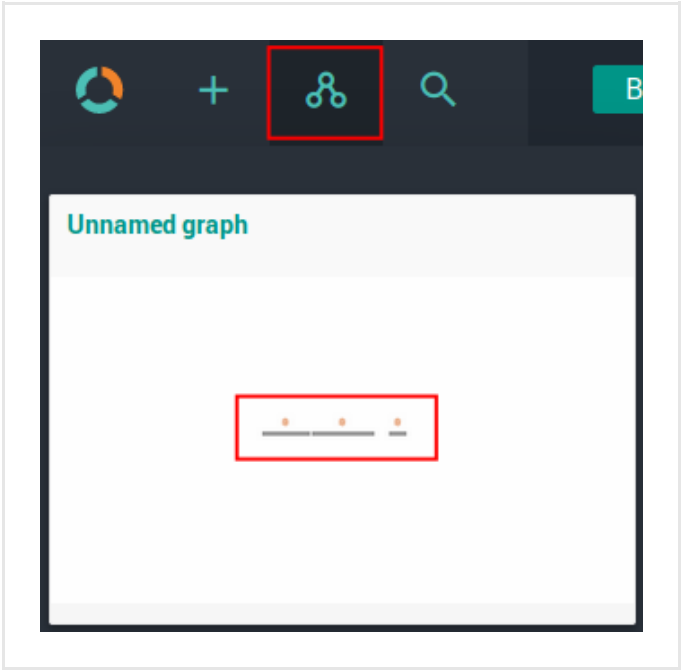

View the graph

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**.

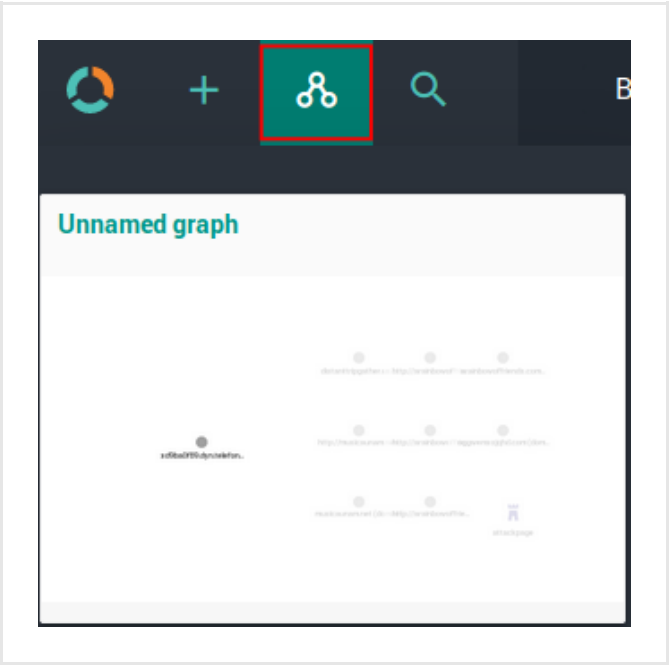

Open the graph from the Saved graphs section in a workspace

- 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**.

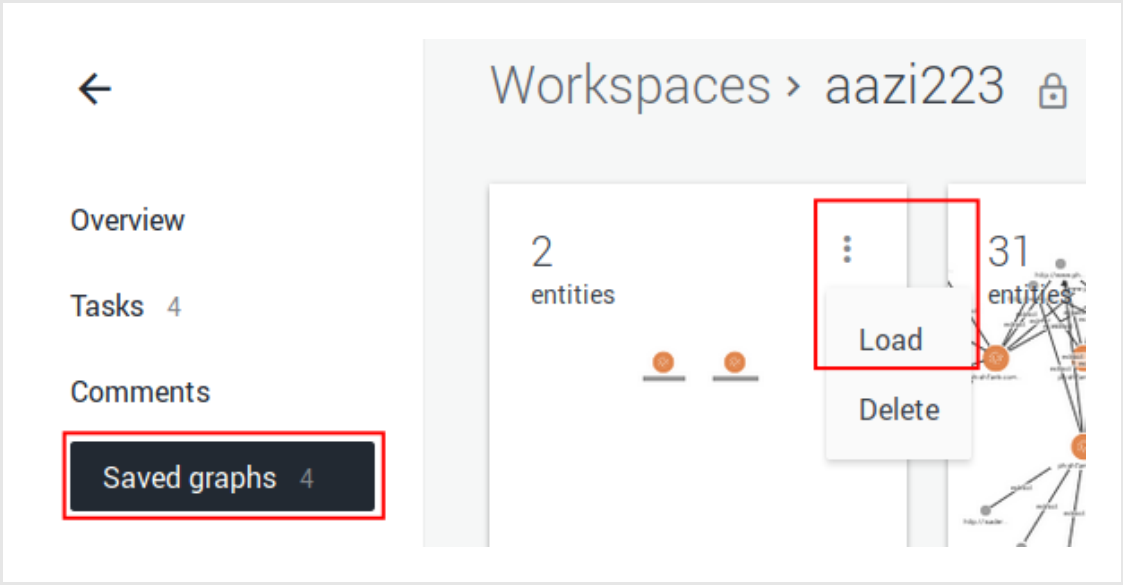

Open the graph from the Neighborhood tab on the entity detail pane

- 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 onto 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.

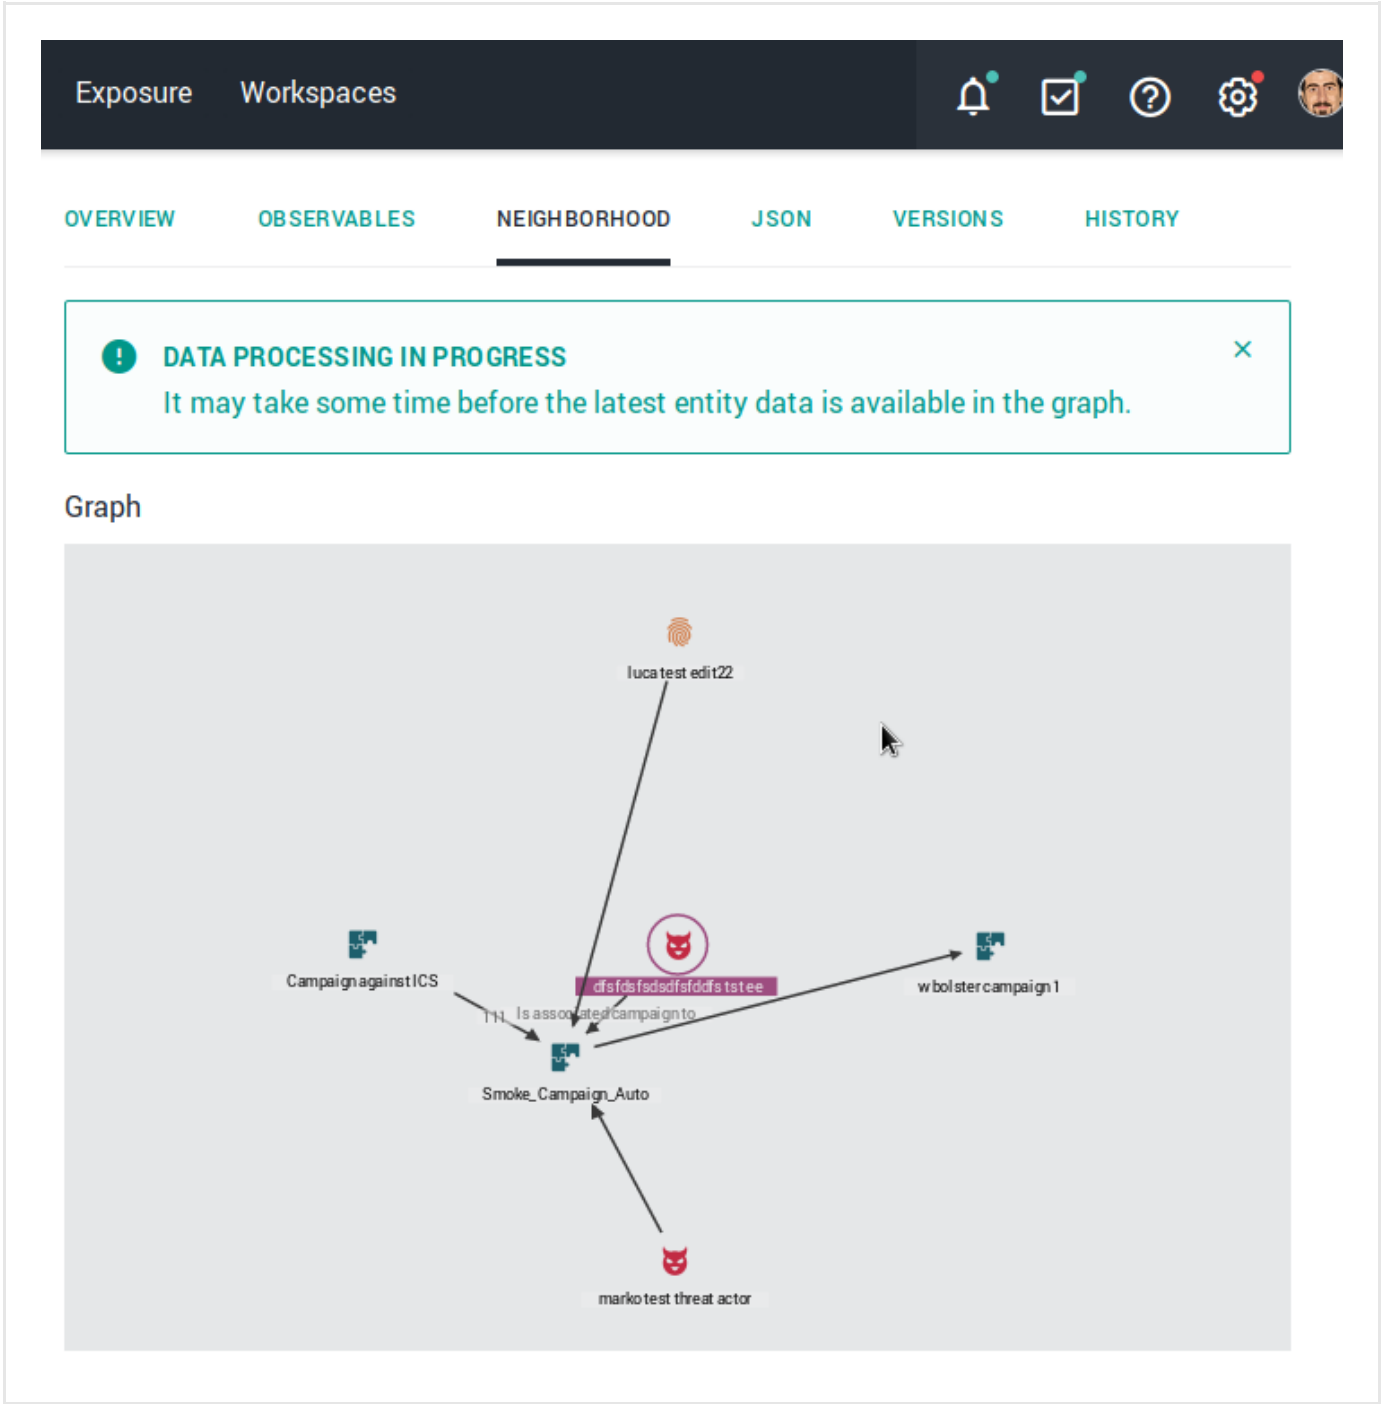

Examine entities on the graph

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.

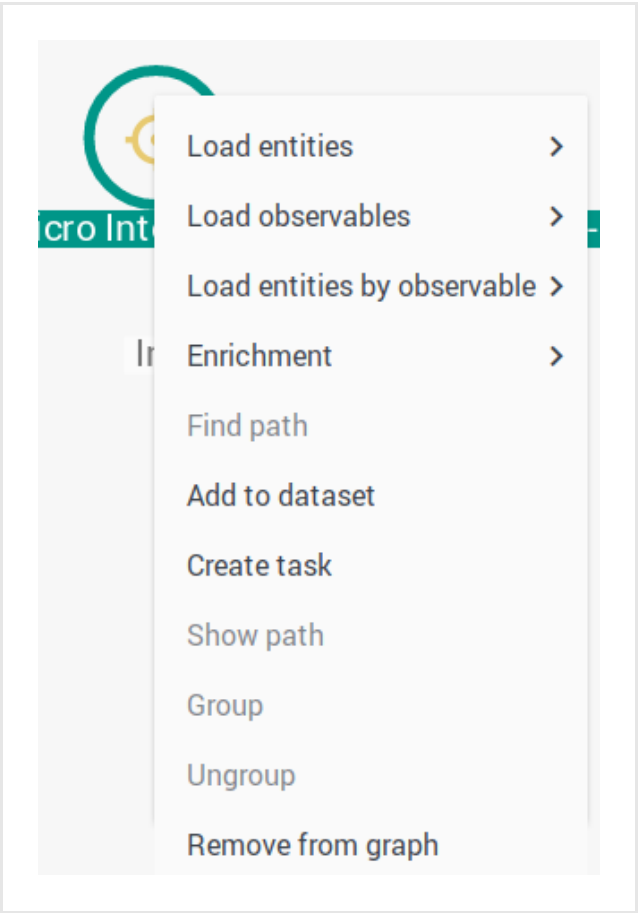

Look for and find relationships among entities and observables loaded on the graph

- **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.

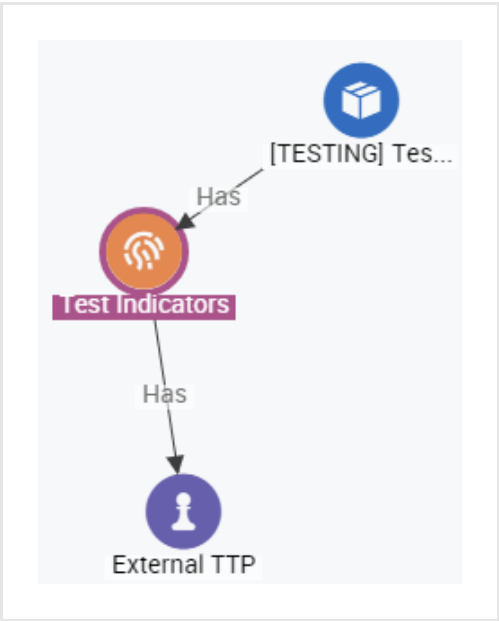

- **Load observables**: loads any observables related to the selected entity; for example, an IP address, or an email address related to the entity.
  From the **Load observables** submenu you can choose whether you want to load all related observables, or only specific related observable types.

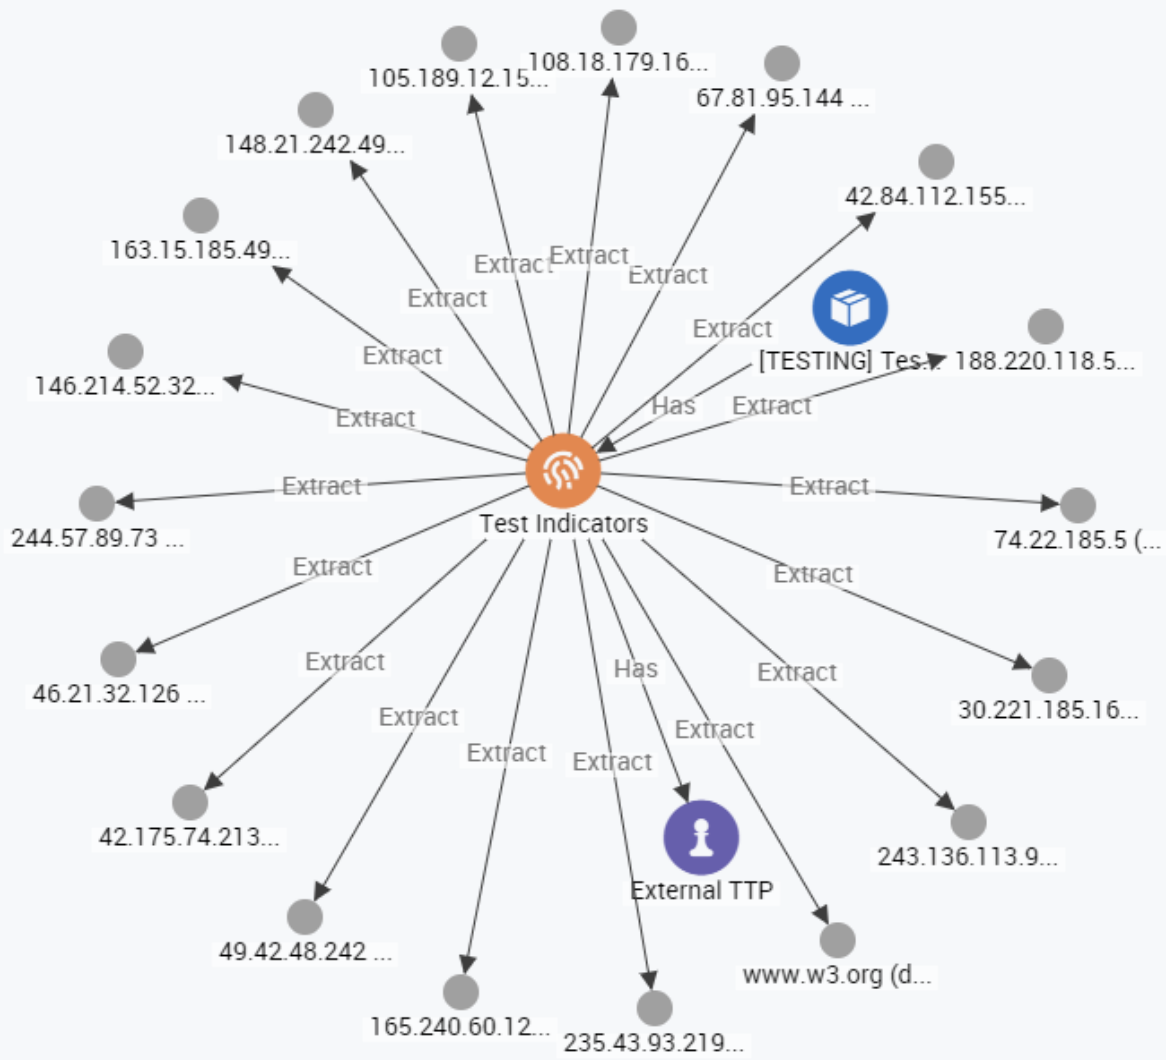

- **Load entities by observable**: loads entities that are *indirectly related* to the selected one through one intermediate observable node.
  From the **Load entities by observable** submenu you can choose whether you want to load indirectly related entities based on all the available observables,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

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.

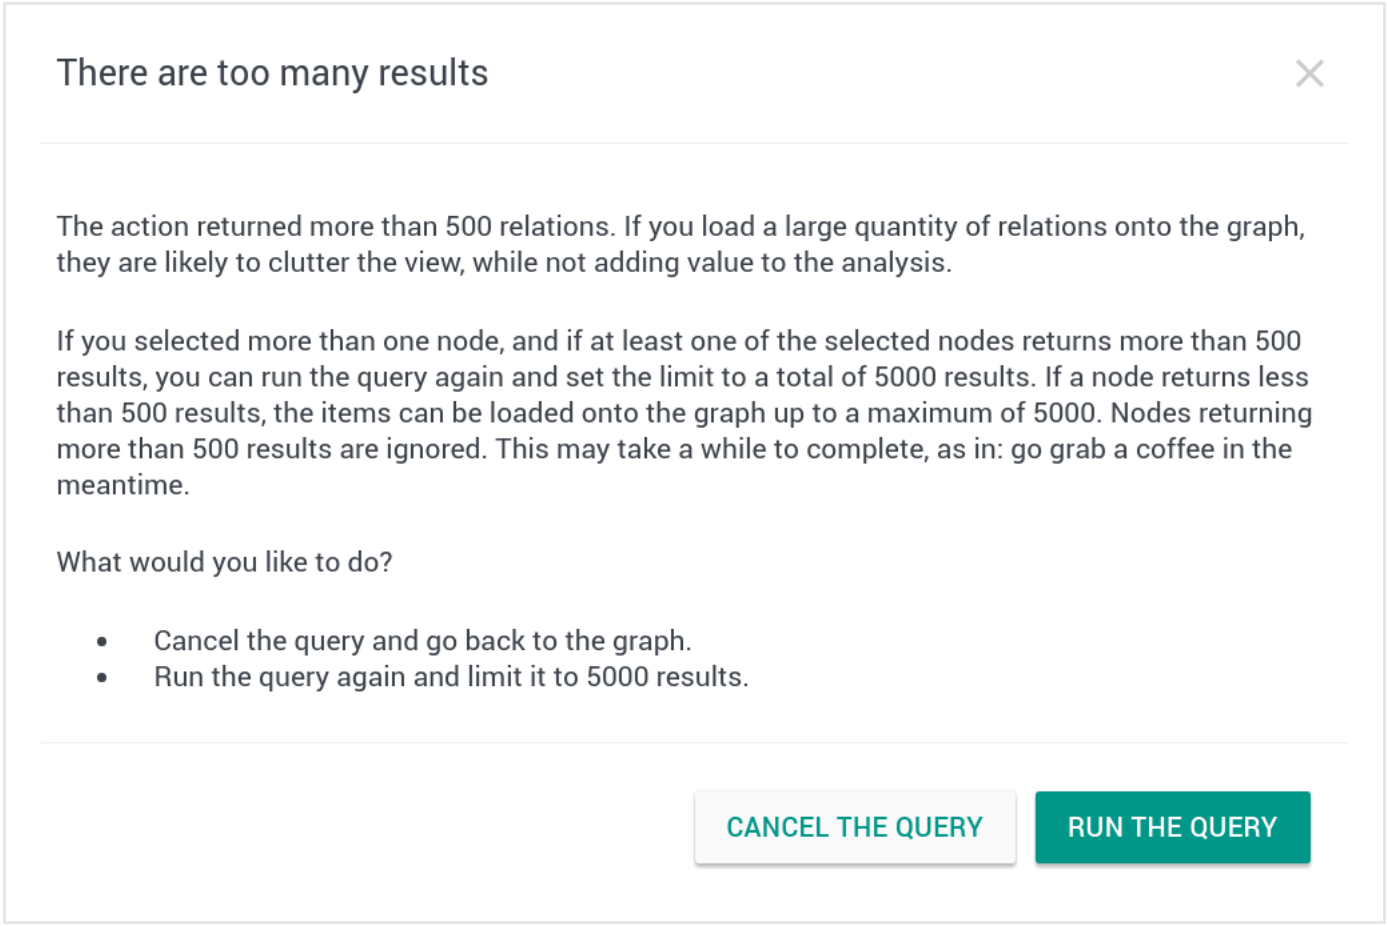

Enrich entities and observables on the fly

- **Enrichment**: manually triggers enricher tasks on the selected entities and observables. Any new enrichment observables are automatically loaded onto 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 asn 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.

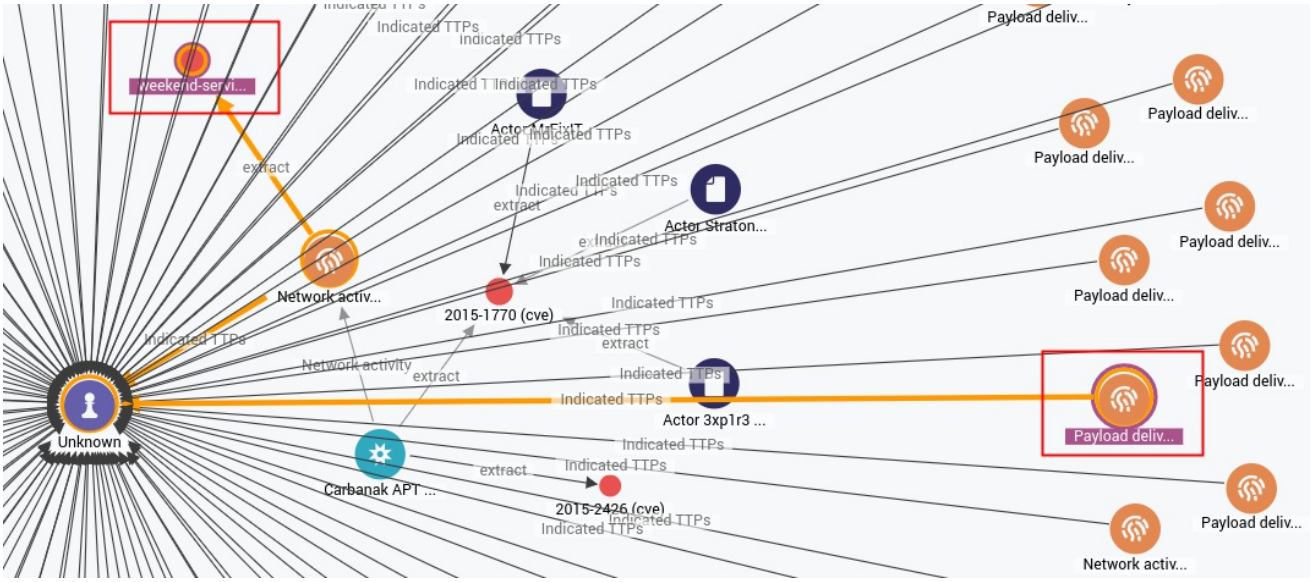

Include selected entities in a workflow

- **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.

Change visualization layout

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.

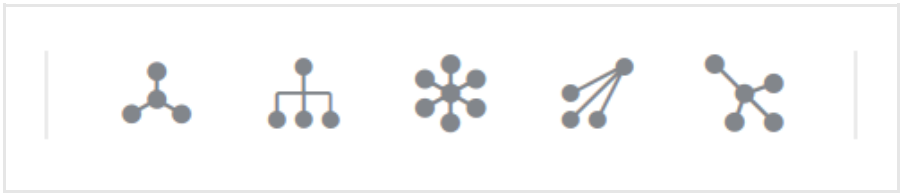

Layout type	Description
Standard	In the standard layout, links on the graph are a consistent length. Nodes and edges overlap as little as possible, and they are evenly distributed on the graph surface. It offers a consistent and clean view. It is a good starting point to begin analyzing any kind of data 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 onto 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.

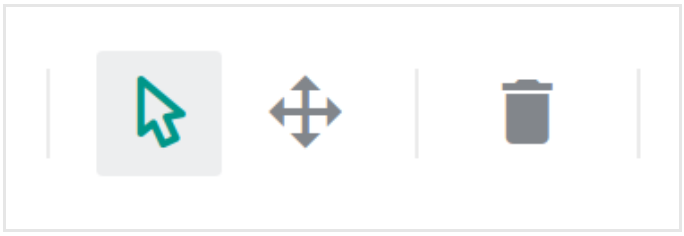

Datasets

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

A dataset is an arbitrary, unordered data collection: its content can be edited or deleted 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.

Create a dataset

To create a dataset, do the following:

- On the top navigation bar click **✚ > Data management > Dataset** .

Alternatively:

- On the top navigation bar click **⚙ > Data management > Datasets** .
- On the **Data management > Datasets** page, click **✚ > Dataset**
- On the **Data management > Datasets > Create** page, enter a name for the dataset under **Dataset name**.

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
```

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 or delete a dataset

To edit an existing dataset, do the following:

- On the top navigation bar click ⚙ **> Data management > Datasets** .
- On the **Data management > Datasets** page, click the click the ⋮ icon on the row corresponding to the dataset you want to modify.

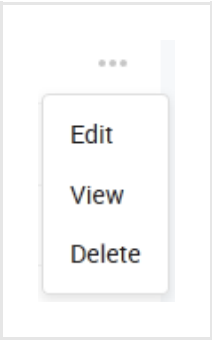

- From the drop-down menu select **View** to inspect the content of the dataset. It shows an overview where you can review the dataset entities and edit them, if necessary.
- From the drop-down menu select **Edit** to rename an existing dataset or to enable/disable the static/dynamic property by selecting/deselecting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

Tasks

Use tasks to assign and manage activities among your collaborators and to create a transparent workflow.

Tasks allow you to manage and distribute workload among your collaborators, keep track of progress, identify any bottlenecks, and create workflows to clearly document and control activities.

Create a task

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 to 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 to 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 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.

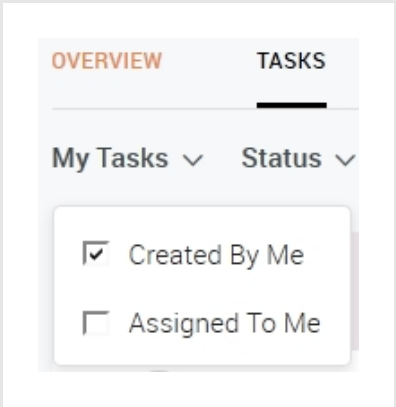

View tasks by status

- 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.

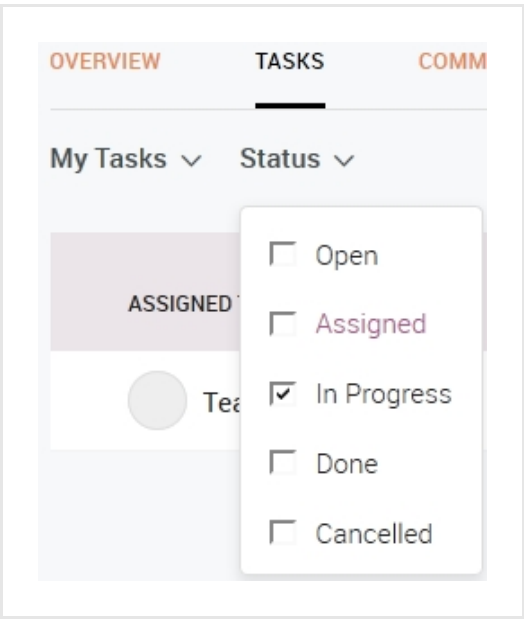

View task details

- On the 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**.

View 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.

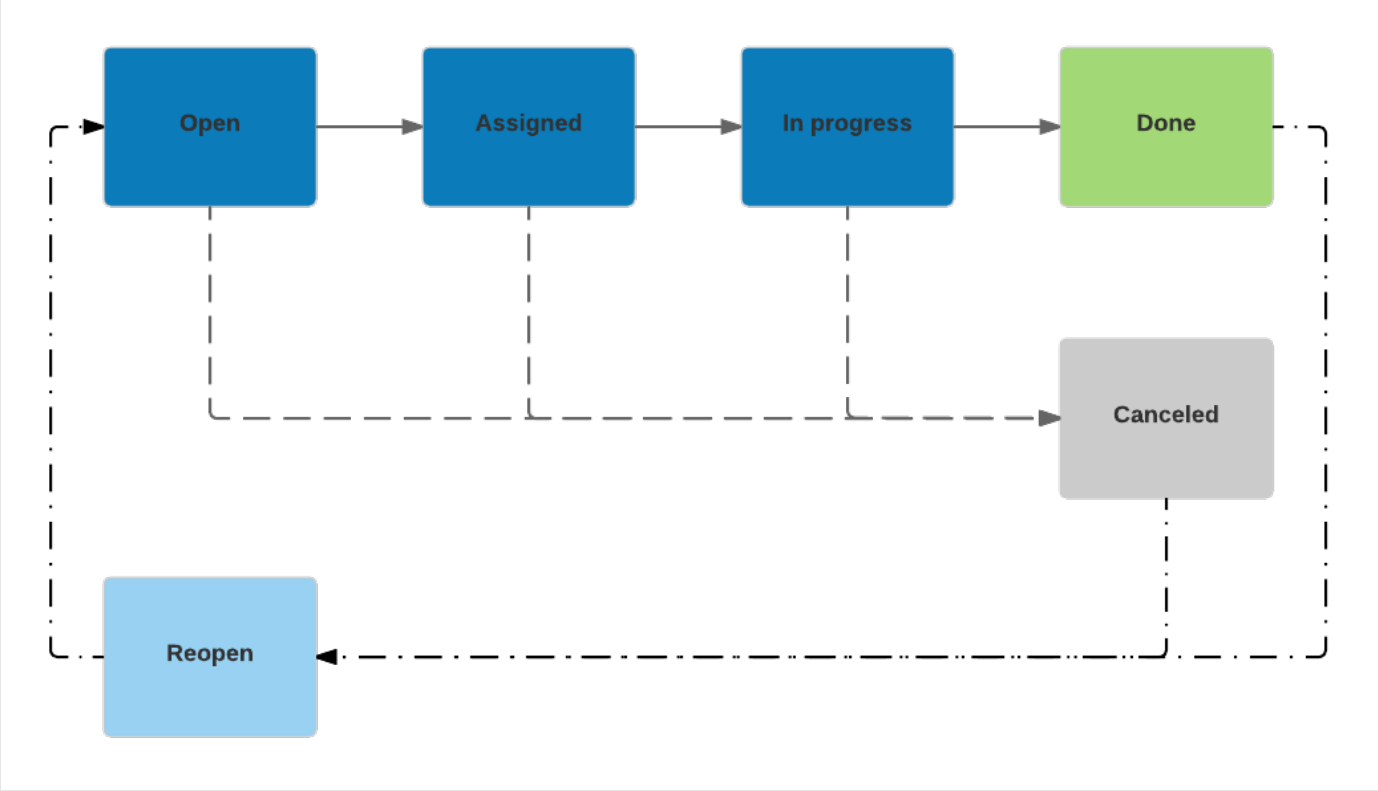

A standard task status flow

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.

Editor

In the entity editor you can create, edit and update detailed information about entities in a user-friendly form interface.

The entity editor helps you tell your cyber threat story by organizing and structuring your narrative around a cyber threat. You can use the editor to create and update entity details, add metadata, create or remove relationships with other entities, as well as assign entities to a workflow. The editor leverages the STIX standard without exposing it, so that you can concentrate on threat analysis and investigation, while the editor handles the underlying complexity.

All the information you submit through the entity editor is available for review in the entity detail pane.

Go to the entity editor

- On the top navigation bar click **✚ > Intelligence > **.

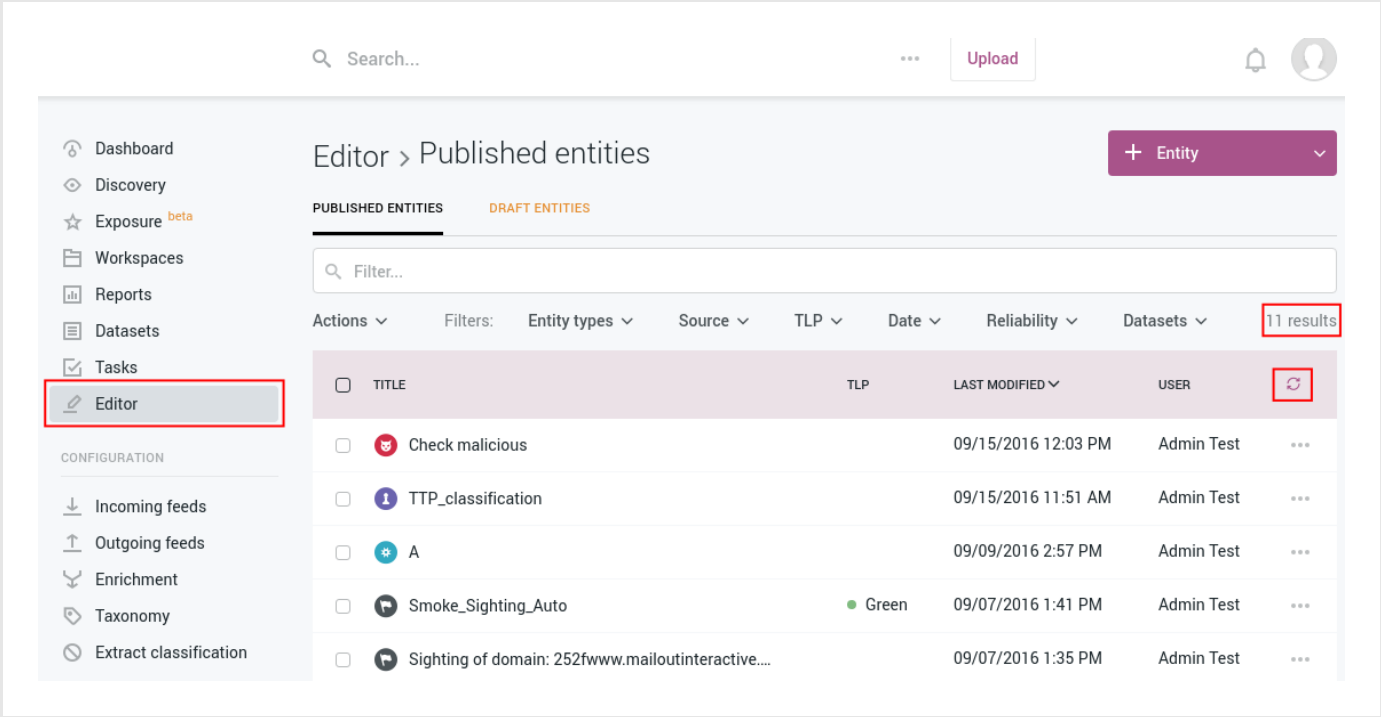

The entity editor opens at **Browse > Create **, and you can start adding details to describe the .

View entity details

To view detailed information about an entity, do the following:

- Go to the entity editor published entity overview.
- Click anywhere on the row corresponding to the entity you want to inspect.
- An overlay slides in from the side of the screen.

The entity name is always visible at the top of the entity detail pane. Click the edit icon next to the entity name to edit it. Below the entity name, you can see summary metadata details:

- Ingestion time
- Entity author user name
- User group the entity author belongs to
- **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).

The TLP value is displayed on the bottom-right corner of the header section. Click it to change its value, if necessary.

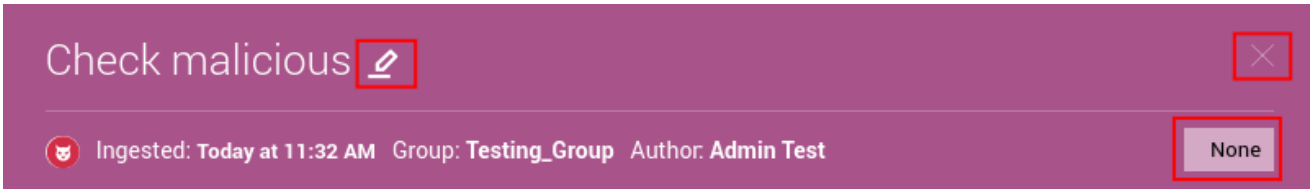

Overview

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 or 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 correspoding flag is **Yes**.

- **Tasks**: actionable user tasks associated to 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.

Enrichments

The **Observables** tab shows an overview of all observable enrichment observables related to the entity. Enrichers poll external data sources; enricher rules process the data and create the relationships between entities and retrieved observables.

If any updates are available from the enricher sources, the tab is populated with the results.

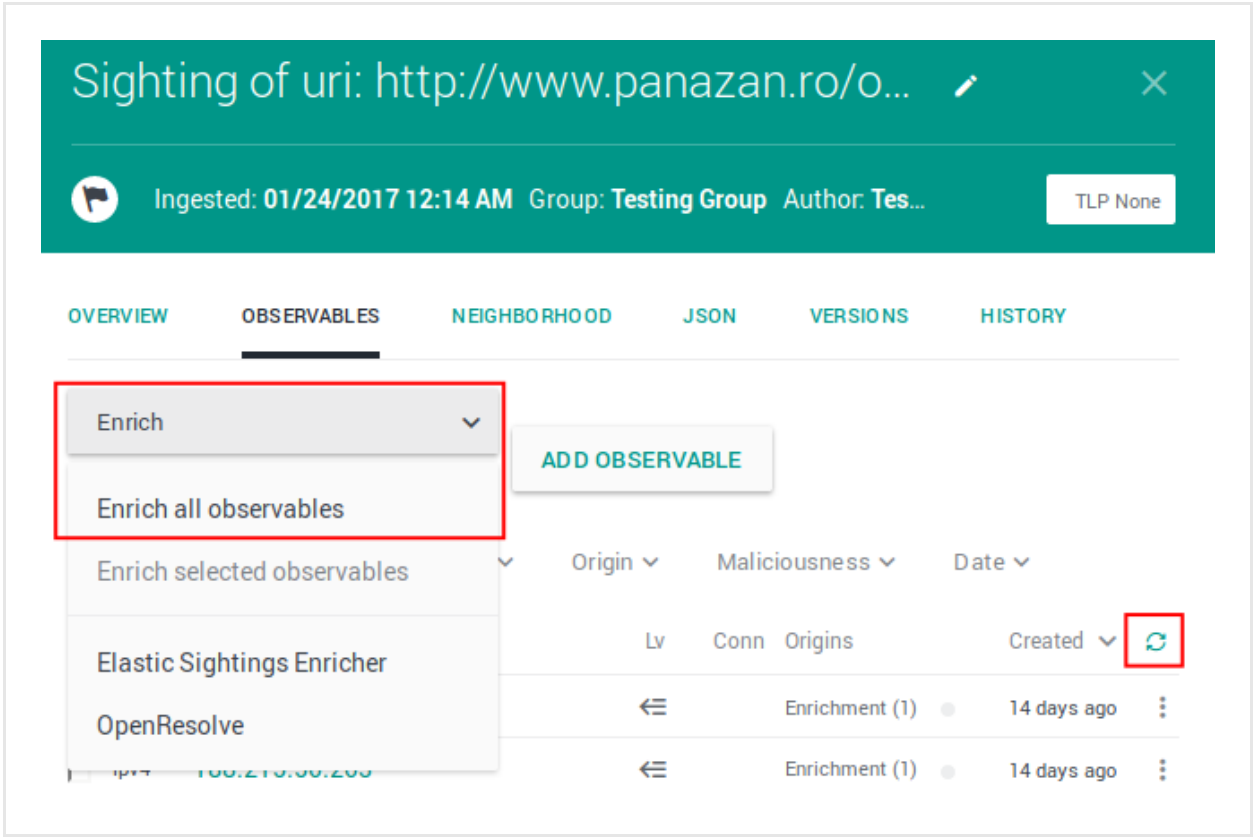

Neighborhood

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.

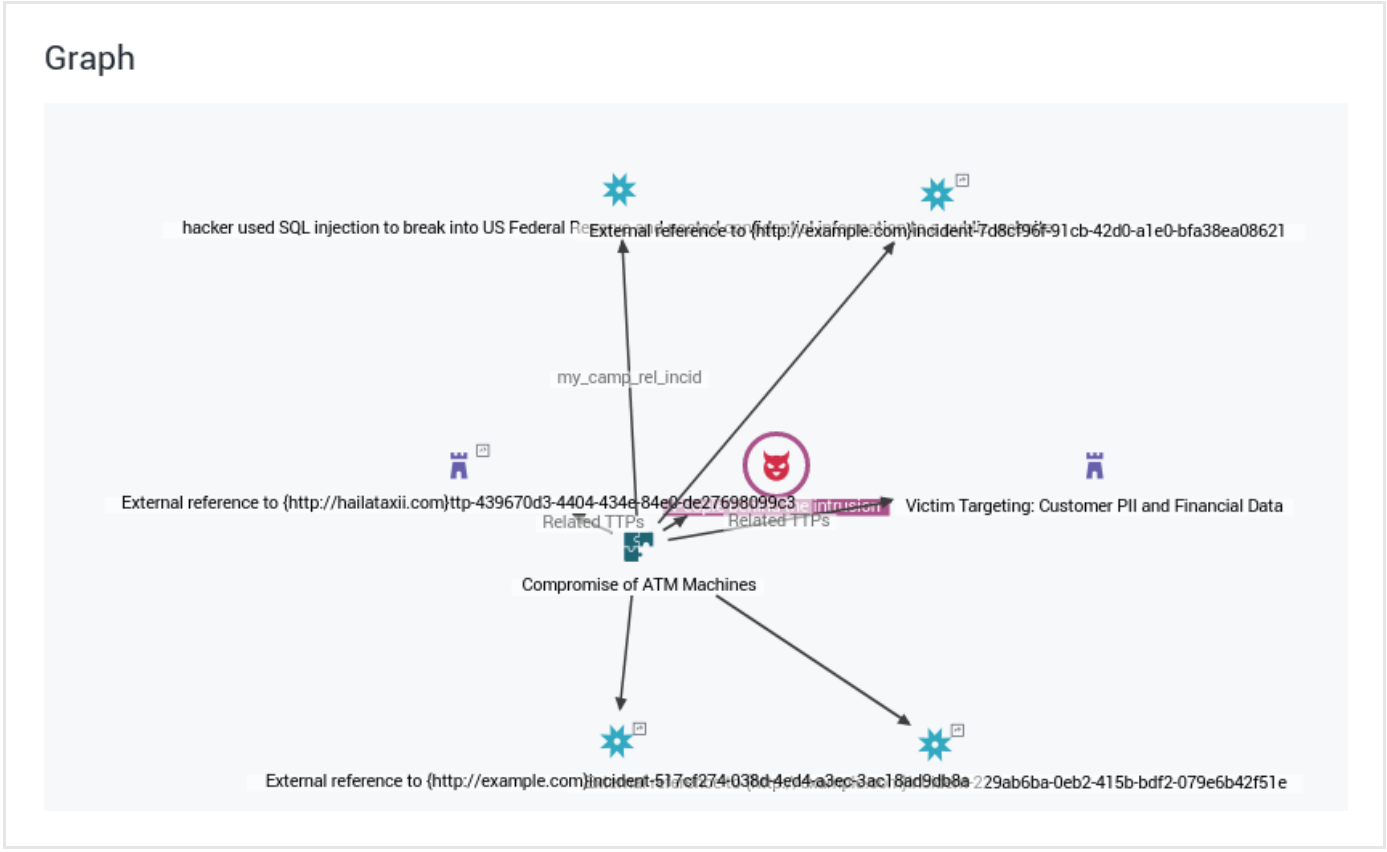

Click the embedded graph to load the entity and its neighborhood relationships onto the graph canvas, where you can further analyze the data.

The **Neighborhood** graph focuses on the immediate context around the entity. If the entity has more than 100 relationships, the **Neighborhood** graph displays only the 30 most recently created relationships. In this case, a notification message is displayed to inform the user:

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

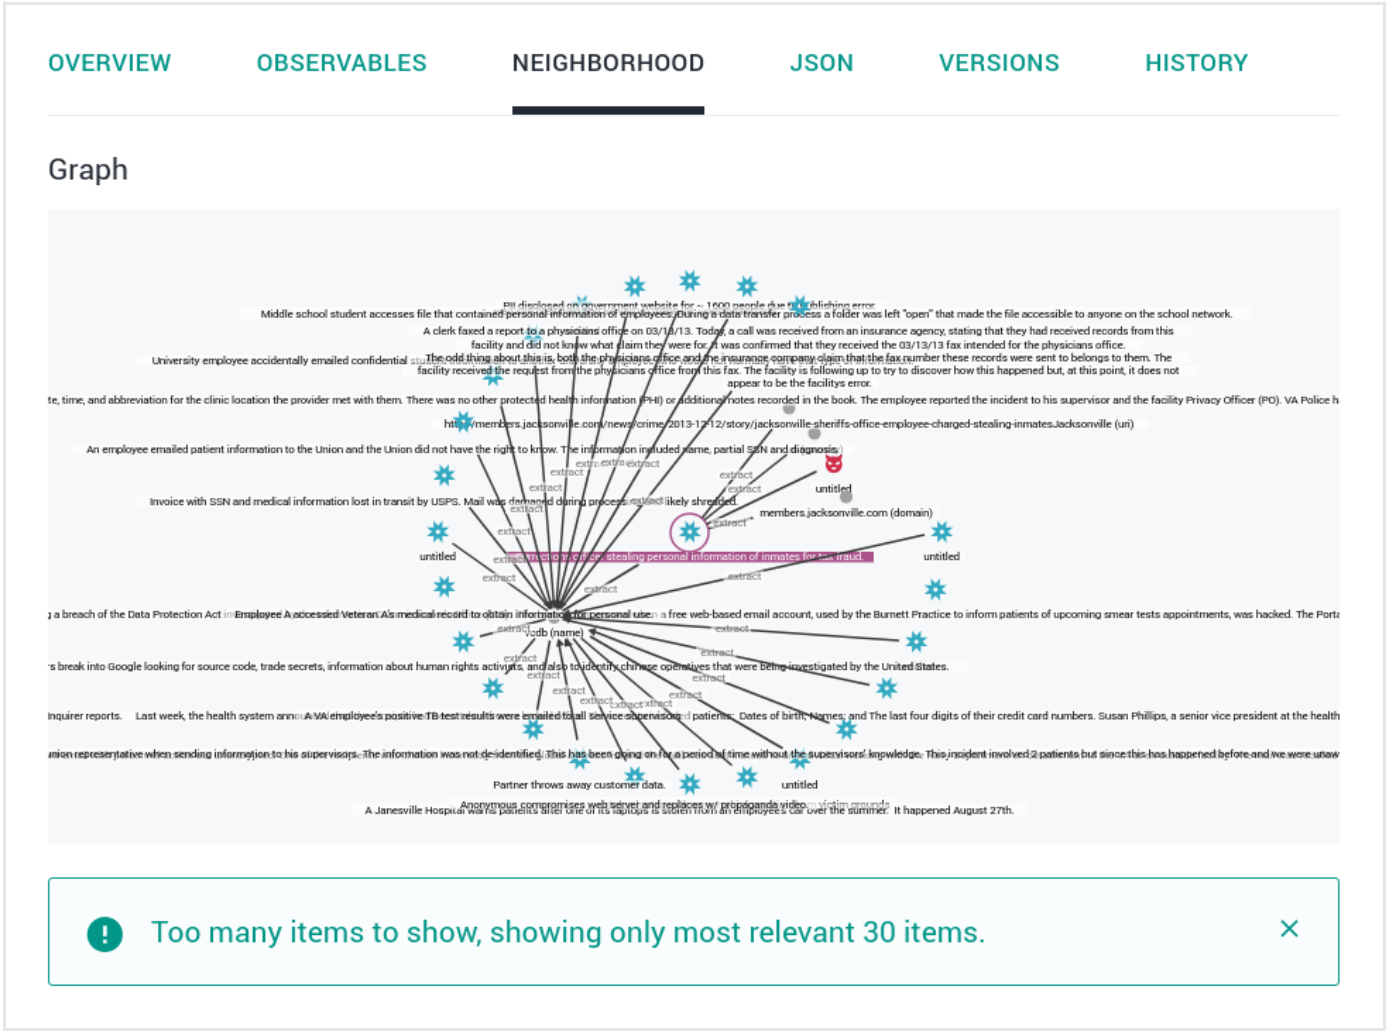

The embedded graph is a snapshot of the graph canvas view. The embedded snapshot is refreshed when accessing the **Neighborhood** tab, but it is not updated in real time. When the entity relationship landscape changes, for example, after adding or removing relationships, the embedded graph is not in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

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

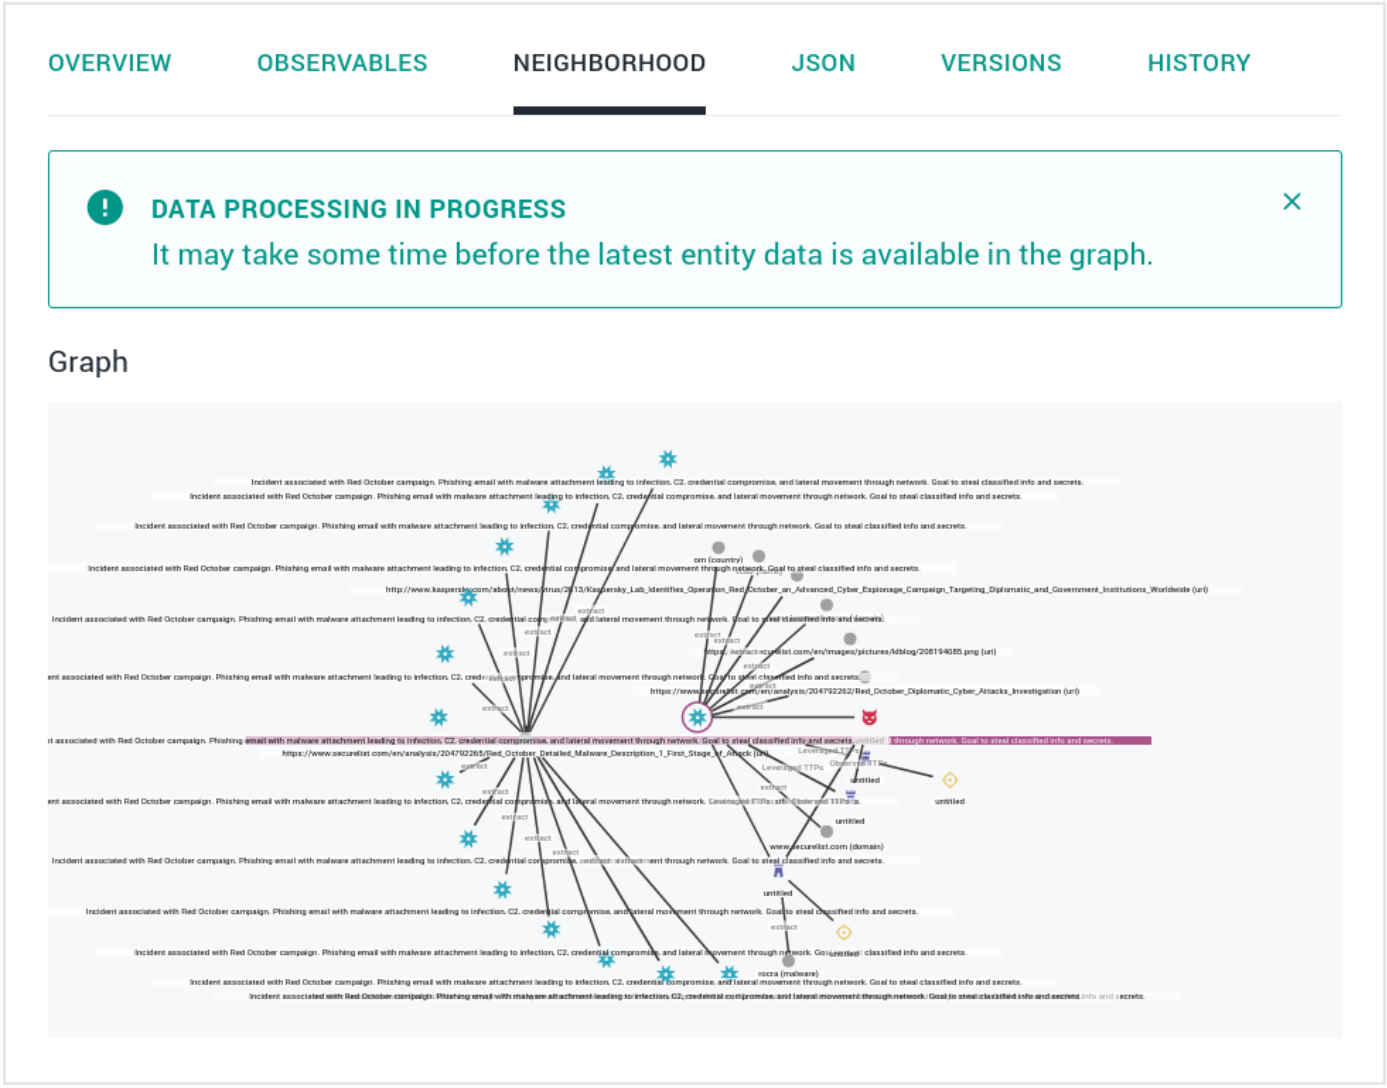

After the platform graph completes indexing, the embedded graph is back in sync. The time this task requires varies, depending on the size of the graph queue.

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.

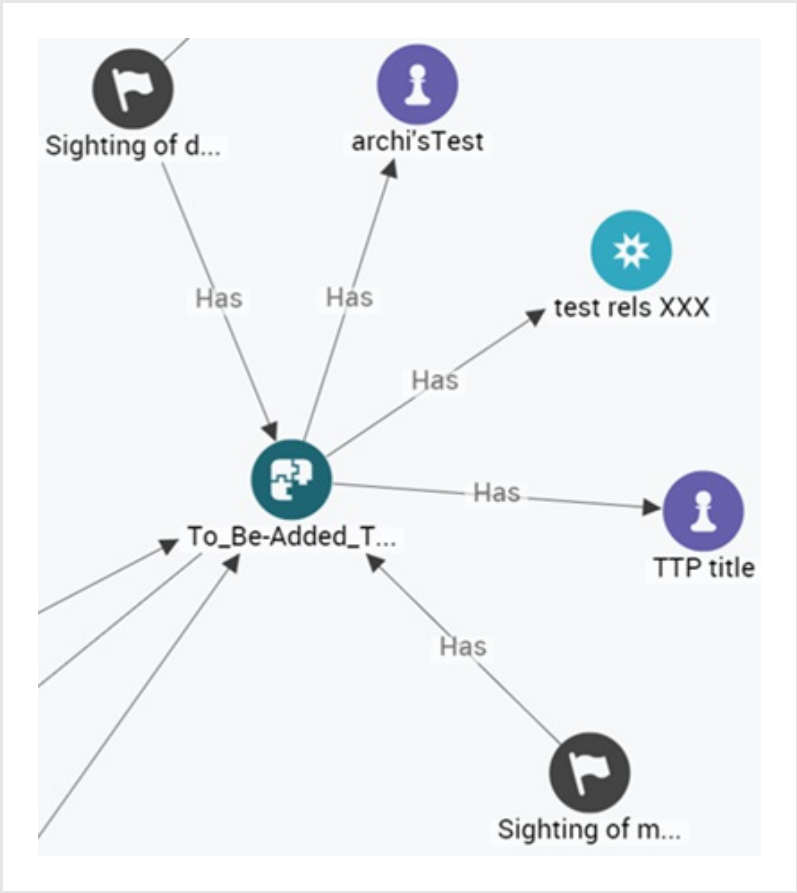

Entities with direct relationships

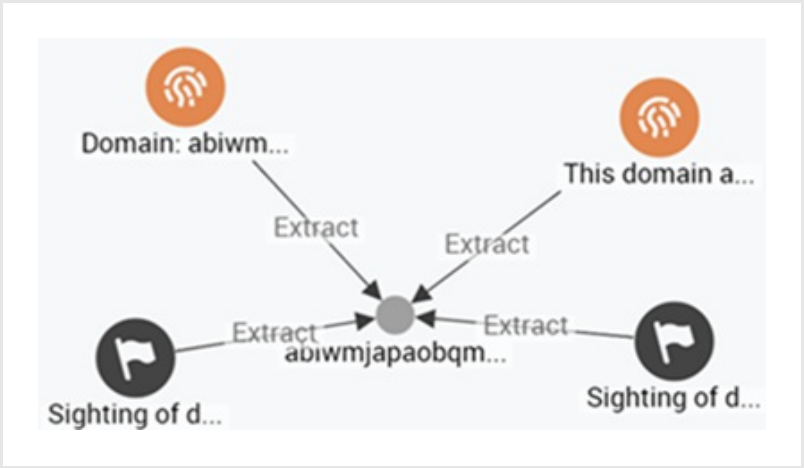

Entities with indirect relationships

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**.

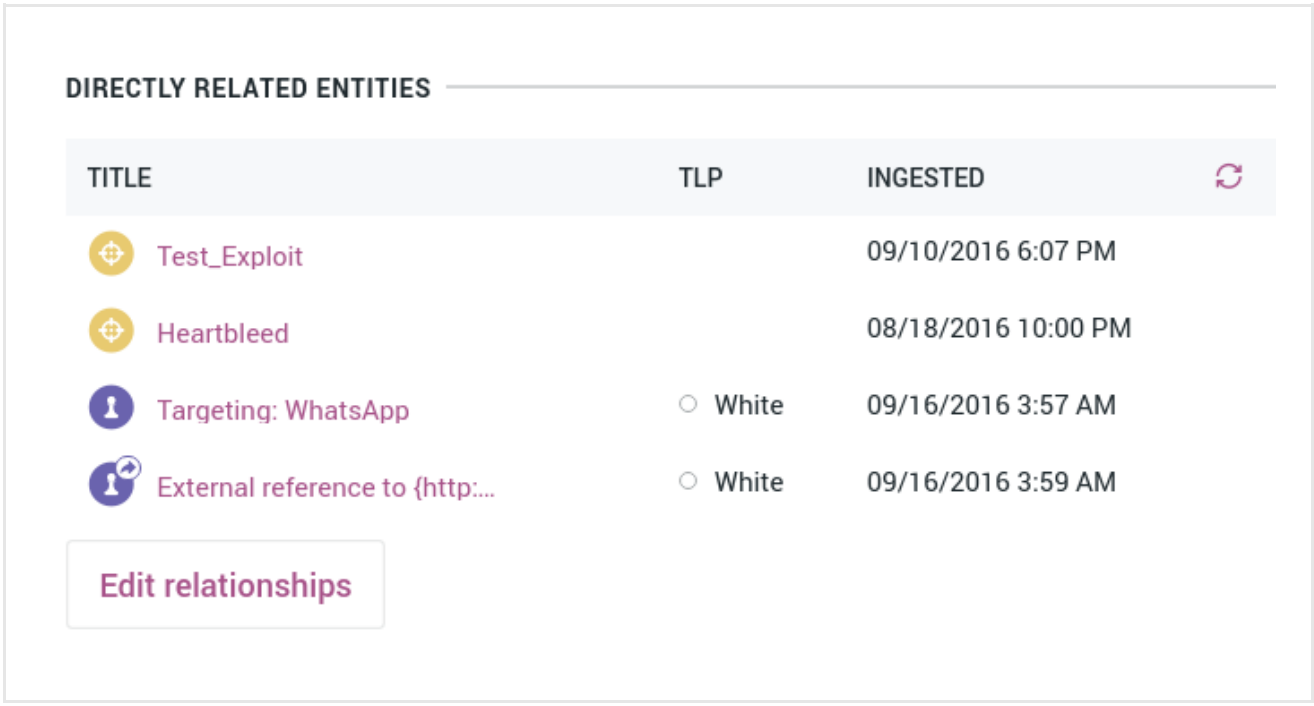

Entities related through observables

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.

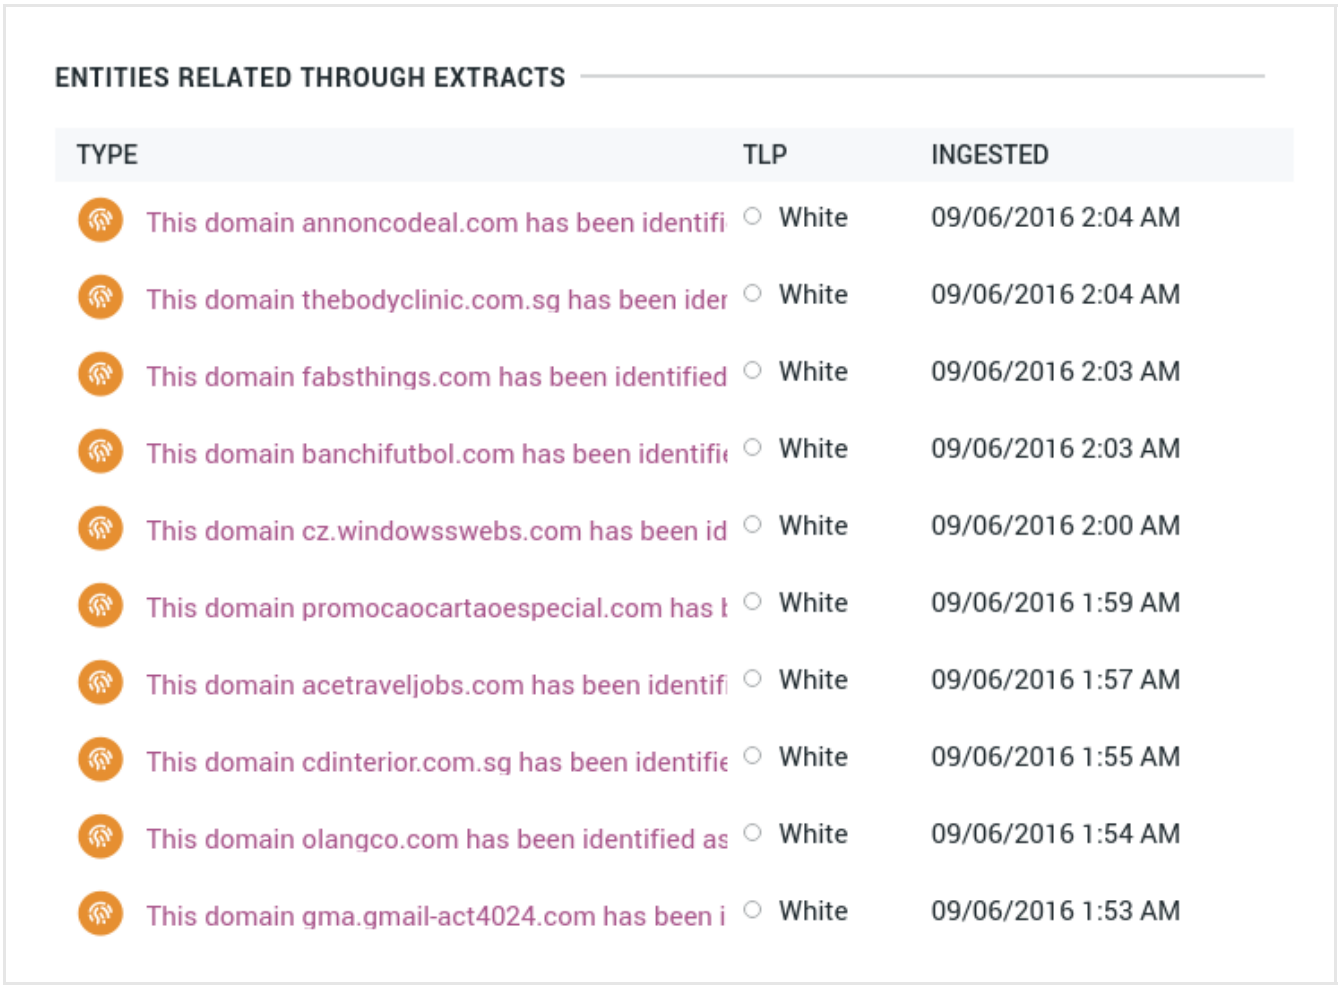

JSON

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.

Versions

If there are multiple versions of the entity in the platform, they are listed here for reference and comparison.

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).

Entity types

In the editor you can work with the following entity types:

Entity type	Description
Campaign (https://stixproject.github.io/data-model/1.2/campaign/campaigntype/)	A campaign is a series of planned actions aiming at achieving a specific goal. It groups a set of related threat actors, TTPs, and incidents sharing a common intent or goal.
Course of action (https://stixproject.github.io/data-model/1.2/coa/courseofactiontype/)	A course of action details a set of clear, specific recommendations and measures to mitigate an incident, address affected exploit targets, and effectively respond to a cyber threat.
Exploit target (https://stixproject.github.io/data-model/1.2/et/exploittargettype/)	An exploit target is a vulnerability or a weakness in software, hardware, systems, or networks that a threat actor can leverage and take advantage of to intrude or carry out an attack.
Incident (https://stixproject.github.io/data-model/1.2/incident/incidenttype/)	An incident describes a specific occurrence of one or more indicators affecting an organization. It includes information on threat actors, tools or skills, timeframes, techniques, as well as impact assessment and the recommended response course of action.

Entity type	Description
Indicator (https://stixproject.github.io/data-model/1.2/indicator/indicatortype/)	An occurrence or a sign that an incident may have occurred or may be in progress. See also the definition provided in the **Cybersecurity Information Sharing Act of 2015 (CISA)** (https://www.congress.gov/bill/114th-congress/senate-bill/754/text).
Report (https://stixproject.github.io/data-model/1.2/report/reporttype/)	A detailed account of an indicator of compromise (IOC), a threat, a campaign or other threat activity as a result of an investigation or an analysis. A report tells a story about a piece of threat intelligence by providing background, context, and by pulling threads together to weave a clear and meaningful description of a security breach, a cyber attack, or a series of attacks.
Sighting ()	A sighting records a specific observation of a malicious indicator by matching fingerprints. For example, it can record the occurrence of a malicious IP address at a specific date and time,
Threat actor (https://stixproject.github.io/data-model/1.2/ta/threatactortype/)	An individual or a group carrying out or planning to execute malicious activities. Threat actors include information on their identity, suspected motivation, and suspected intended effect.
TTP (https://stixproject.github.io/data-model/1.2/ttp/ttptype/)	Tactics, Techniques and Procedures. Sometimes referred to also as Tools, Techniques, Procedures. TTPs describe the behavior of cyber adversaries. Tactics describe *"the employment and ordered arrangement of forces in relation to each other"*. Techniques are *"non-prescriptive ways or methods used to perform missions, functions, or tasks."* Procedures are *"standard, detailed steps that prescribe how to perform specific tasks."* (definitions from *"Joint Publication 1-02, Department of Defense Dictionary of Military and Associated Terms, 8 November 2010 (as amended through 15 February 2016)"*)
Package (https://stixproject.github.io/data-model/1.2/stix/stixtype/)	A package is a wrapper containing one or more STIX objects such as indicators, threat actors, TTPs, and so on. When the platform ingests packages, it extracts the STIX objects and it converts them to its internal JSON data model.

For further information on building and structuring entities, see the **STIX data model** (http://stixproject.github.io/data-model/) and the recommendations about using a **controlled vocabulary** (http://stixproject.github.io/documentation/concepts/controlled-vocabularies/).

Create an entity

To create a new entity, do the following:

- On the left-hand navigation sidebar, click **Editor**.
- Click the ✚ **Entity** button.
- From the drop-down menu select the entity type you want to create.

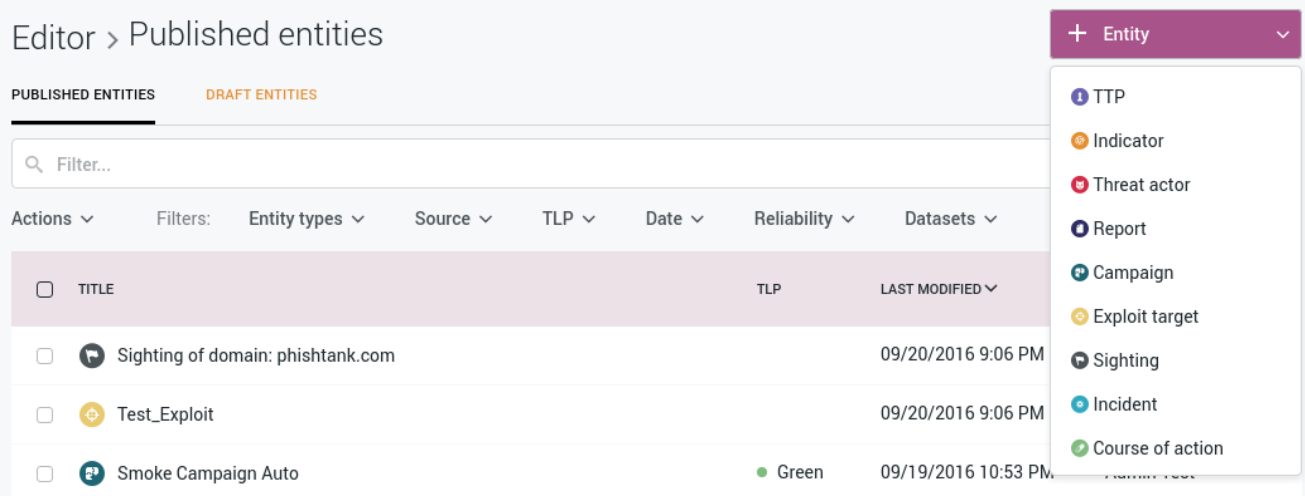

The entity editor is displayed, and you can proceed to create a new entity.

✔ On the forms, input fields marked with an asterisk are required.

Title

Applies to the following entities: *all entity types*
Specify the name of the new entity. It should be descriptive and easy to remember.

Analysis

Applies to the following entities: *all entity types but Report*
it is a free-text input field to include non-structured information such as additional context, references, links, and so on.

Description

Applies to the following entities: *Report*
You can enter one or more free-text short descriptions to sum up the report content.
To add a description field, click the **+ More** link.
To remove a description field, click the corresponding 🗑 icon.

Upload attachments

Applies to the following entities: *Report*
You can upload files by dragging and dropping them onto the highlighted upload area.
Alternatively, click anywhere on the upload area, browse to the location where the file you want to upload is stored, and then select it.
To remove an uploaded file from the attachment list, click the **Remove file** link.

Intents

Applies to the following entities: *Report*
From the drop-down menu select one or more options to specify the purpose of the report and the threat scope it focuses on.
Available values:

- **Collective Threat Intelligence**
- **Threat Report**
- **Indicators**

- **Indicators - Phishing**
- **Indicators - Watchlist**
- **Indicators - Malware Artifacts**
- **Indicators - Network Activity**
- **Indicators - Endpoint Characteristics**
- **Campaign Characterization**
- **Threat Actor Characterization**
- **Exploit Characterization**
- **Attack Pattern Characterization**
- **Malware Characterization**
- **TTP - Infrastructure**
- **TTP - Tools**
- **Courses of Action**
- **Incident**
- **Observations**
- **Observations - Email**
- **Malware Samples**

Types

Applies to the following entities: *Indicator*, *Threat actor*

Indicator types

From the drop-down menu select one or more options to provide additional information on the type of indicator you are creating, for example an indicator with guidelines and recommendations concerning an observable or a TTP. Available values:

- **Malicious E-mail**
- **IP Watchlist**
- **File Hash Watchlist**
- **Domain Watchlist**
- **URL Watchlist**
- **Malware Artifacts**
- **C2**
- **Anonymization**
- **Exfiltration**
- **Host Characteristics**
- **Compromised PKI Certificate**
- **Login Name**
- **IMEI Watchlist**
- **IMSI Watchlist**

Threat actor types

From the drop-down menu select one or more options to specify the type of threat actor you are describing.
Available values:

- **Cyber Espionage Operations**
- **Hacker**
- **Hacker - White hat**
- **Hacker - Gray hat**
- **Hacker - Black hat**
- **Hacktivist**
- **State Actor / Agency**
- **eCrime Actor - Credential Theft Botnet Operator**
- **eCrime Actor - Credential Theft Botnet Service**
- **eCrime Actor - Malware Developer**
- **eCrime Actor - Money Laundering Network**
- **eCrime Actor - Organized Crime Actor**
- **eCrime Actor - Spam Service**
- **eCrime Actor - Traffic Service**
- **eCrime Actor - Underground Call Service**
- **Insider Threat**
- **Disgruntled Customer / User**

Confidence

Applies to the following entities: *TTP*, *Indicator*, *Threat actor*, *Campaign*, *Sighting*, *Incident*
From the drop-down menu select an option to assign the entity a confidence value.
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 or trustworthiness of the entity information.

Impact / Likely impact

Applies to the following entities: *Indicator*, *Sighting*
From the drop-down menu select an option to assign the entity a value to estimate how heavy the consequences of the threat would be.
Impact provides an estimate of how seriously a threat can affect your organization and/or your infrastructure.
Allowed values:

- **Unknown**
- **None**
- **Low**
- **Medium**
- **High**

Status

Applies to the following entities: *Campaign*, *Incident*
From the drop-down menu select an option to define the current status of a campaign or an incident. Update it as needed to reflect the actual campaign or incident scenario
Campaign status

Available values:

- **Ongoing**
- **Historic**
- **Future**

Incident status

Available values:

- **New**
- **Open**
- **Stalled**
- **Containment Achieved**
- **Restoration Achieved**
- **Incident Reported**
- **Closed**
- **Rejected**
- **Deleted**

Names

Applies to the following entities: *Campaign*
Enter here one or more alternative name aliases the campaign is known by.
To confirm the current input and to display a new input field, press **ENTER**.
To remove an input field from this section, click the corresponding ✖ icon.

Categories

Applies to the following entities: *Incident*
From the drop-down menu select one or more options to specify the type of incident you are describing.
Available values:

- **Exercise/Network Defense Testing**
- **Unauthorized Access**
- **Denial of Service**
- **Malicious Code**
- **Improper Usage**
- **Scans/Probes/Attempted Access**
- **Investigation**

Intended effects

Applies to the following entities: *TTP*, *Campaign*, *Incident*
From the drop-down menu select an option to specify the purpose or the goal the cyber threat aims at achieving.
Available values:

- **Advantage**
- **Advantage - Economic**
- **Advantage - Military**

- **Advantage - Political**
- **Theft**
- **Theft - Intellectual Property**
- **Theft - Credential Theft**
- **Theft - Identity Theft**
- **Theft - Theft of Proprietary Information**
- **Account Takeover**
- **Brand Damage**
- **Competitve Advantage**
- **Degradation of Service**
- **Denial and Deception**
- **Destruction**
- **Disruption**
- **Embarrassment**
- **Exposure**
- **Extortion**
- **Fraud**
- **Harassment**
- **ICS Control**
- **Traffic Diversion**
- **Unauthorized Access**

Security compromise

Applies to the following entities:

Discovery methods

Applies to the following entities:

Characteristic

Applies to the following entities:
This field allows you to add extra details to more accurately describe the entity, for example by specifying the threat type, the resources it uses to spread and reach its target, or any connections with other entities.

Relations :

Estimated observed time

Applies to the following entities: *all entity types*
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.

Estimated threat start time

Applies to the following entities: *all entity types*
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.

Estimated threat end time

Applies to the following entities: *all entity types*
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.

Half life

Applies to the following entities: *all entity types*
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.

Tags

Applies to the following entities: *all entity types*
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**.

Source

Applies to the following entities: *all entity types*
From the drop-down menu select the source of the threat information you are using to create the new entity.

Source reliability

Applies to the following entities: *all entity types*
From the drop-down menu select a value to assess how reliable the source of the threat information is.
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)`.

References

Applies to the following entities:
Enter a URL pointing to relevant reference information on the threat, if available.
This field takes only URLs as input, and one URL per field.
To confirm the current input and to display a new input field, press **ENTER**.
To remove an input field from this section, click the corresponding ✖ icon.

TLP

Applies to the following entities: *all entity types but Report*

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).

Terms of use

Applies to the following entities: *all entity types*
Enter here any legal notes about fair use of the information about the entity.

Workflow

Applies to the following entities: *all entity types*
Add to dataset

Select this checkbox to associate the new entity to an existing workspace, and then from the drop-down menu select the target workspace.

Manually enrich

Select this checkbox to manually enrich the entity with the enricher sources you select from the drop-down menu.

After filling out the necessary input fields to record the new entity, you can save it as a draft or save it and publish it immediately:

- **Save draft**: Click **Save draft** to store your changes, or **Cancel** to discard them.
  The new entry is saved as a draft, but it is not published, i.e. it is inactive. It is available in the entity editor under **Draft entities**.
  If you are creating a new entity and you have not yet saved it for the first time, you can also click the drop-down arrow and select **Save draft and new** to:
  - Save the current populated form as a draft without publishing it to the platform;
  - Create and open a new draft form in the editor.
- **Publish**: Click **Publish** to store your changes, or **Cancel** to discard them.
  The new entry is saved and published to the platform. As soon as it is indexed, it is available in the platform in the entity editor under **Published**.
  Published entities associated to a workspace or included in a dataset are available also through the corresponding workspace and dataset.
  If you are creating a new entity and you have not yet saved it for the first time, you can also click the drop-down arrow and select **Publish and new** to:
  - Save the current populated form and publish it to the platform;
  - Create and open a new form in the editor.

Filter entities

You can filter entities to zero in on specific subsets:

- Use the drop-down menus above the entity list to select the criteria you want to apply to filter the entities.
- You can make multiple selections per menu and across the menus to further refine your results.
- The available menu options may vary, since they are based on the metadata of the entities in the workspace.

The available filter menus are:

Filter menu	Description	

Filter menu	Description	
Entity Types	Filter entities by type.	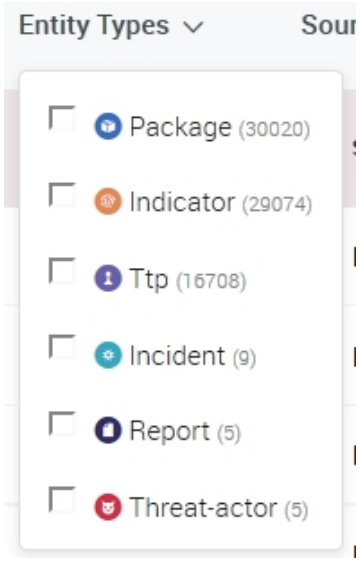
Source Types	Filter entities by source/origin.	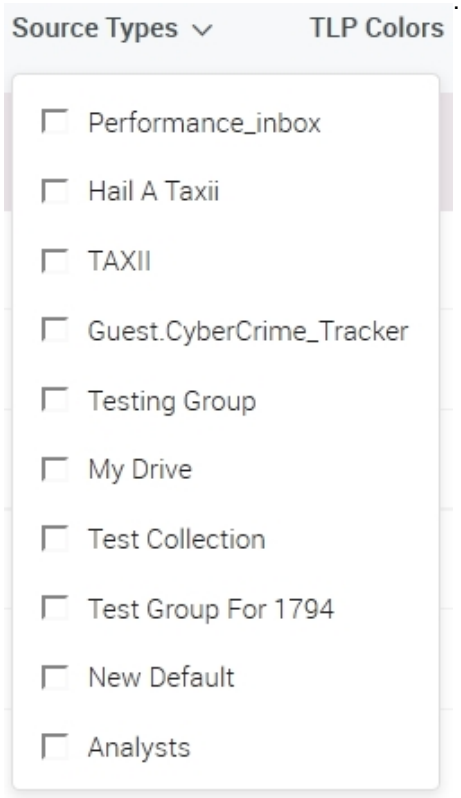
TLP Colors	Filter entities by Traffic Light Protocol color code.	

Filter menu	Description	
Date	Filter entities included in a date range.	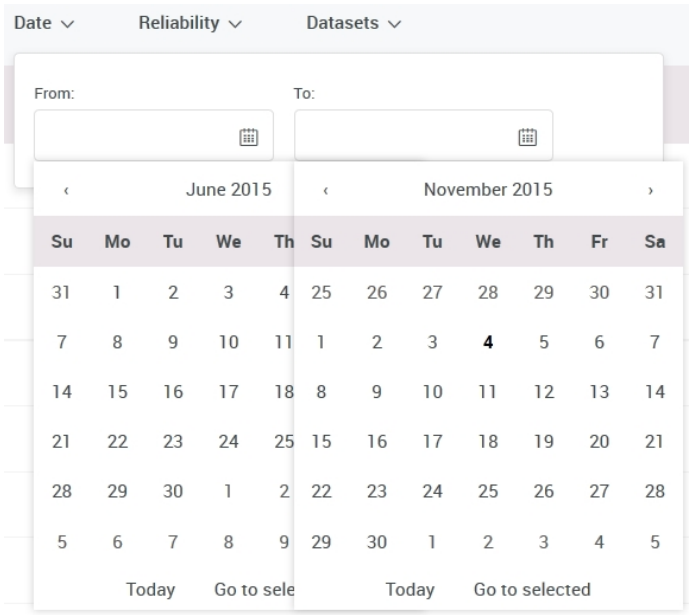
Reliability	Filter entities based on their reliability index.	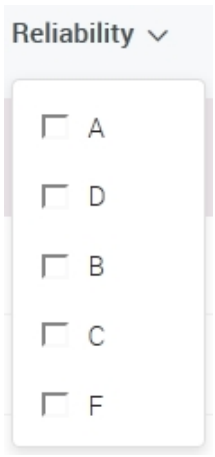
Datasets	Filter entities based on the dataset they belong to.	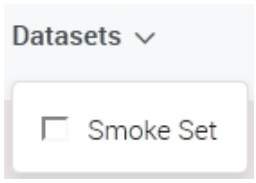

Manage entities

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 onto 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.