Getting started with EclecticIQ Platform

User guide for cyber threat analysts

Last generated: March 08, 2017

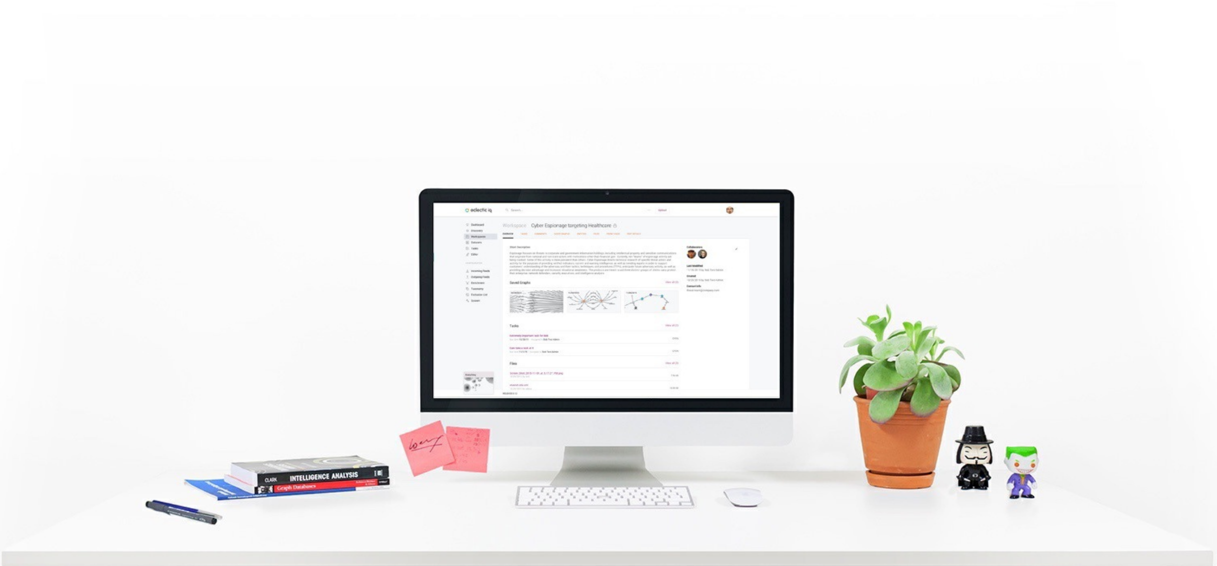

Last generated on Mar 8, 2017

Table of contents

Getting started with EclecticIQ Platform

Summary: 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

Last generated on Mar 8, 2017

Launch the platform

Summary: 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`
- On the login page, enter the appropriate credentials.

⚠ **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

ℹ 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 addr
```

The `inet addr` IP address is the one you need to use to access the platform.

Go to the platform

- In your host machine, launch a web browser.
- In the address bar, enter the VM IP address.
- The platform login screen is displayed.
- Log 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

Summary: 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.

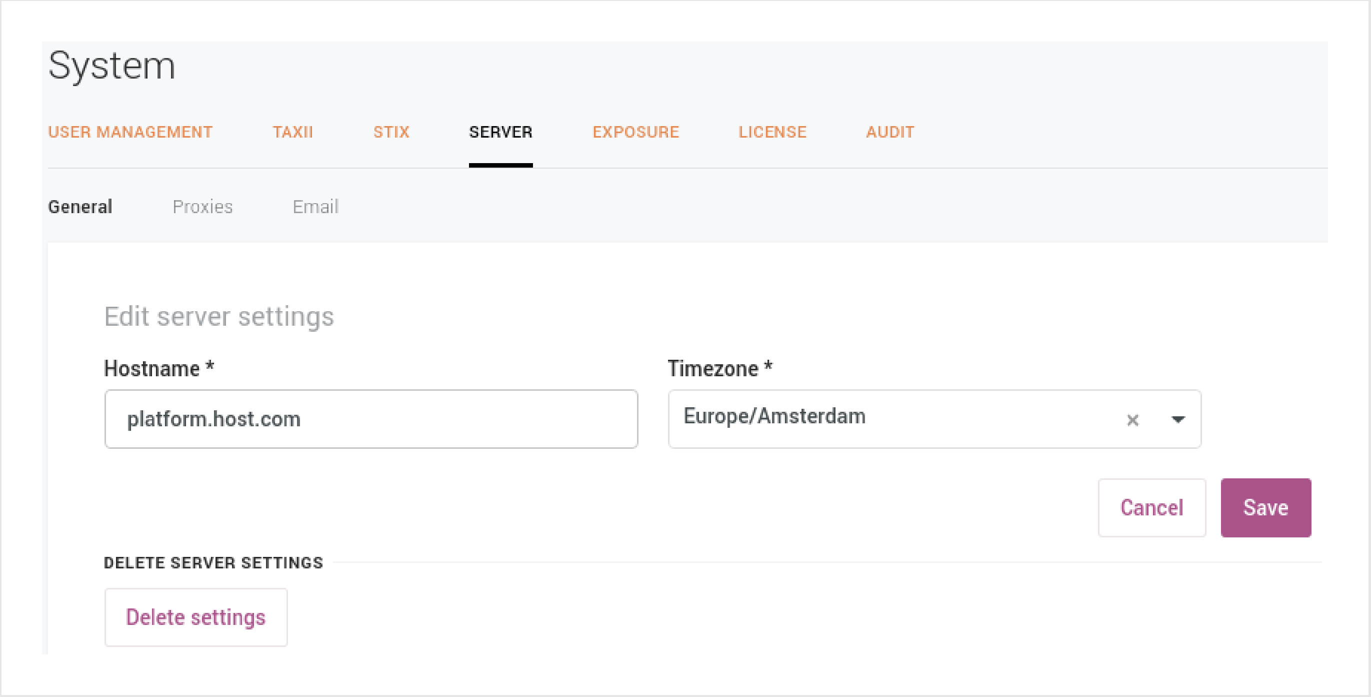

Hostname *: platform.host.com

Timezone *: Europe/Amsterdam

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 settings:
  - **Server**: the proxy server domain name, for example *host.com*.
  - **Port**: the proxy server access port, for example *9999*.
  - **Username**: valid user name credentials to authenticate and receive authorization to access the resource(s). For example, *nigeltufnel*..
  - **Password**: valid password credentials to authenticate and receive authorization to access the resource(s). For example, *s3cr3tp@SSw0rd_*..
  - **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.

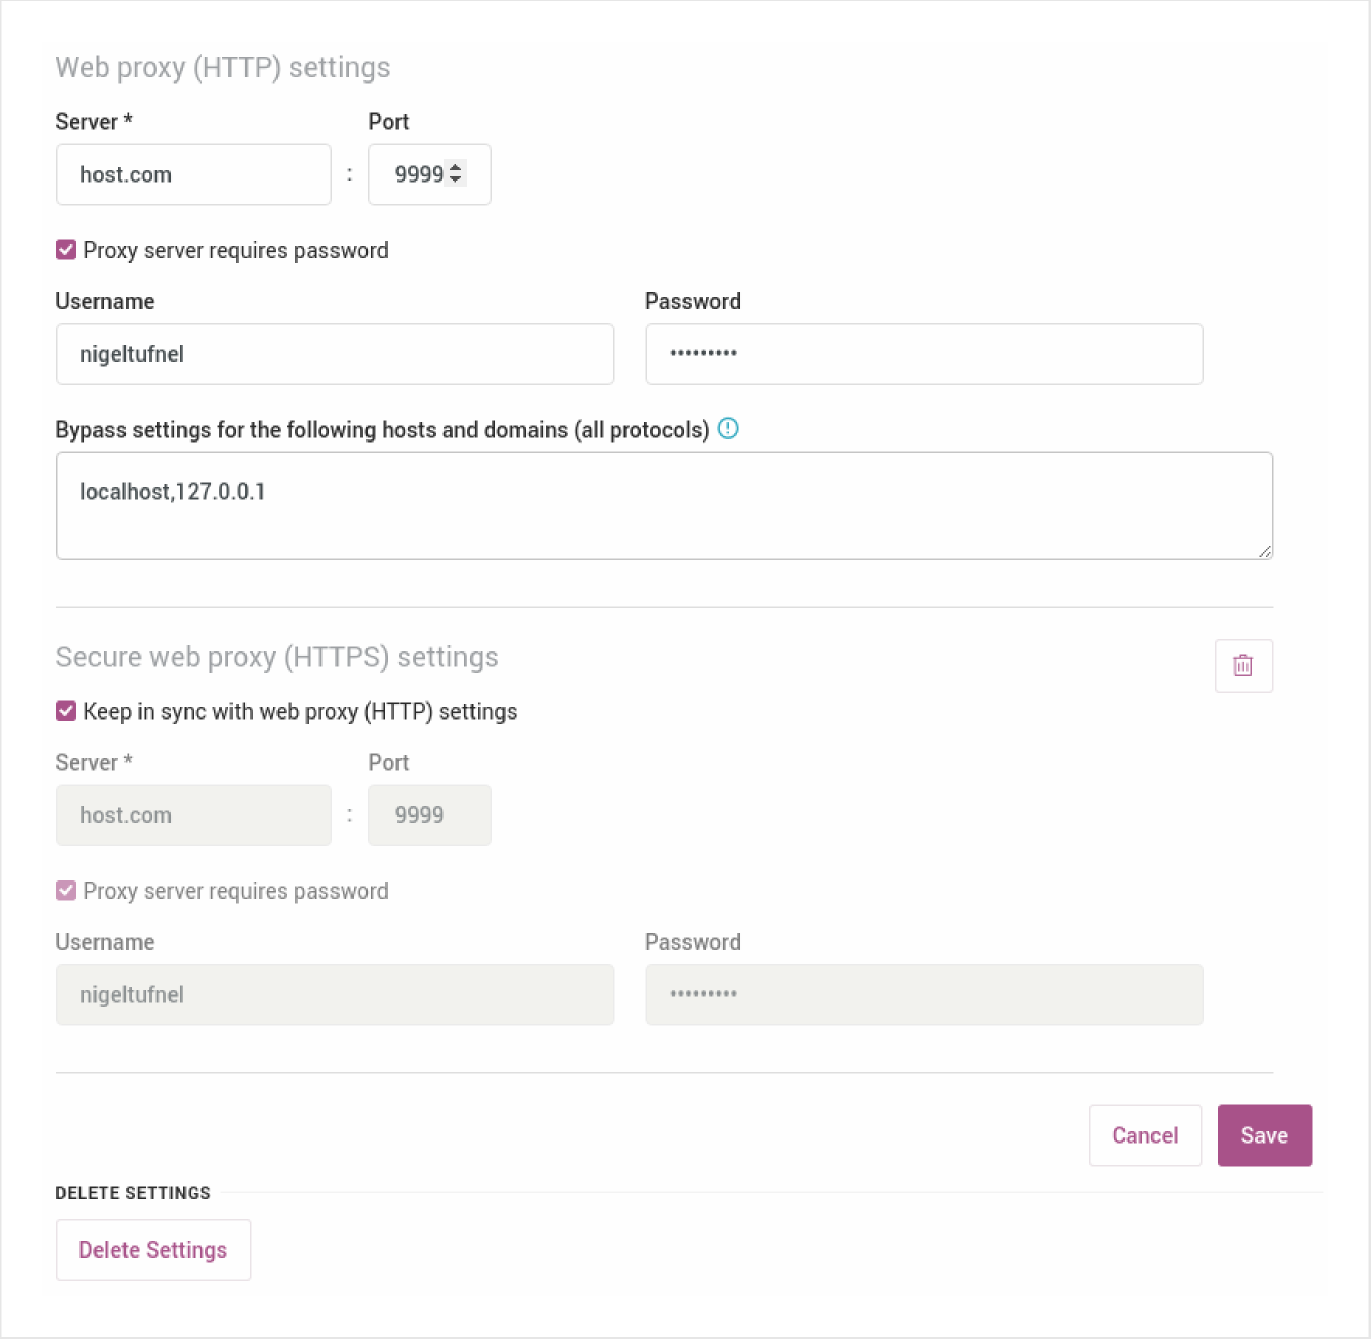

Web proxy (HTTP) settings

Server *: host.com
Port: 9999

☑ Proxy server requires password

Username: nigeltufnel
Password: •••••••••

Bypass settings for the following hosts and domains (all protocols): localhost,127.0.0.1

Secure web proxy (HTTPS) settings

☑ Keep in sync with web proxy (HTTP) settings

Server *: host.com
Port: 9999

☑ Proxy server requires password

Username: nigeltufnel
Password: •••••••••

DELETE SETTINGS

Proxy update

When proxy settings are modified or updated, 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 `supervisord` processes for the changes to become effective.
To do so, run the following command(s):

```
$ sudo service supervisord restart
```

or:

```
$ systemctl restart supervisord
```

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 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**: valid password credentials to authenticate and receive authorization to access the resource(s). For example, *s3cr3tp@SSw0rd_*..
  - **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.

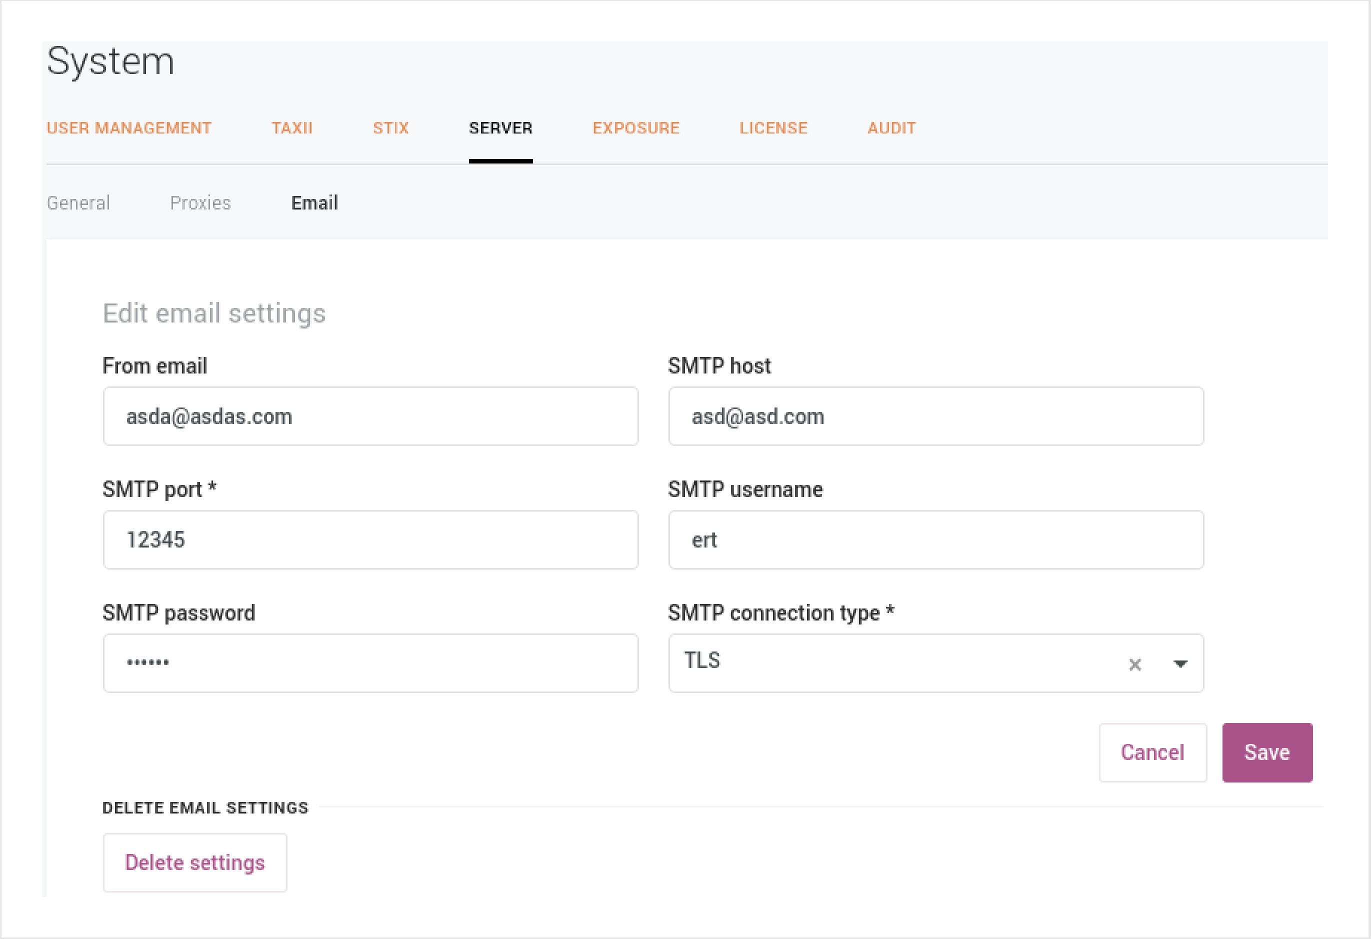

From email: asda@asdas.com

SMTP host: asd@asd.com

SMTP port *: 12345

SMTP username: ert

SMTP password: ••••••

SMTP connection type *: TLS

STIX

- On the left-hand navigation sidebar, click **System**.
- Click **STIX**, and then **Edit settings**.
- Under **Edit STIX settings**, define the following 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.

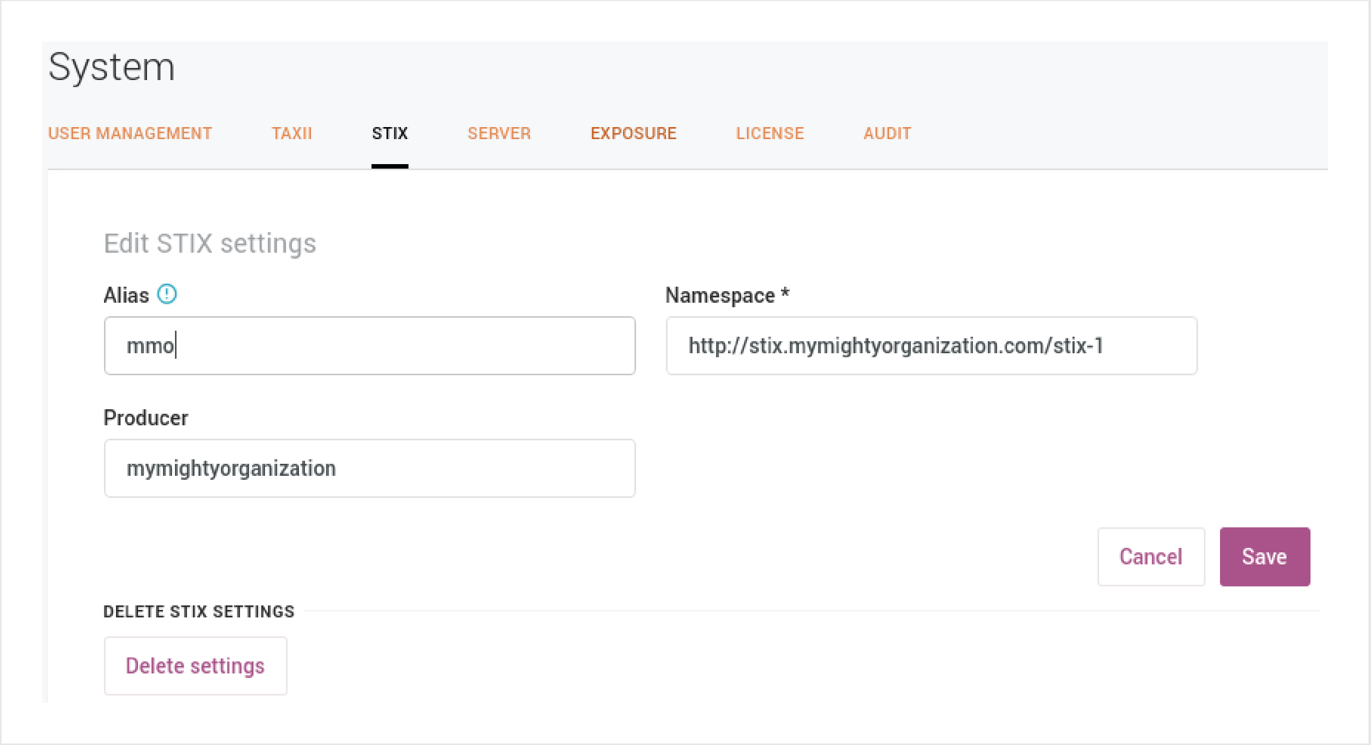

Edit STIX settings

Alias: mmo

Namespace *: http://stix.mymightyorganization.com/stix-1

Producer: mymightyorganization

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.

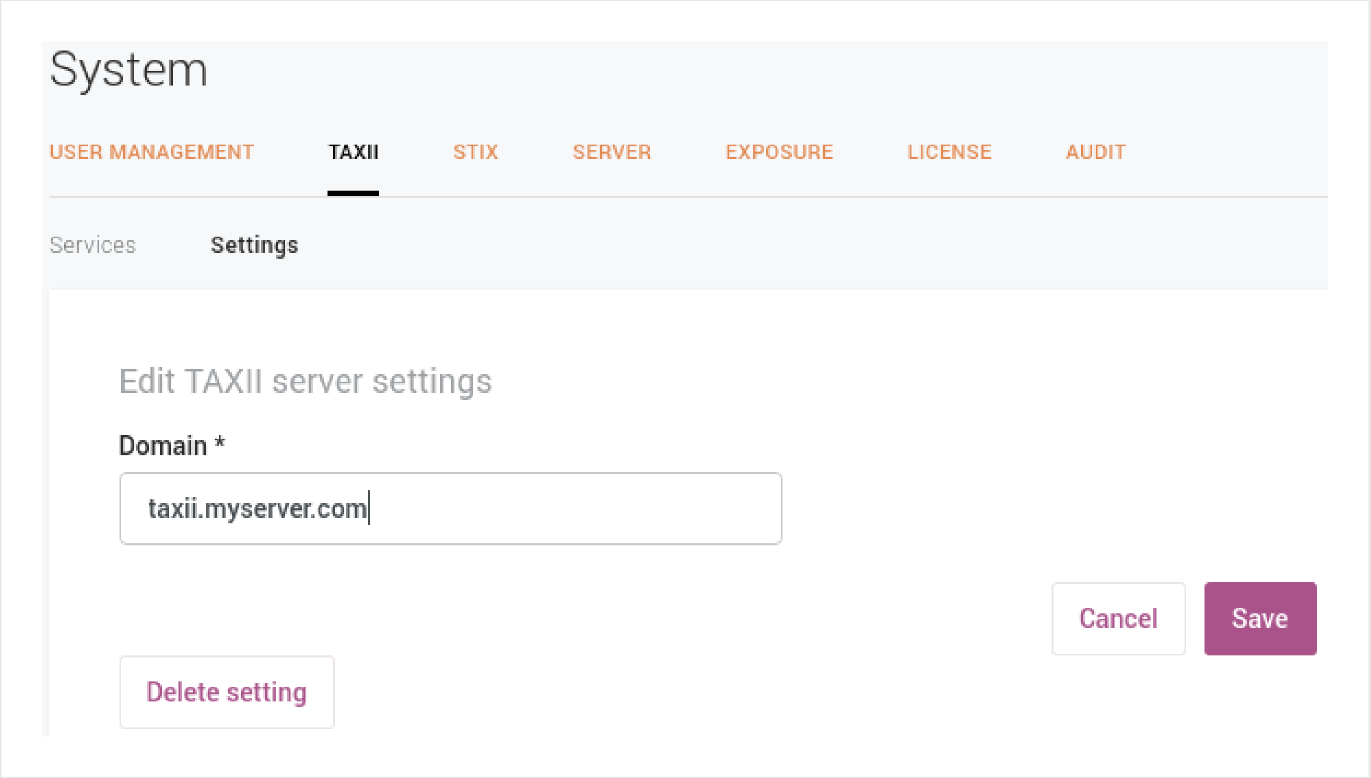

Domain *

taxii.myserver.com

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 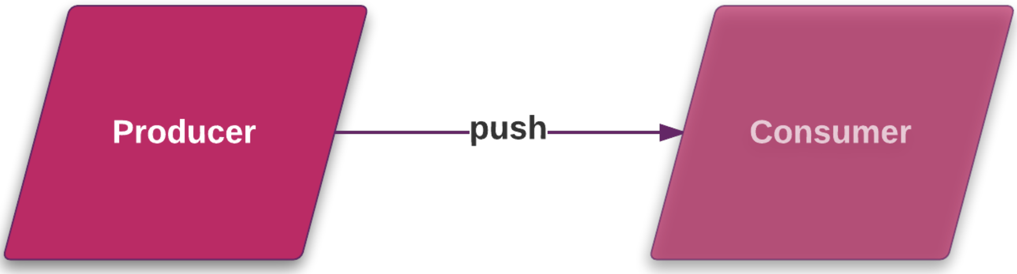

Service type	**Description**
Polling service	The TAXII poll service allows TAXII consumers to request TAXII data collection content from a TAXII producer, usually through TAXII outgoing feeds. Binding: TAXII HTTP 1.0, TAXII HTTPS 1.0 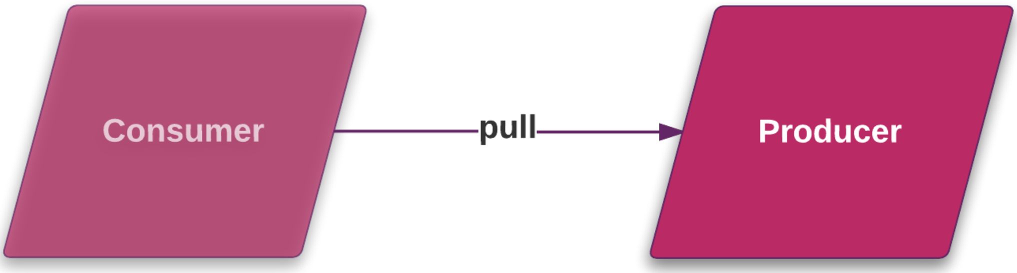

TAXII data collections (TAXII data feeds and TAXII datasets) are a typical example of inbox and polling service content.

Add a TAXII service

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

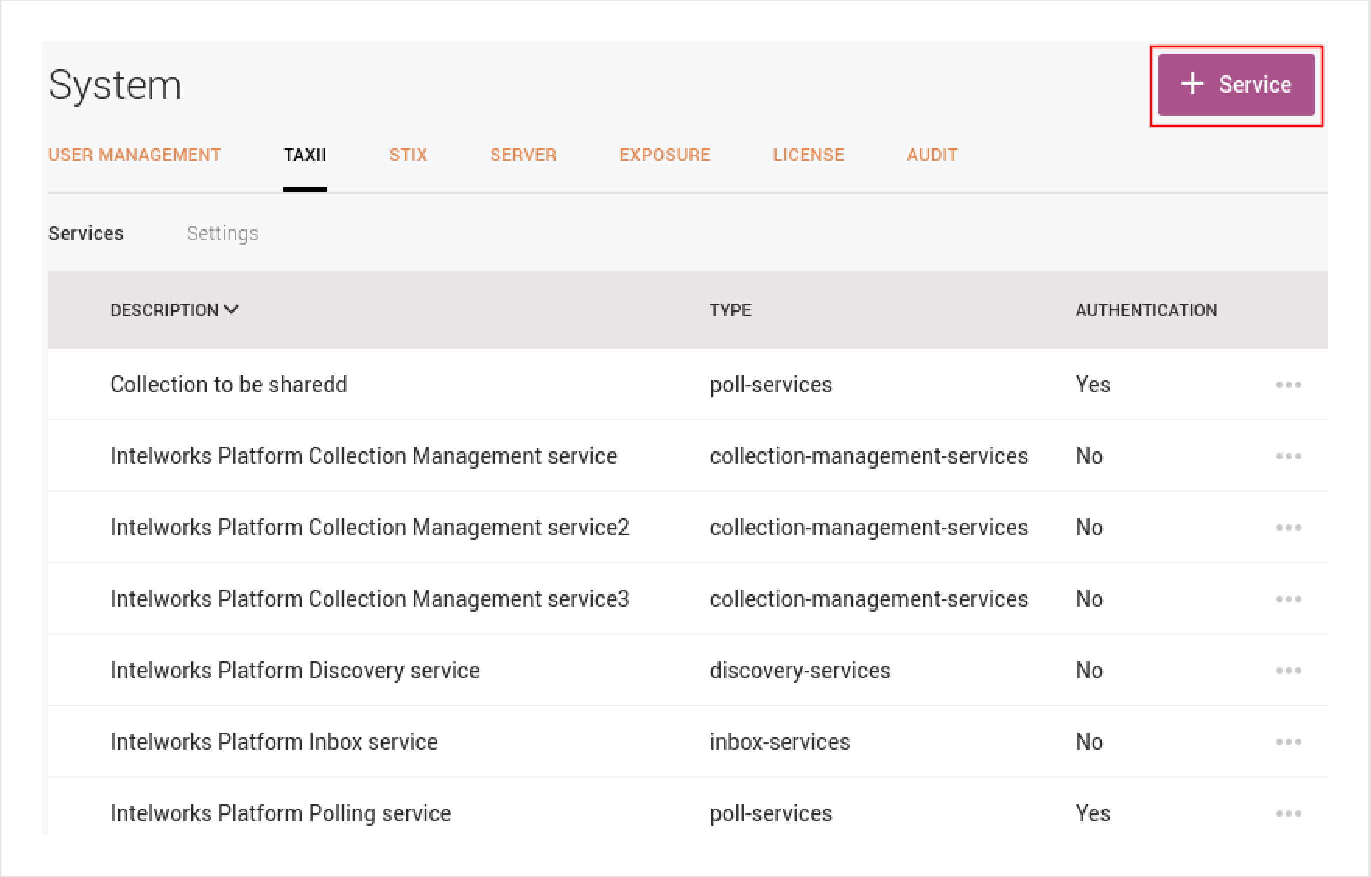

DESCRIPTION	TYPE	AUTHENTICATION
Collection to be sharedd	poll-services	Yes
Intelworks Platform Collection Management service	collection-management-services	No
Intelworks Platform Collection Management service2	collection-management-services	No
Intelworks Platform Collection Management service3	collection-management-services	No
Intelworks Platform Discovery service	discovery-services	No
Intelworks Platform Inbox service	inbox-services	No
Intelworks Platform Polling service	poll-services	Yes

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

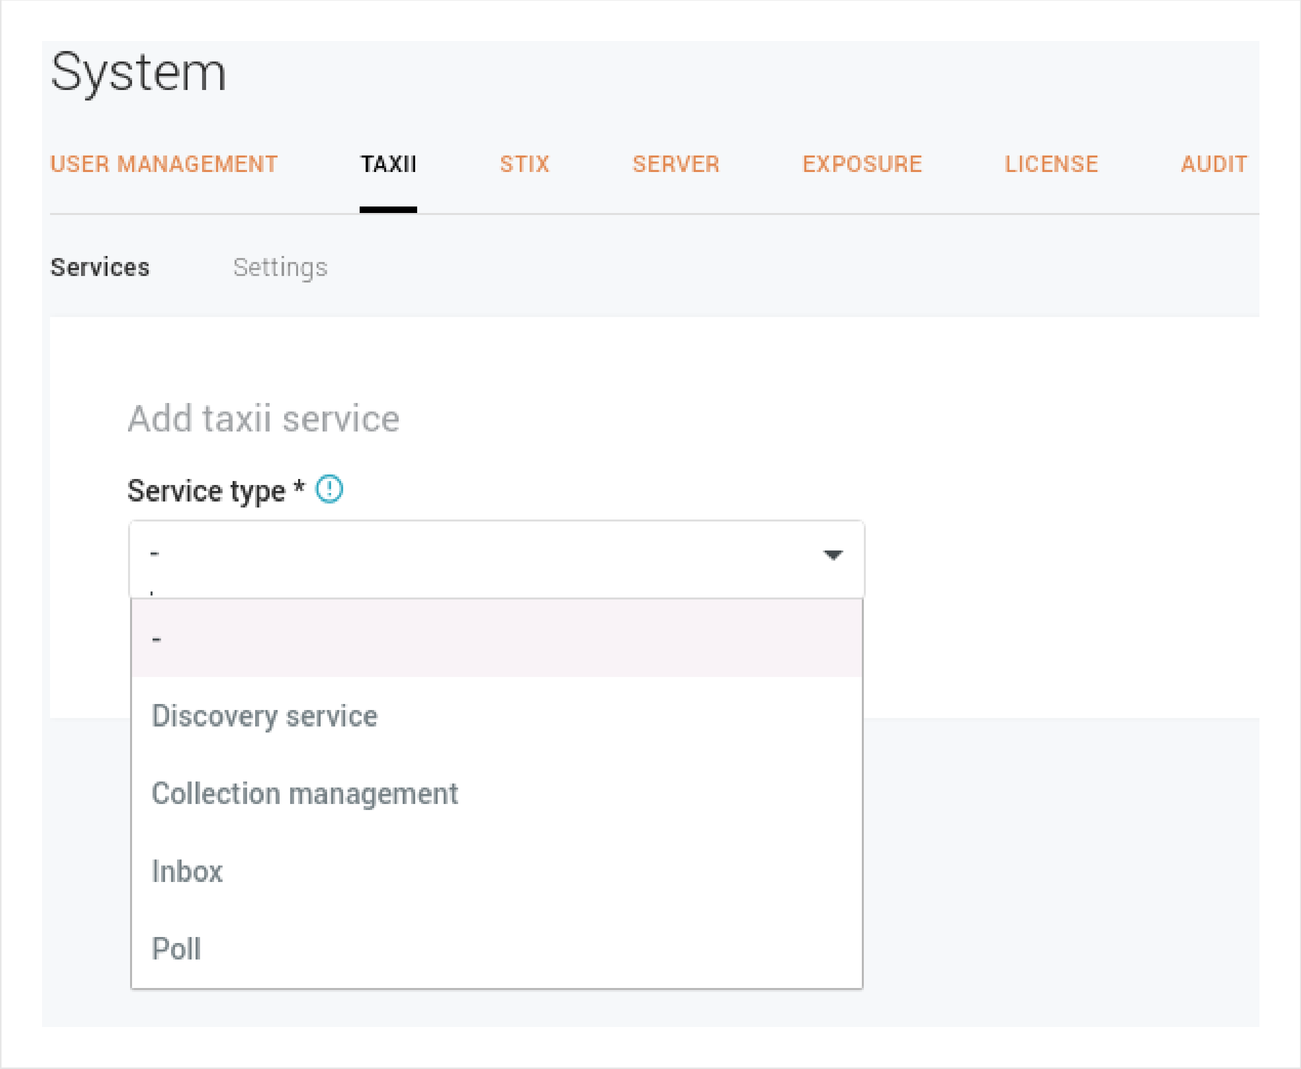

- Fill out the required fields:
  - **Description**: a free-text description of the service. It should be descriptive and easy to remember. Example: *Polling from XYZ*.
  - **Address**: the public endpoint the service can be reached at. Example: */taxii/services/polling*.
  - **Protocol bindings**: the 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 selection 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 selection 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 table view by column. 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 dotted menu icon on the row corresponding to the service.

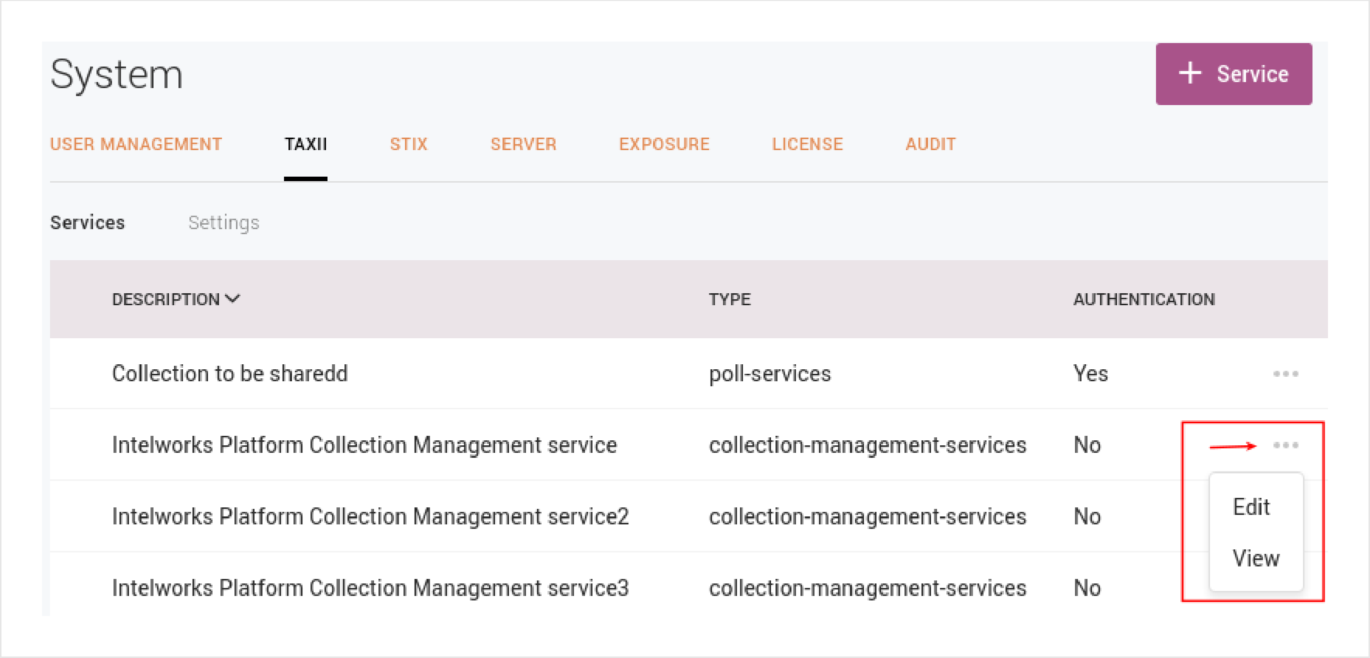

DESCRIPTION	TYPE	AUTHENTICATION
Collection to be sharedd	poll-services	Yes
Intelworks Platform Collection Management service	collection-management-services	No
Intelworks Platform Collection Management service2	collection-management-services	No
Intelworks Platform Collection Management service3	collection-management-services	No

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**.

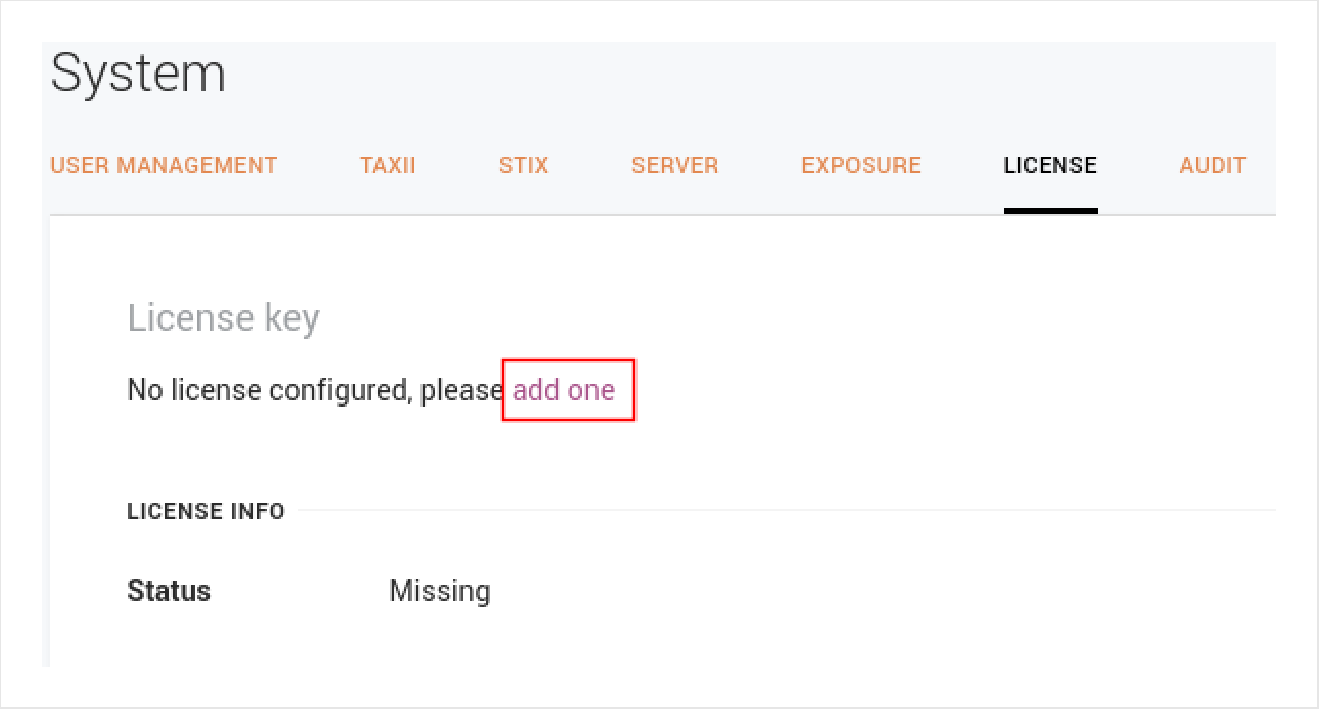

No license configured, please add one

Status Missing

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

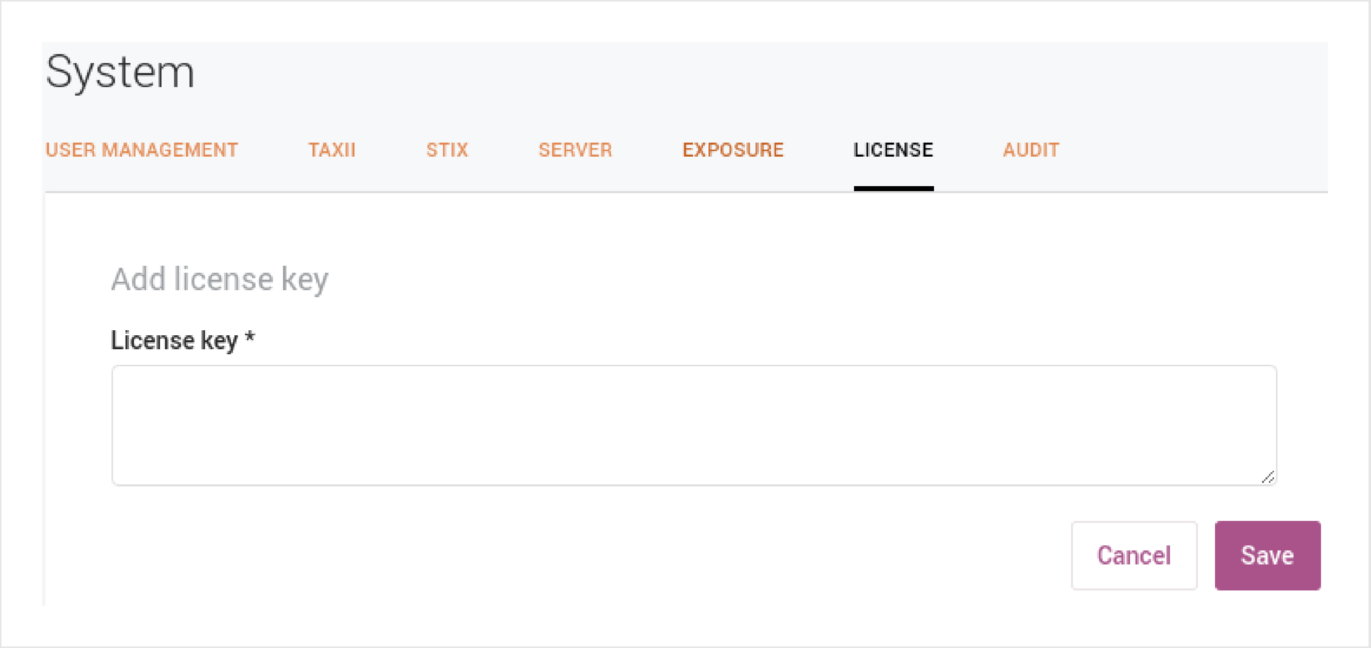

Valid license information populates the license information section:

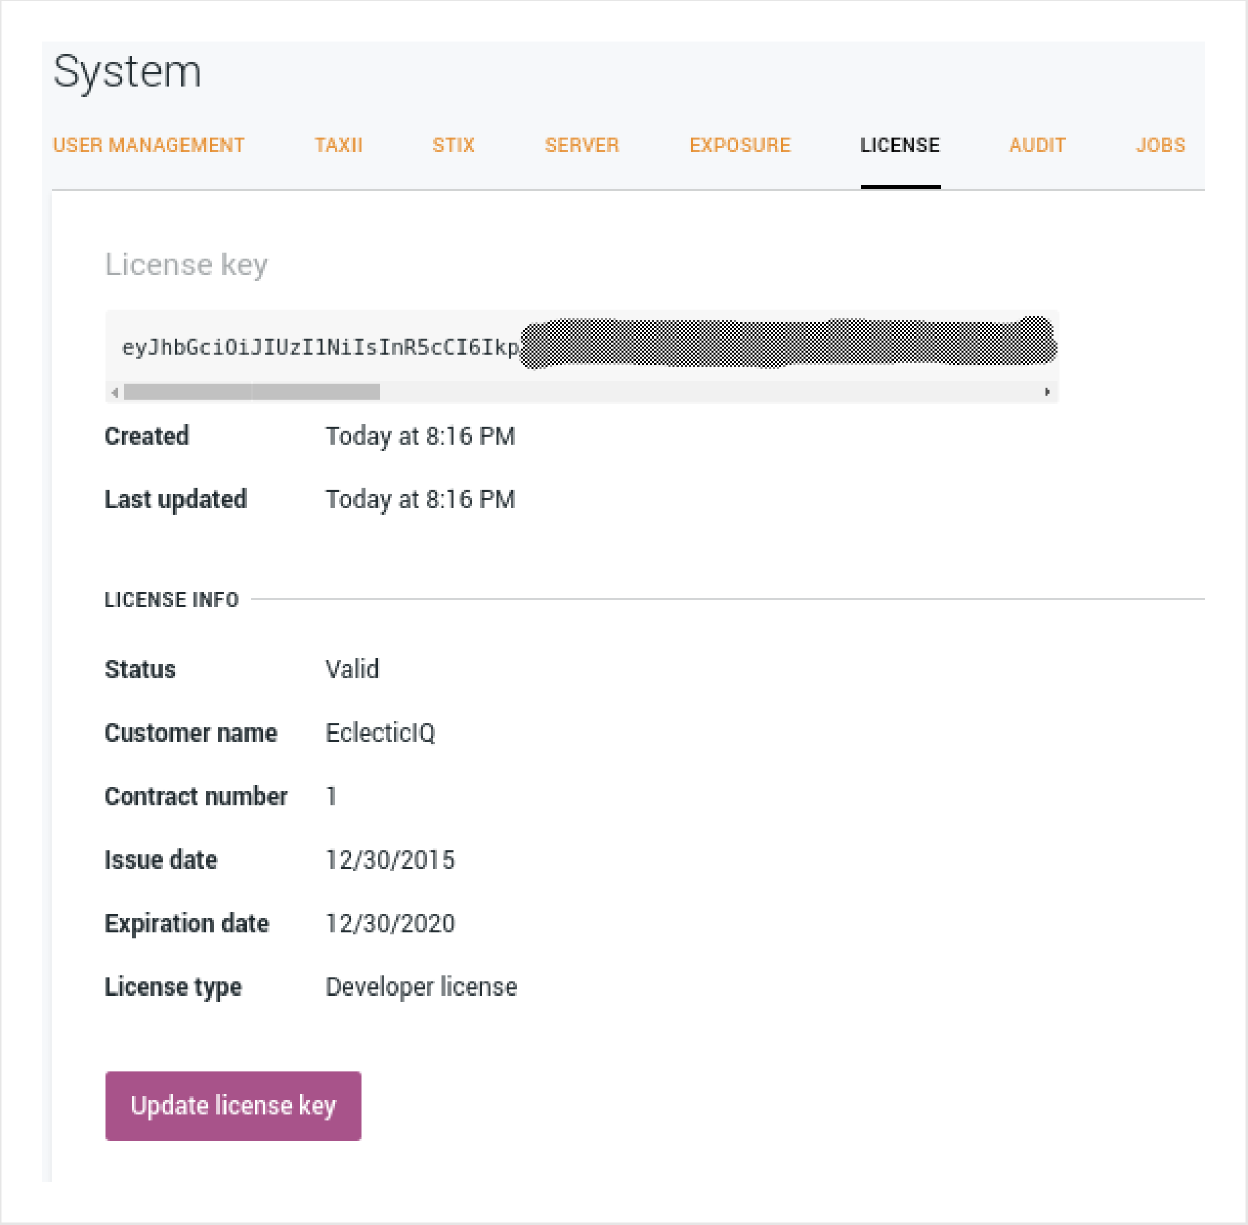

License key

`eyJhbGciOiJIUzI1NiIsInR5cCI6Ikp`[redacted]

Created	Today at 8:16 PM
Last updated	Today at 8:16 PM

LICENSE INFO

Status	Valid
Customer name	EclecticIQ
Contract number	1
Issue date	12/30/2015
Expiration date	12/30/2020
License type	Developer license

The license type is displayed on the status bar:

Developer license

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

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.
- **Extract 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 the selection you specify here.
  This filter furter limits the scope of the exposure configuration by watching only entities with at least one observable type matching your selection(s) under **Extract types**. Moreover, the maliciousness confidence level of the matching observable type(s) needs to correspond to least one of the values specified here.
  Confidence corresponds to the value you set under **Rules > Extract > ✚ 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 the 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.
- **Only active entities**: if you select this checkbox, only published entities are tracked to assess their exposure.
  Draft entities are excluded.

- **Show enrichment extracts**: if you select this checkbox, enrichment observables are included and displayed, when available.
- When you are done, 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**.
- Select the **Audit** tab.
- If audit logging is enabled, and if the audit log file is populated, the matching audit log records are returned. You can sort the table view by column. 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 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, 5xx.

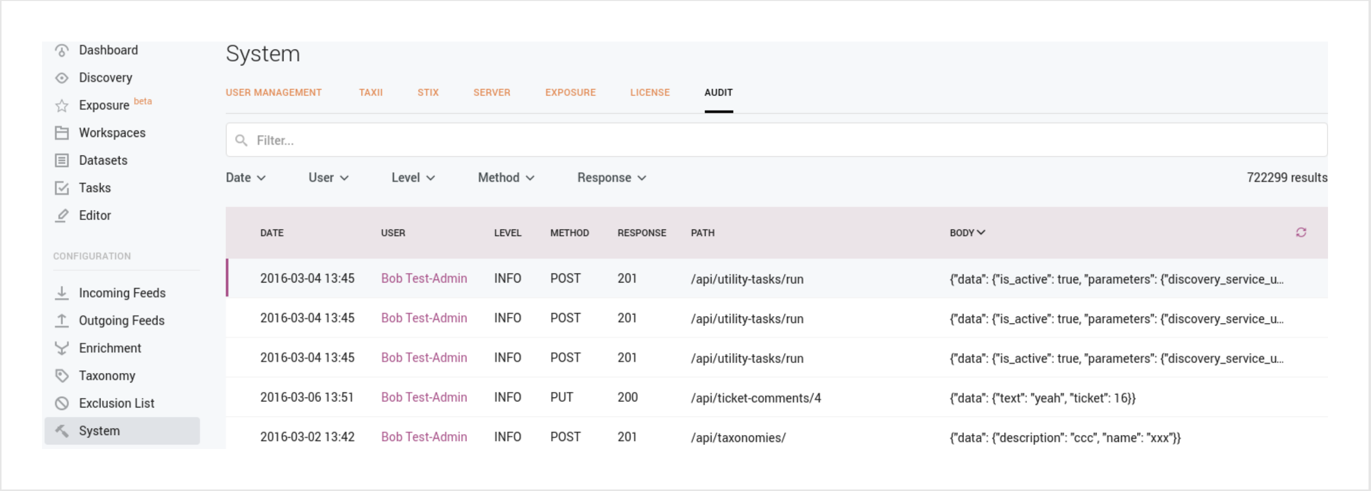

DATE	USER	LEVEL	METHOD	RESPONSE	PATH	BODY
2016-03-04 13:45	Bob Test-Admin	INFO	POST	201	/api/utility-tasks/run	{"data": {"is_active": true, "parameters": {"discovery_service_u...
2016-03-04 13:45	Bob Test-Admin	INFO	POST	201	/api/utility-tasks/run	{"data": {"is_active": true, "parameters": {"discovery_service_u...
2016-03-04 13:45	Bob Test-Admin	INFO	POST	201	/api/utility-tasks/run	{"data": {"is_active": true, "parameters": {"discovery_service_u...
2016-03-06 13:51	Bob Test-Admin	INFO	PUT	200	/api/ticket-comments/4	{"data": {"text": "yeah", "ticket": 16}}
2016-03-02 13:42	Bob Test-Admin	INFO	POST	201	/api/taxonomies/	{"data": {"description": "ccc", "name": "xxx"}}

Jobs

To display an overview of the platform jobs, do the following:

- On the left-hand navigation sidebar click **System**.
- Select the **Jobs** tab.
- By default, the job overview table displays only currently running jobs.
  You can sort the table view by column. 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 selecting/deselecting the job status checkboxes under **Status** to narrow down or to include more results in the overview:

Filter	Description
Running	Displays currently running, i.e. active and not yet completed, jobs.

Filter	Description
Success	Displays successfully completed jobs.
Failure	Displays failed jobs, i.e. jobs that failed to successfully complete because one or more errors occurred.
Revoked	Displays revoked jobs, i.e. jobs that were manually terminated.

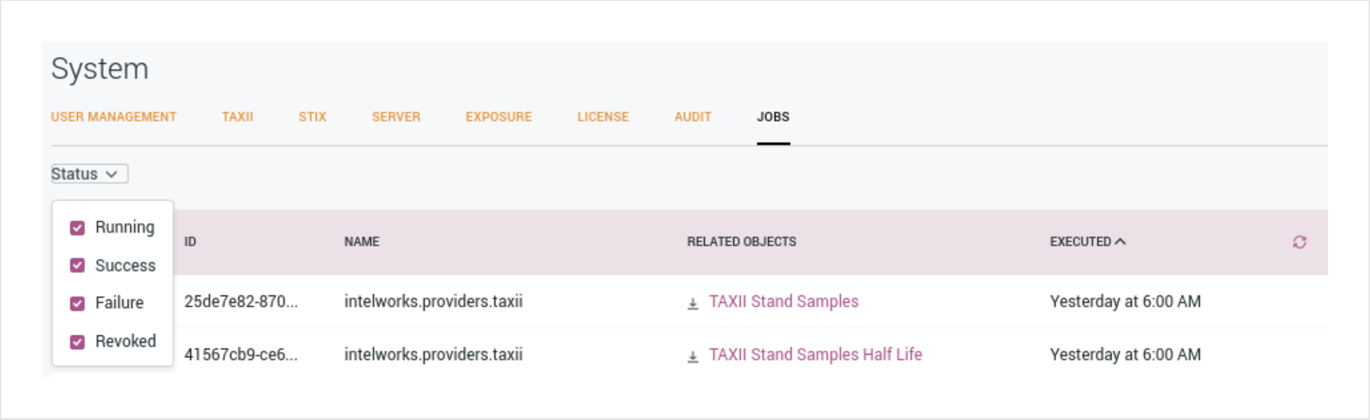

☑ Running
☑ Success
☑ Failure
☑ Revoked

ID	NAME	RELATED OBJECTS	EXECUTED
25de7e82-870...	intelworks.providers.taxii	TAXII Stand Samples	Yesterday at 6:00 AM
41567cb9-ce6...	intelworks.providers.taxii	TAXII Stand Samples Half Life	Yesterday at 6:00 AM

Under the **Related objects** column you can view the platform objects affected by a task. In this context, the objects are the channels the platform uses to ingest and publish information: incoming and outgoing feeds, discovery and entity or extract rules. The object names in this column are manually assigned to the feeds and/or rules in question upon creation.
On the job detail pane you can click a related object name to go to the corresponding detail pane, where you can more closely inspect the selected feed or rule.

- To inspect a job more closely, click the row corresponding to the job you want to examine. An overlay slides in from the side of the screen.
- The job detail pane shows job details like the job/task name, any related platform objects like feeds or rules 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 a job

You can terminate a running task in one of the following ways:

- In the table overview displaying the running jobs, click 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:

- In the table overview displaying the running jobs, 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

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

The 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 accounts.
- **Roles**: the expected functions of users. Roles define typical tasks and behaviors related to the functions described.
- **Permissions**: constrain user scope. Permissions delimit scope by defining:
  - *What* the users are allowed to do.
  - *Where* they can carry out the allowed actions, by defining areas in the platform where users can perform the tasks and behaviors that comply with their assigned roles.
- **Groups**: several users with the same roles and permissions.

Configure users

To add a new user, do the following:

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

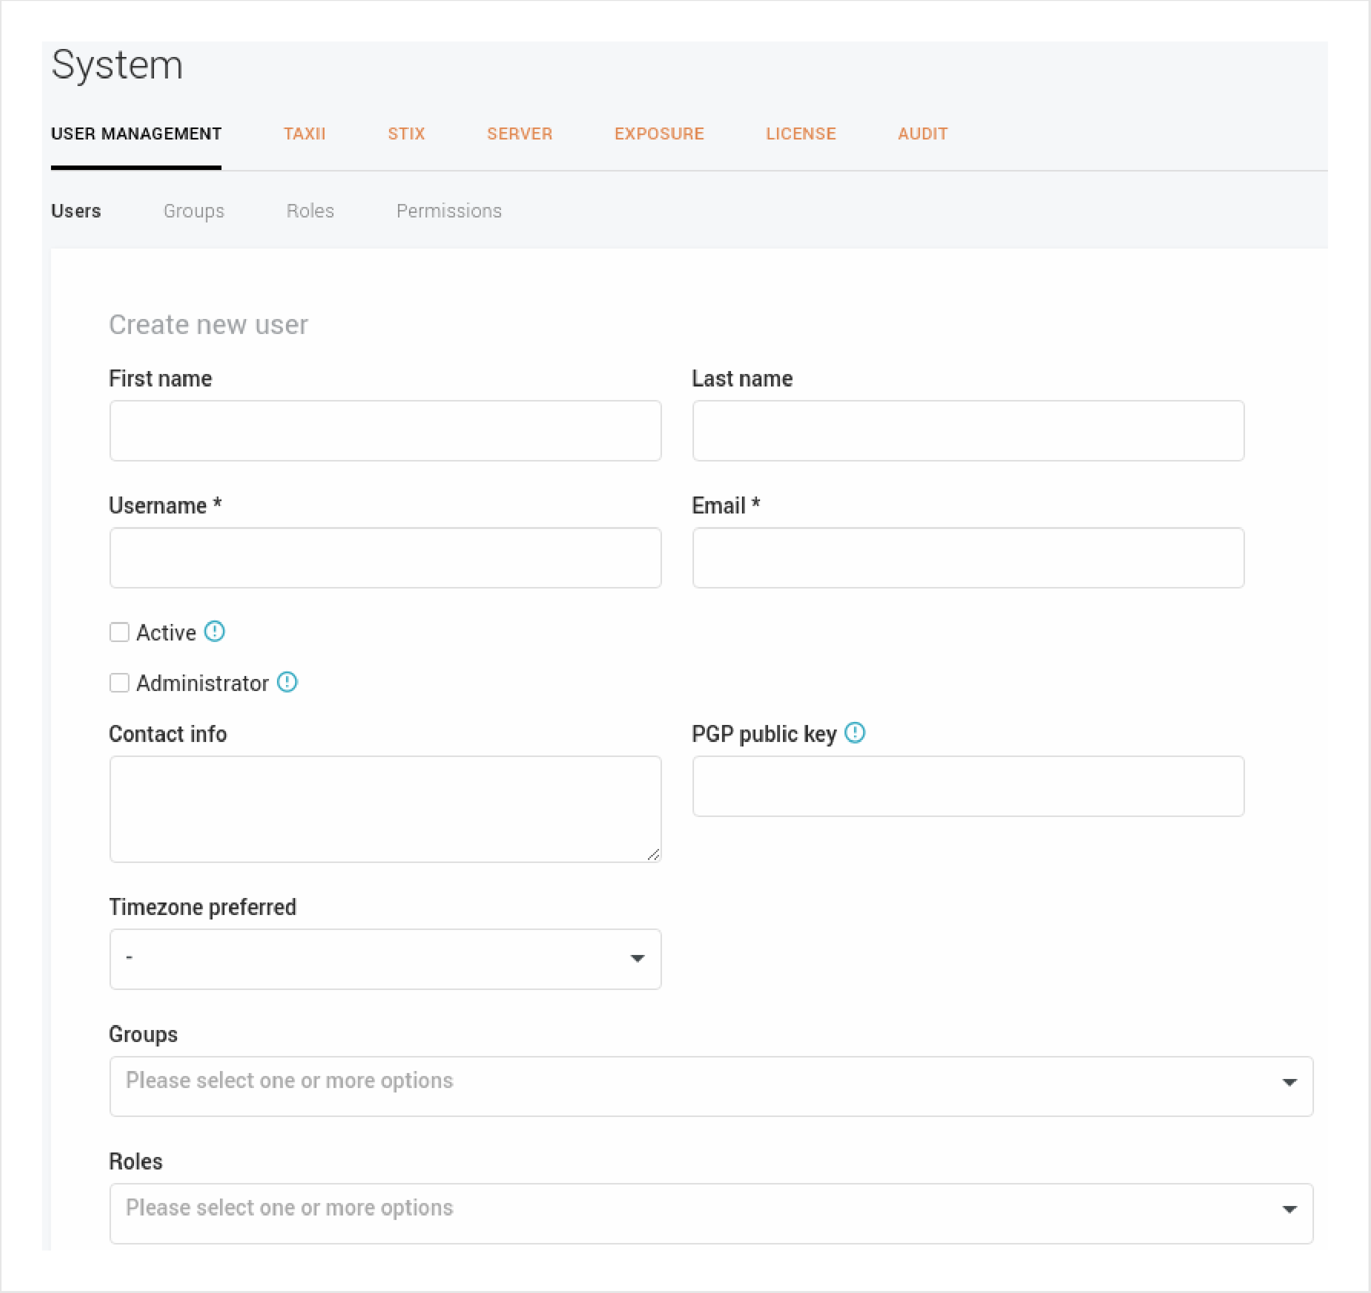

System

Create new user

First name

Last name

Username *

Email *

☐ Active

☐ Administrator

Contact info

PGP public key

Timezone preferred: -

Groups: Please select one or more options

Roles: Please select one or more options

- Under **Create new user**, define the following settings:
  - **First name**: the user's first/given name.
  - **Last name**: the user's last/family name.
  - **Username**: the designated user name to identify the user, when signed in to the platform.
  - **Email**: the user's valid email address.
  - **Active**: select this checkbox to enable/activate the user immediately after saving the newly created user profile.
  - **Administrator**: select this checkbox to elevate the user's role to administrator. When the checkbox is selected, the user has administrator rights and permissions.
  - **Contact info**: the user's contact details, like address or phone number.
  - **PGP public key**: the user's **PGP public key** `(http://www.pgpi.org/doc/pgpintro/#p9)`.
  - **Groups**: this pane lists all available groups the new user can be assigned to.
    - From the drop-down menu select one or more groups to assign the user to.
    - Alternatively, start 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.
    - When you are done, 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

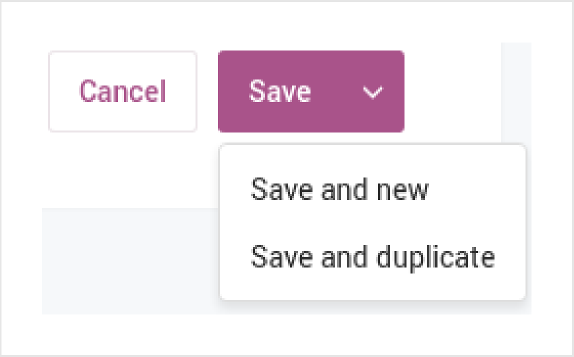

View users

To get an overview of the active platform users, do the following:

- On the left-hand navigation sidebar, click **System**.
- Under **User management > Users**, the user overview is displayed in table format, where each user is assigned a row.

By default, only active users are included in the overview. To view inactive users, do the following:

- From the **Show** drop-down menu, select the **Inactive** checkbox. When selected, this options shows only the inactive platform users.
- To view user details, click an area on a user row.
- An overlay slides in from the side of the screen. It displays detailed user information in a flash-card format. Here you can carry out some actions to modify the user profile:
  - Click **Change Password** to modify the user's password.
  - Click **Edit** to modify the user's profile. The **Edit user** form is identical to the **Create new user** one.

To revoke a user's ability to access the platform, do the following:

- On the left-hand navigation sidebar, click **System**.
- Under **User management > Users**, click an area on a user row.
- On the selected user's detail pane, click **Edit**.
- On the **Edit user** page, deselect the **Active** checkbox.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

To bypass the slide-in user detail pane view and jump directly to the user edit options, click the dotted menu icon on the row corresponding to the user profile you want to modify:

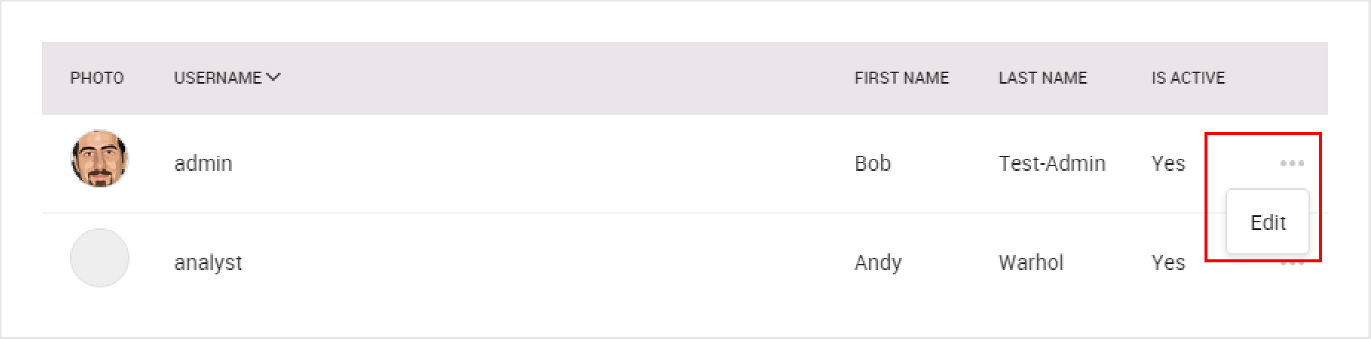

PHOTO	USERNAME	FIRST NAME	LAST NAME	IS ACTIVE	
	admin	Bob	Test-Admin	Yes	•••
	analyst	Andy	Warhol	Yes	•••

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.

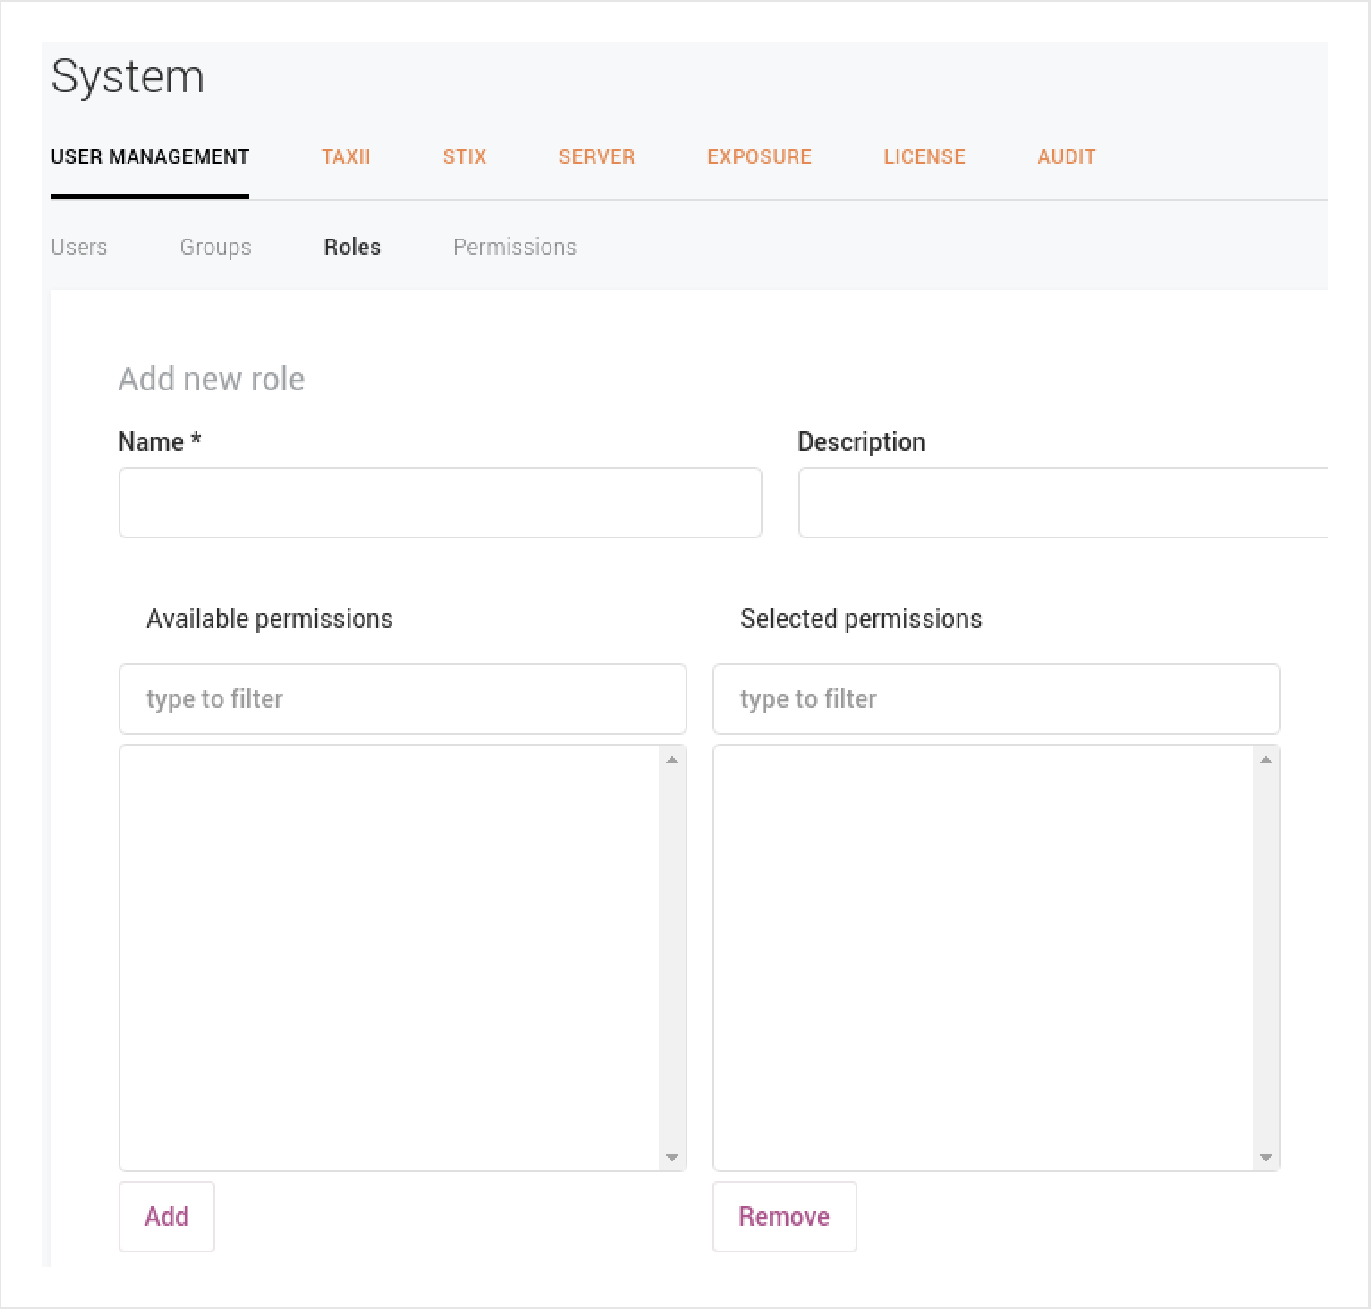

- Under **Add new role**, define the following settings:
  - **Name**: a descriptive name for the role.
  - **Description**: a short description of the role and its purpose.
  - **Available permissions**: this pane lists all available permissions the new role can be granted.
    - Select one or more 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

User permissions are predefined in the platform, and they are not editable or configurable. You can either assign them to user roles, or revoke them.

Permission names try to be as self-explanatory as possible:

- Format: `<type of action> <object of the action>`
- Example: *modify entities*

There are two permission actions:

- **modify**: a modification permission allows write operations.
- **read**: a read permission grants access to the data without allowing any modifications.

To get an overview of the available permissions on the platform, do the following:

- On the left-hand navigation sidebar, click **System**.
- Under **User management > Permissions**, the permission overview is displayed as a table, where each permission is assigned a row.
- To view permission details, click an area on a row.
- An overlay slides in from the side of the screen. It displays permission information in a flash-card format.

Configure groups

To add a new user group, do the following:

- On the left-hand navigation sidebar, click **System**.
- Under **User management**, click **Groups**.
- Under **Groups**, click the ✚ **Group** button.

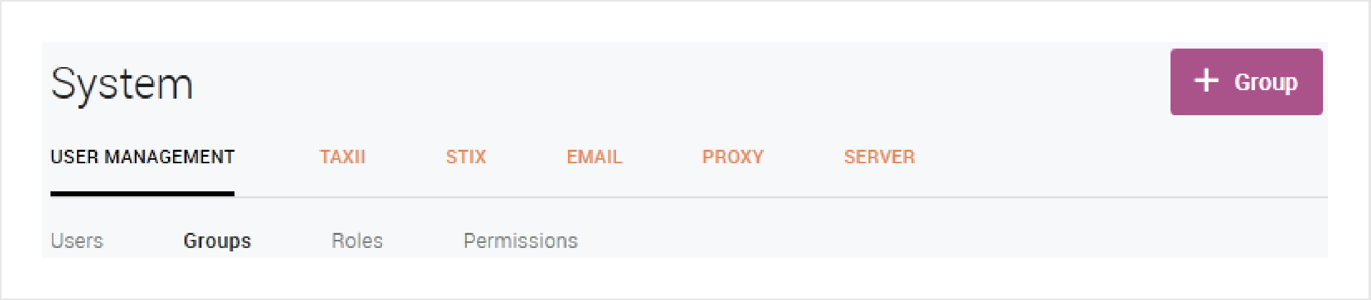

- Under **Add new group**, define the following settings:
  - **Name**: a descriptive name for the group.
  - **Description**: a short description of the group and its purpose.
  - **Allowed sources**: defines the cyber threat intelligence sources the group is allowed to access.
    - Click the ✚ **add** link.
    - From the **Source** drop-down menu, choose a source you want to make available to the group.
    - From the **TLP** drop-down menu, choose a **Traffic Light Protocol** `(https://www.us-cert.gov/tlp)` color to filter the source data accordingly.
    - Click the ✚ **more** link to specify additional sources.
  - 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

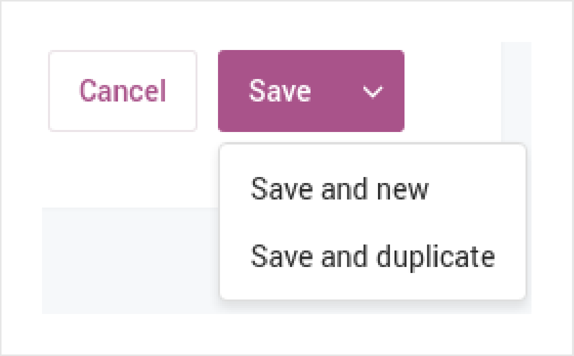

View groups

To get an overview of the platform user groups, do the following:

- On the left-hand navigation sidebar, click **System**.
- Under **System > User Management**, click **Groups**. It shows an overview of the available groups where each group is assigned a row.
- To view group details, click an area on a group row. An overlay slides in from the side of the screen. It displays detailed group information in a flash-card format.
  Here you can carry out some actions to modify the group configuration:
  - Click **Edit** to modify the group configuration. The **Edit Group** form is identical to the **Add new group** one.
  - Click **Delete** to remove the group from the platform. This may impact users belonging to that group.

To bypass the slide-in group detail pane view and jump directly to the group edit or delete options, click the dotted menu icon on the row corresponding to the group you want to modify:

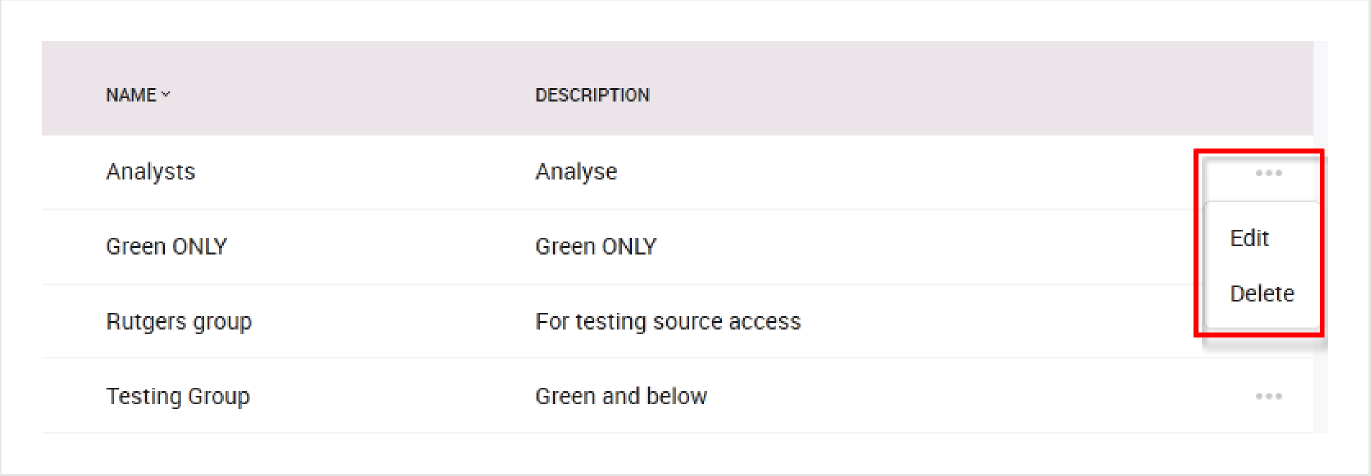

NAME	DESCRIPTION
Analysts	Analyse
Green ONLY	Green ONLY
Rutgers group	For testing source access
Testing Group	Green and below

Incoming feeds

Summary: 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.
Cisco AMP 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.

Content type	Feed type	Description
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.
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.
PDF	in	Standard PDF format, preferably native (not scanned).
STIX 1.0	in	STIX data model **v. 1.0** `(http://stixproject.github.io/data-model/1.0/).`
STIX 1.1	in	STIX data model **v. 1.1** `(http://stixproject.github.io/data-model/1.1/).`
STIX 1.1.1	in	STIX data model **v. 1.1.1** `(http://stixproject.github.io/data-model/1.1.1/).`
STIX 1.2	in	STIX data model **v. 1.2** `(http://stixproject.github.io/data-model/1.2/).`
Text	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.

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 AMP Threat Grid Curated Feed	in	Provides data on compromised IP addresses, domains, hashes, registry keys, and network streams.
Cisco AMP 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.

Transport type	Feed type	Description
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, out	Custom feed ingesting data through HTTP.
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, and actors.
Mount point download	in	Custom feed using a local or a network drive as a data source.
TAXII inbox	in, out	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.

Configure incoming feeds

Summary: 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, click the ✚ *plus* button.
- On the **Create new** sidebar, click **Data management > Incoming feed**.

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button.
- Under **Configuration** on the drop-down menu, click **Data management**, and then **Incoming feeds**.

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*

Field	Type	Description	Example
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*
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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

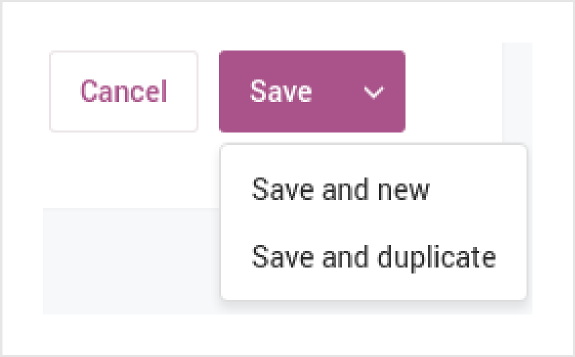
Cancel
Save
Save and new
Save and duplicate

Run feeds

Summary: 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, click the ✚ *plus* button.
  - On the **Create new** sidebar, click **Data management > Incoming feed**.

</ul>

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button.
- Under **Configuration** on the drop-down menu, click **Data management**, and then **Incoming feeds**.
- 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

Summary: 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.

- On the top navigation bar, click the ✚ *plus* button.
- On the **Create new** sidebar, click **Data management > Outgoing feed**.

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button.
- Under **Configuration** on the drop-down menu, click **Data management**, and then **Outgoing feeds**.

The **Outgoing feeds** page displays an overview of the configured outgoing feeds to publish and distribute selected intel from the platform to external parties, services, and systems.

On the top-right corner of the screen, click the **+ Outgoing feed** button. - The **✚ > Data management > Outgoing feeds > Create** form page includes several input fields to help you define *what* you want to share 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 the 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, new data is appended to the existing data.
      When new intelligence is made available through the outgoing feed, it is added after/below the existing data from the same feed.
    - **Replace** every time the outgoing feed task runs, existing data is deleted and it is replaced with new data.
      When new intelligence is made available through the outgoing feed, it overwrites and replaces existing data from the same feed.

Set a schedule

- Under **Execution schedule** you can define how often you want to run the outgoing feed task:
- **None**: no schedule is defined. You need to manually trigger the task to publish data through the outgoing feed.
- **Minute**: the outgoing feed task run automatically.
  You define the execution interval in 5-minute increments from the corresponding drop-down menu.
- **Hour**: the outgoing feed task task run automatically every hour.
  You define how long after the beginning of an hour the task should run from the corresponding drop-down menu.
- **Day**: the outgoing feed task task run automatically once a day.
  You define the time of the day when the task should run from the corresponding drop-down menu.
- **Week**: the outgoing feed task task run 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 outgoing feed task task run 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.
- Select the **Active** checkbox to make the feed available immediately after creating it.

Set a TLP override

- The **Override TLP** options overwrite the **TLP** `(https://www.us-cert.gov/tlp)` color code associated with the outgoing feed entities with the one you set here. The selected TLP value is assigned to all the entities in the outgoing 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** radio buttons allow including in the outgoing 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 level of reliability of the source. It helps analysts assess how much confidence they can reasonably have in an intel source.
  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 outgoing feed data only the entities whose relevancy value is higher than the one defined here.

Set extract filters

- **Allowed extract states**: from the drop-down menu select one or more extract states to include in the outgoing feed data only the entities whose extract states match the selections defined here.
- **Extract types**: from the drop-down menu select one or more extract types to include in the outgoing feed data only the entities whose extracts types match the selections defined here.
- **Enrichment extract types**: from the drop-down menu select one or more enrichment extract types to include in the outgoing feed data only the entities whose enrichment extracts types match the selections defined here.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

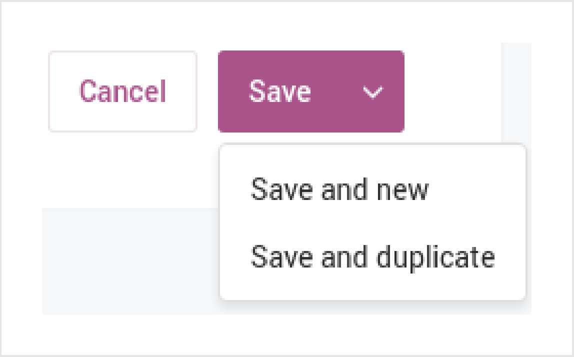

Depending on the selected transport type, you may need to specify a URI:

- 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

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

Enriching entities via extracts

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.

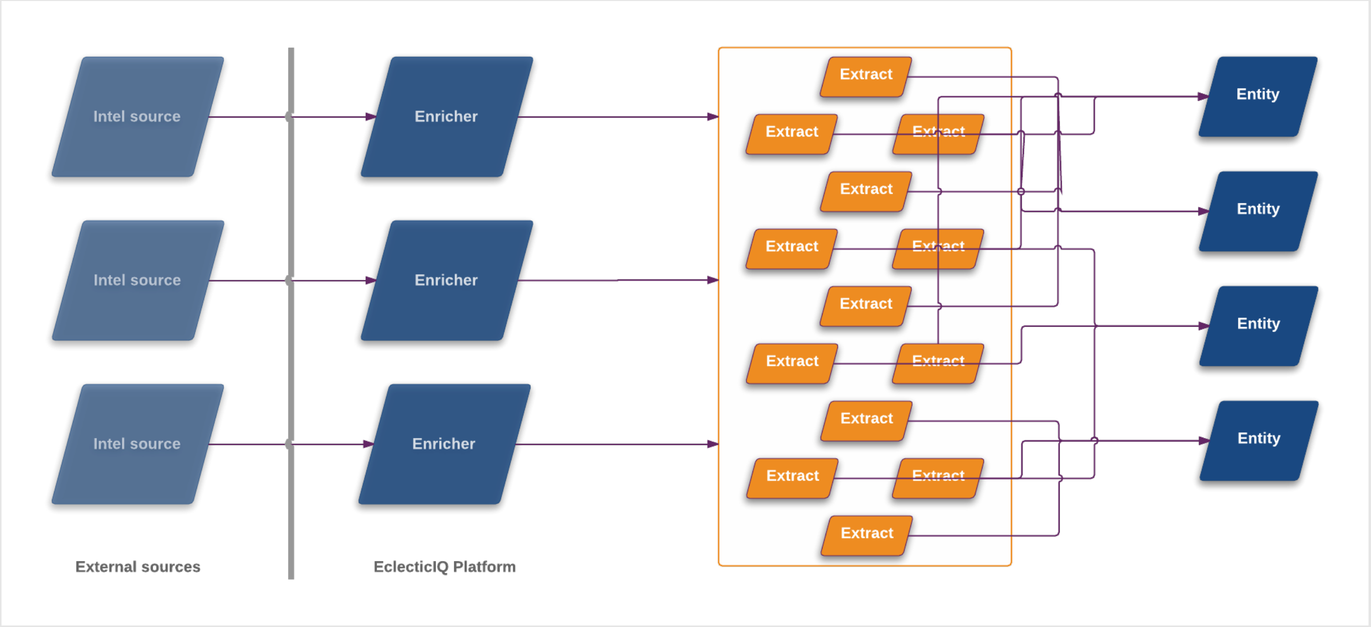

For example, let’s assume a scenario where an analyst is investigating a threat actor entity. The entity includes some observables, and one of them is a DNS domain name.
The analyst looks up the domain name by running it through a whois service. The lookup results include an email address.
During the 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.

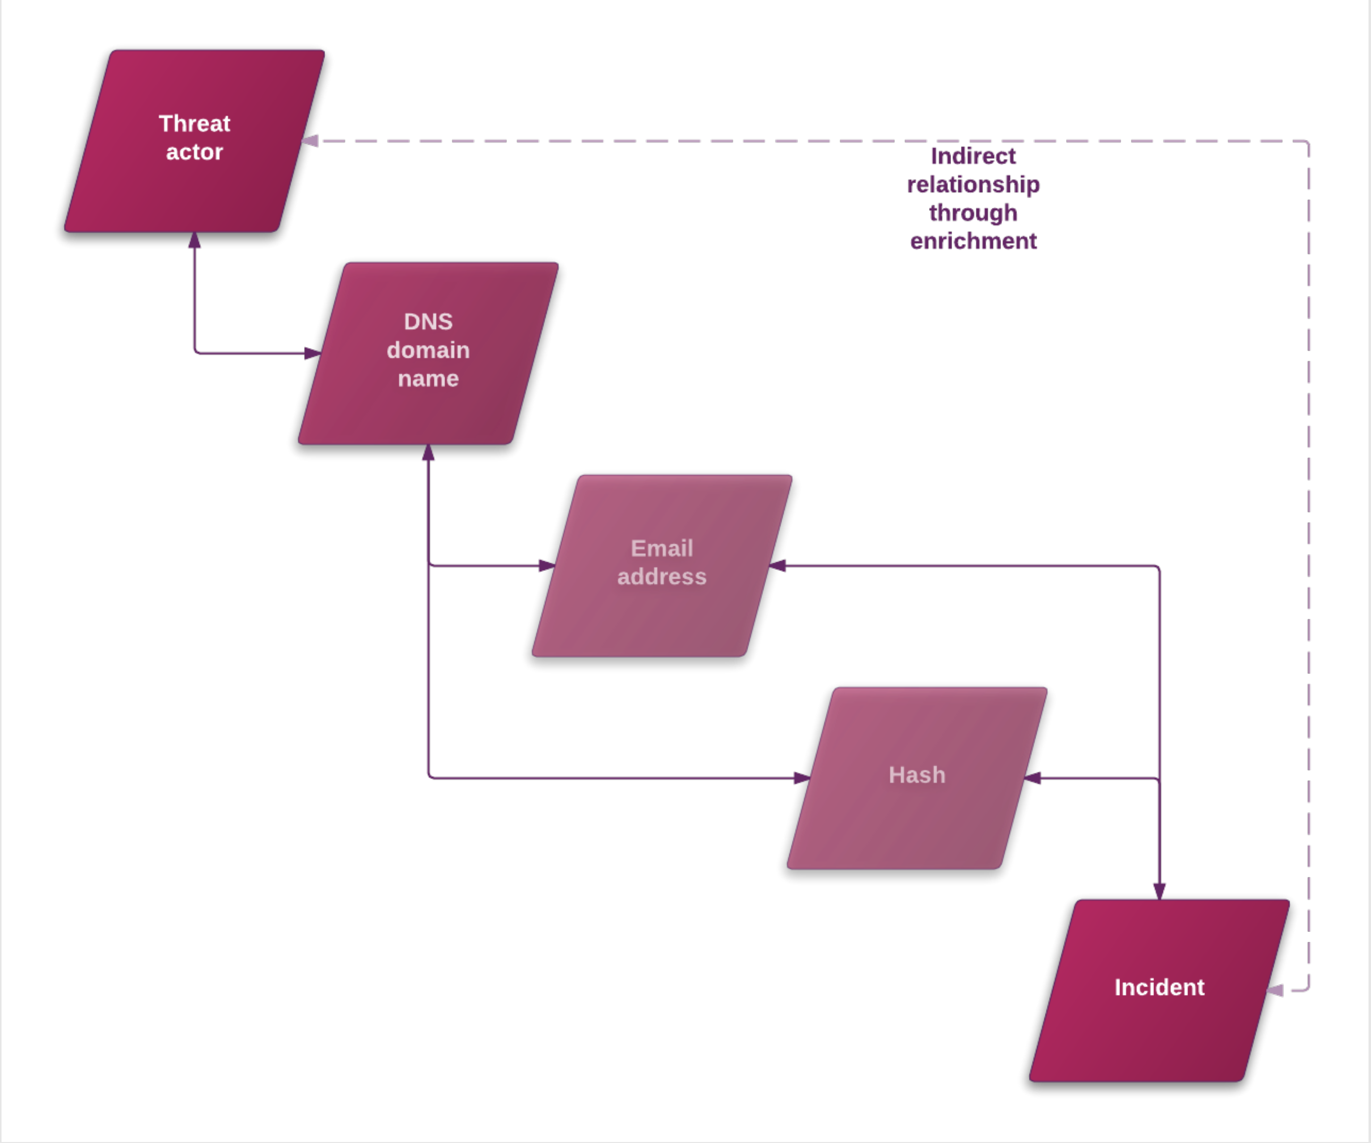

Enrichments augment extracts with raw data information related to entities:

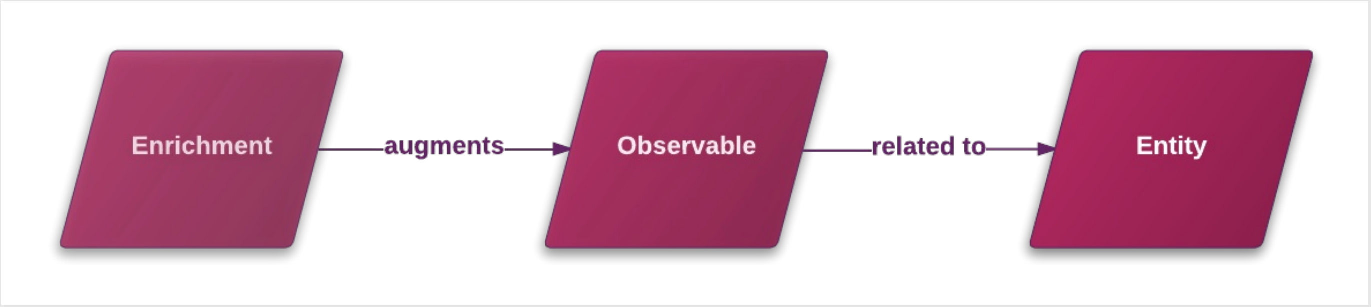

Enrich entities

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 extract types. You can set multiple filters to cover usage scenarios as needed. You can then examine the returned enrichment extract 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 editing 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 dotted menu icon to display the context menu.
- From the drop-down menu select **Edit**.
- At the bottom of the entity editor page, select the **Manually enrich** checkbox.
  A new input field with a drop-down menu becomes available.
- From the drop-down menu select one or more enrichers you want to apply to the entity.

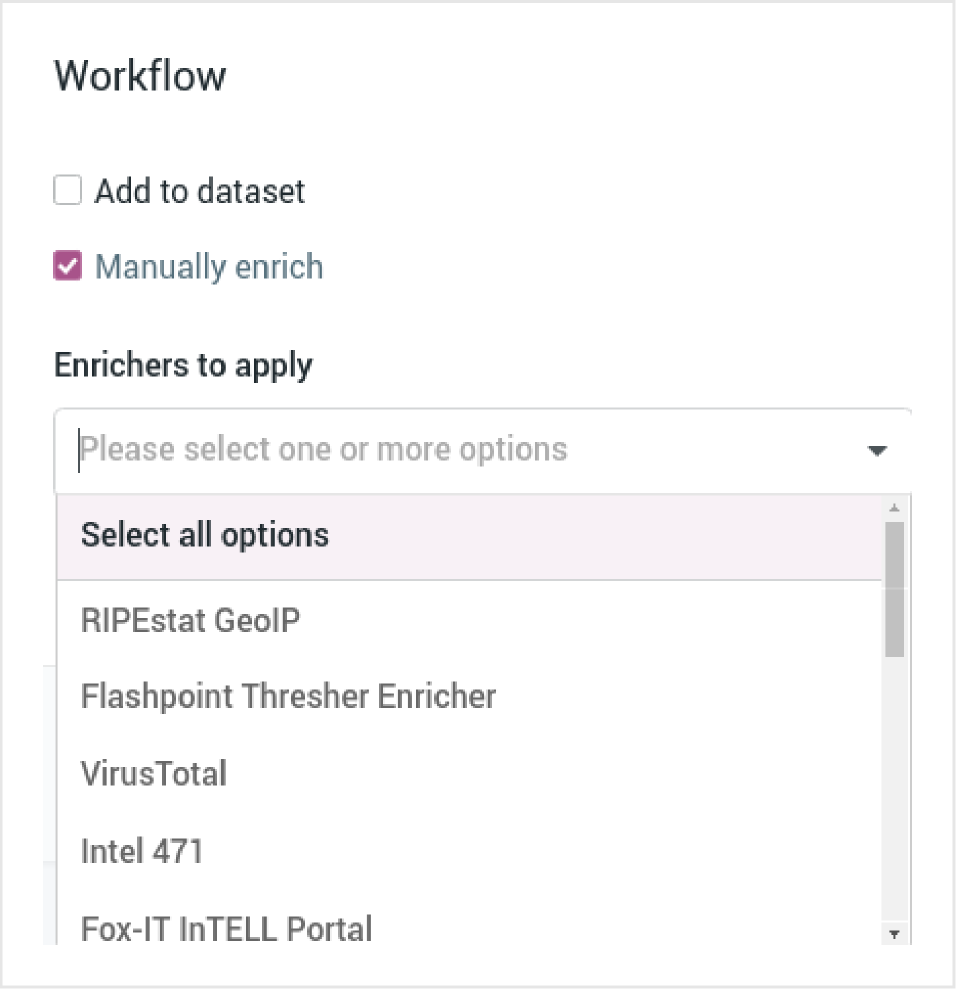

Workflow

☐ Add to dataset

☑ Manually enrich

Enrichers to apply

- When you are done, 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.

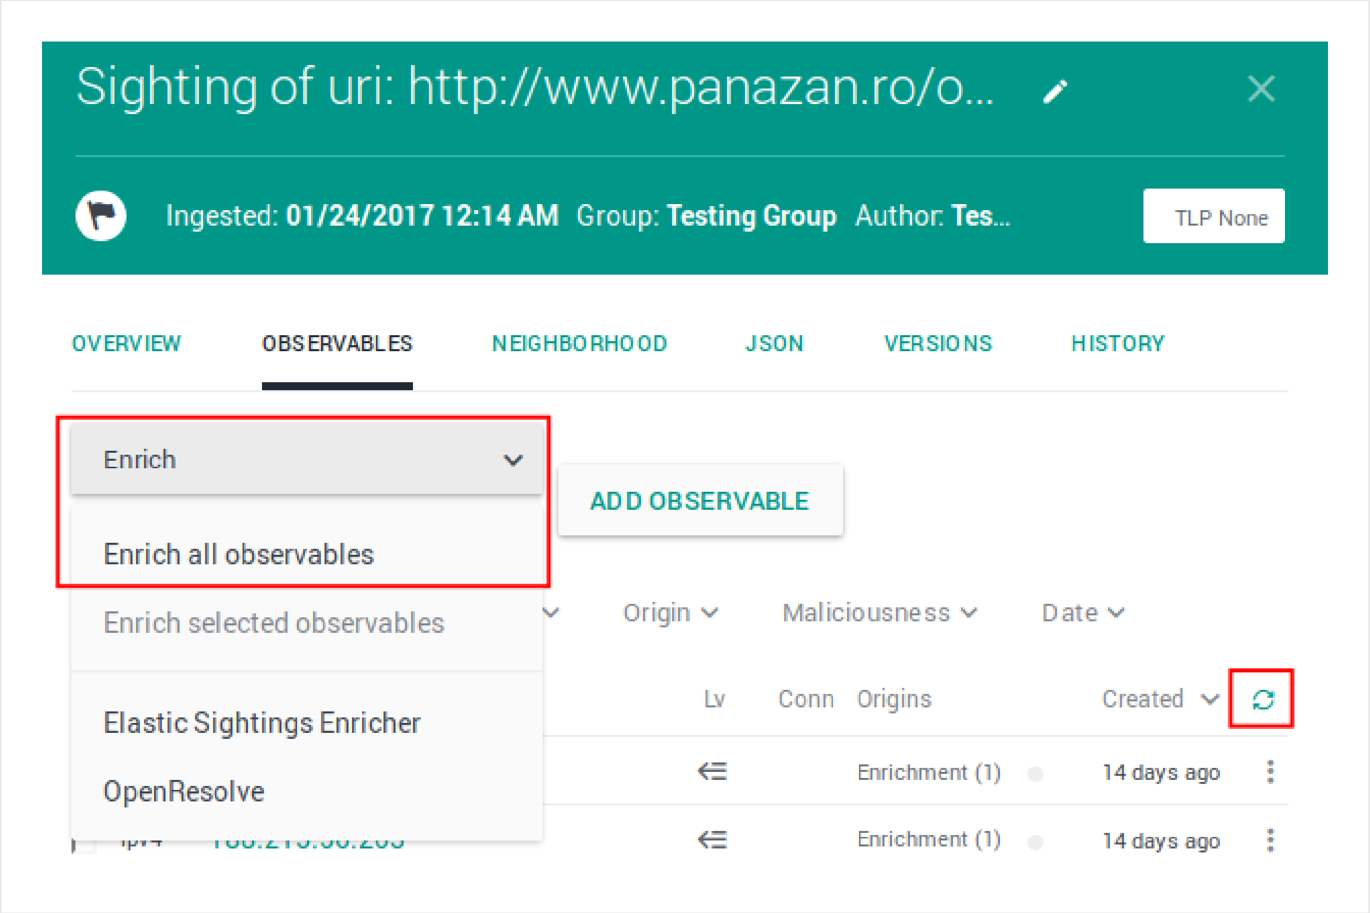

Sighting of uri: http://www.panazan.ro/o...

Ingested: 01/24/2017 12:14 AM Group: Testing Group Author: Tes...

Lv	Conn	Origins	Created
		Enrichment (1)	14 days ago
		Enrichment (1)	14 days ago

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.

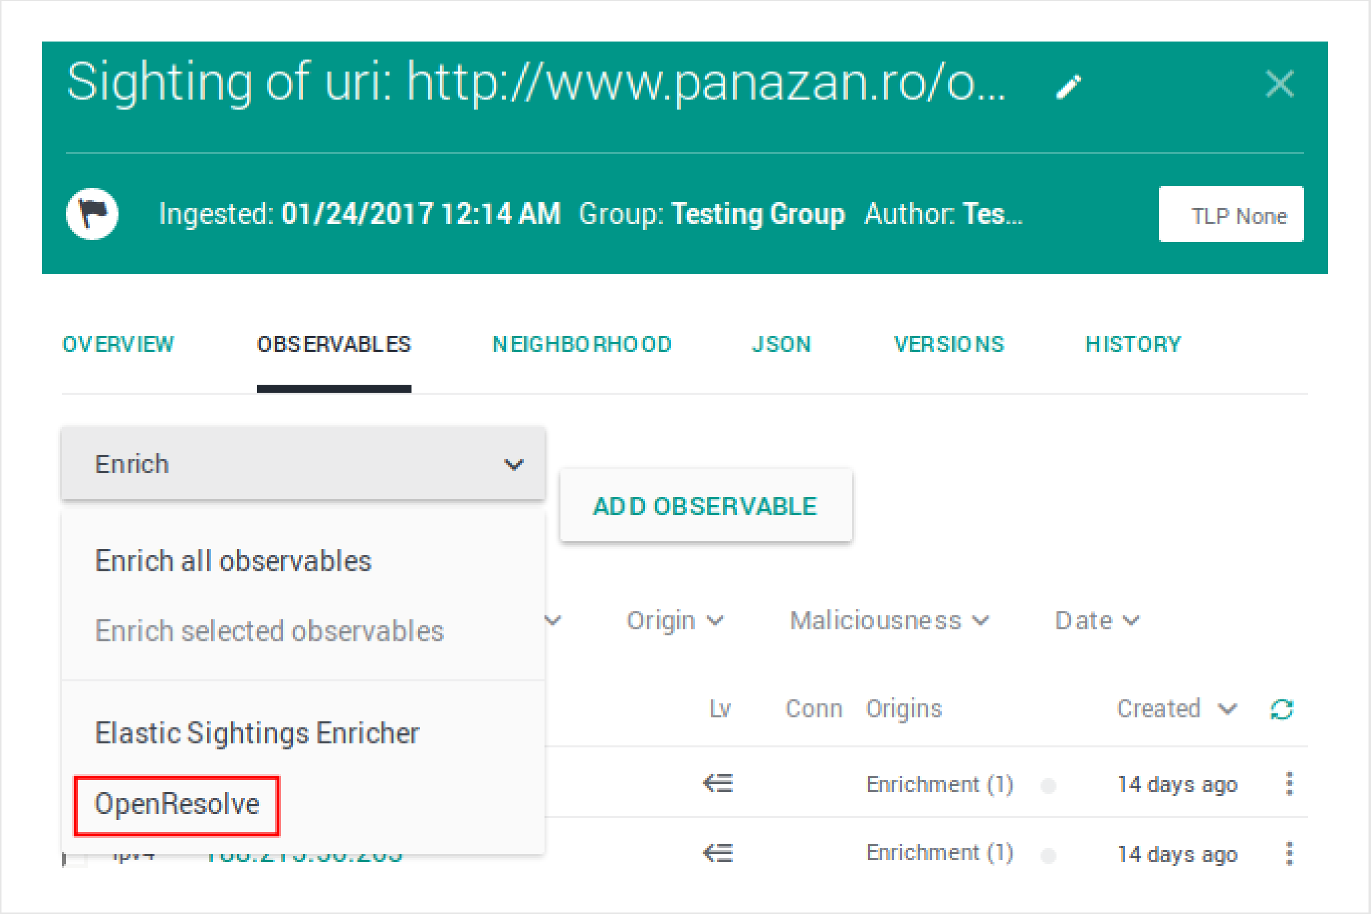

To enrich only specific observables:

- 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.

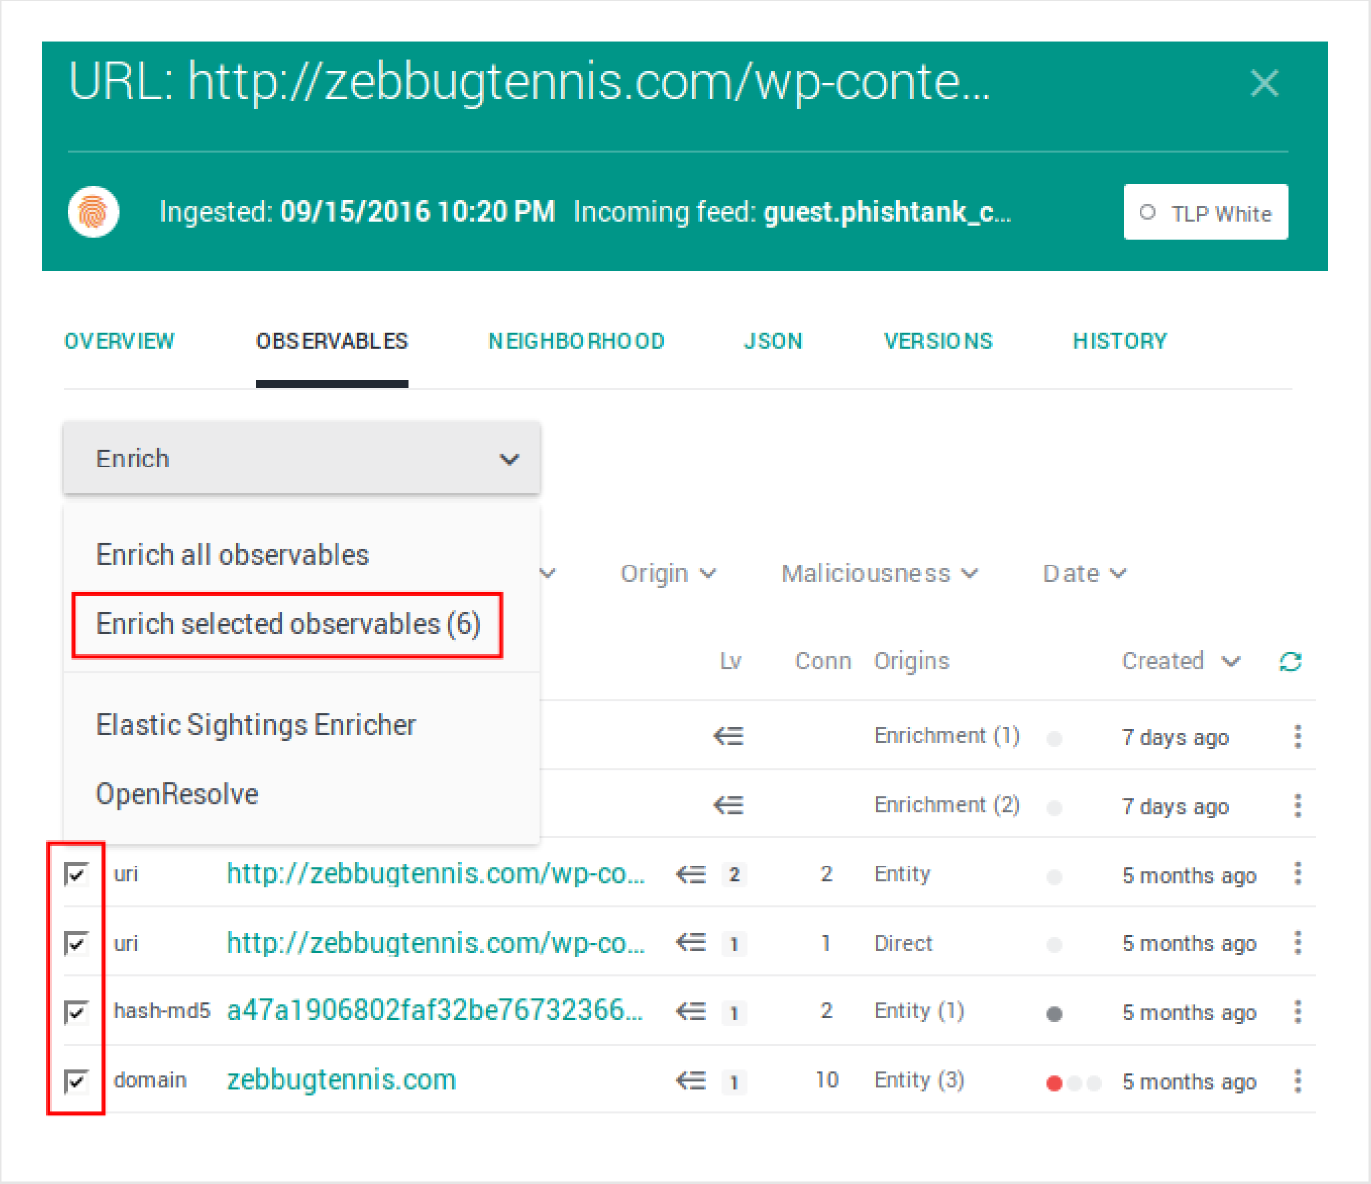
URL: http://zebbugtennis.com/wp-conte...
Ingested: 09/15/2016 10:20 PM
Incoming feed: guest.phishtank_c...
TLP White
OVERVIEW
OBSERVABLES
NEIGHBORHOOD
JSON
VERSIONS
HISTORY
Enrich
Enrich all observables
Enrich selected observables (6)
Elastic Sightings Enricher
OpenResolve
Origin
Maliciousness
Date
Lv
Conn
Origins
Created

			Lv	Conn	Origins		Created
					Enrichment (1)		7 days ago
					Enrichment (2)		7 days ago
☑	uri	http://zebbugtennis.com/wp-co...	2	2	Entity		5 months ago
☑	uri	http://zebbugtennis.com/wp-co...	1	1	Direct		5 months ago
☑	hash-md5	a47a1906802faf32be76732366...	1	2	Entity (1)	●	5 months ago
☑	domain	zebbugtennis.com	1	10	Entity (3)	●	5 months ago

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.

Work with enricher rules

View enricher rules

To view enricher rules, do the following:

- On the top navigation bar click **+ > Rules > Enrichment**

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button next to the user avatar image.
- From the drop-down menu select **Rules**.
- On the left-hand navigation sidebar click **Enrichment**.
- The default view is **Rules > Enrichment**. It shows an overview of the configured enricher rules.
  You can sort the table view by column. To do so, click the column header you want to base the data sorting on. An upward-pointing (**^**) or a downward-pointing (**v**) 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 ⚙ *configuration and settings* button next to the user avatar image.
- From the drop-down menu select **Rules**.
- On the left-hand navigation sidebar click **Enrichment**.
- The default view is **Rules > Enrichment**. It shows an overview of the configured enricher rules.
  You can sort the table view by column. To do so, click the column header you want to base the data sorting on. An upward-pointing (**^**) or a downward-pointing (**v**) 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 **Active** checkbox to enable the rule immediately after creating it.
- When you are done, 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

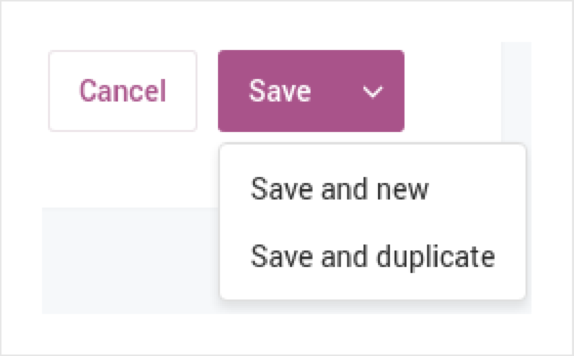

Edit enricher rules

To edit enricher rules, do the following:

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

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button next to the user avatar image.
- From the drop-down menu select **Rules**.
- On the left-hand navigation sidebar click **Enrichment**.

- The default view is **Rules > Enrichment**. It shows an overview of the configured enricher rules.
  You can sort the table view by column. 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 dotted menu 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 **Active** checkbox to enable the rule immediately after creating it.
- When you are done, 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 **✚ > Rules > Enrichment**

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button next to the user avatar image.
- From the drop-down menu select **Rules**.
- On the left-hand navigation sidebar click **Enrichment**.
- The default view is **Rules > Enrichment**. It shows an overview of the configured enricher rules.
  You can sort the table view by column. 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 dotted menu 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.

Work with enricher tasks

View enricher tasks

To view enricher tasks, do the following:

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

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button 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.

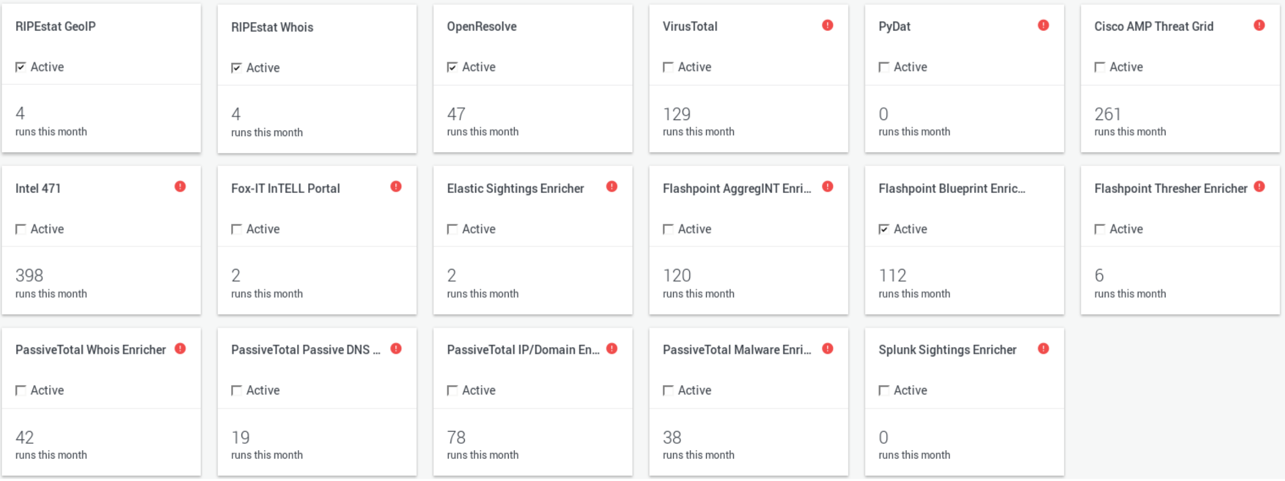

RIPEstat GeoIP
☑ Active
4
runs this month

RIPEstat Whois
☑ Active
4
runs this month

OpenResolve
☑ Active
47
runs this month

VirusTotal
☐ Active
129
runs this month

PyDat
☐ Active
0
runs this month

Cisco AMP Threat Grid
☐ Active
261
runs this month

Intel 471
☐ Active
398
runs this month

Fox-IT InTELL Portal
☐ Active
2
runs this month

Elastic Sightings Enricher
☐ Active
2
runs this month

Flashpoint AggregINT Enri...
☐ Active
120
runs this month

Flashpoint Blueprint Enric...
☑ Active
112
runs this month

Flashpoint Thresher Enricher
☐ Active
6
runs this month

PassiveTotal Whois Enricher
☐ Active
42
runs this month

PassiveTotal Passive DNS ...
☐ Active
19
runs this month

PassiveTotal IP/Domain En...
☐ Active
78
runs this month

PassiveTotal Malware Enri...
☐ Active
38
runs this month

Splunk Sightings Enricher
☐ Active
0
runs this month

Edit enricher tasks

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

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

Alternatively:

- On the top navigation bar, click the ⚙ *configuration and settings* button 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

Summary: 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. Categorization criteria take into account document content and the meaning it conveys. Taxonomies make it easier for you to maintain content, and they help users find what they're looking for. They provide a hierarchical framework to structure tags and to point out relationships among tags. In turn, tag relationships form 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 seach results.

The Taxonomy feature

You can use the **Taxonomy** feature to define specific categories to organize entity tags. You can create as many taxonomies as you need, and you can hierarchically relate taxonomy entries with parent-child relationships.

Predefined taxonomies

The EclecticIQ Platform ships with the following predefined taxonomy entry sets:

- Admiralty code: based on the **two-character Admiralty System code** `(https://en.wikipedia.org/wiki/admiralty_code)`, it assesses the reliability of the source, and the accuracy level of the information.
- Kill chain phase: defines 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 taxonomy to label entities with tags that define the level of reliability of the entity source, and the level of accuracy of the entity data. The Admiralty Code taxonomy makes it easier to filter entities and information based on criteria like relevance and credibility.

Entity source reliability	**Entity data accuracy**
Reliable	Confirmed by other sources
Usually reliable	Probably True
Fairly reliable	Possibly True
Not usually reliable	Doubtful
Unreliable	Improbable
Cannot be judged	Truth cannot be judged

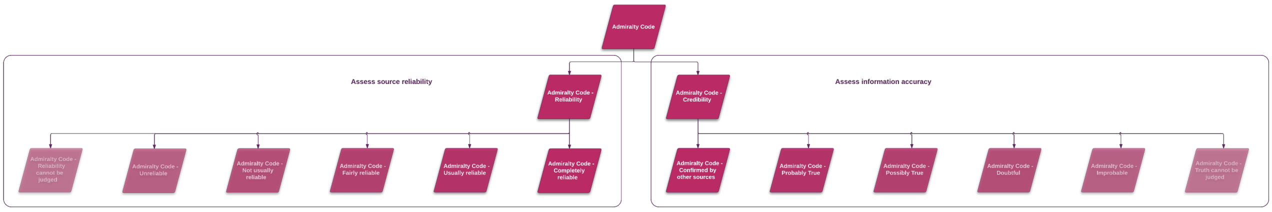
Admiralty Code
Assess source reliability
Admiralty Code - Reliability
Admiralty Code - Reliability cannot be judged
Admiralty Code - Unreliable
Admiralty Code - Not usually reliable
Admiralty Code - Fairly reliable
Admiralty Code - Usually reliable
Admiralty Code - Completely reliable
Assess information accuracy
Admiralty Code - Credibility
Admiralty Code - Confirmed by other sources
Admiralty Code - Probably True
Admiralty Code - Possibly True
Admiralty Code - Doubtful
Admiralty Code - Improbable
Admiralty Code - Truth cannot be judged

Kill chain

In the context of cyber threat defense, a 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 model to:

- Identify the root cause of an intrusion, and quantify the damage it causes.
- Plan a defensive course of action to neutralize it.

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.

Kill chain phase	Description
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*, 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	For example use analytics, auditing, logging tools, and intrusion detection systems (IDS) to detect the intrusion.
Deny	For example use patching, firewall rules, access control lists (ACL), and intrusion prevention systems (IPS) to deny exploitation.
Disrupt	For example use data execution prevention (DEP) and intrusion prevention systems to block or otherwise disturb exploitation.
Degrade	For example use queuing or a tarpit to hinder or otherwise reduce exploitation.
Deceive	For example use DNS redirection or a honeypot to divert exploitation to a decoy.
Destroy	Take control of the attacker's system, and completely neutralize it.

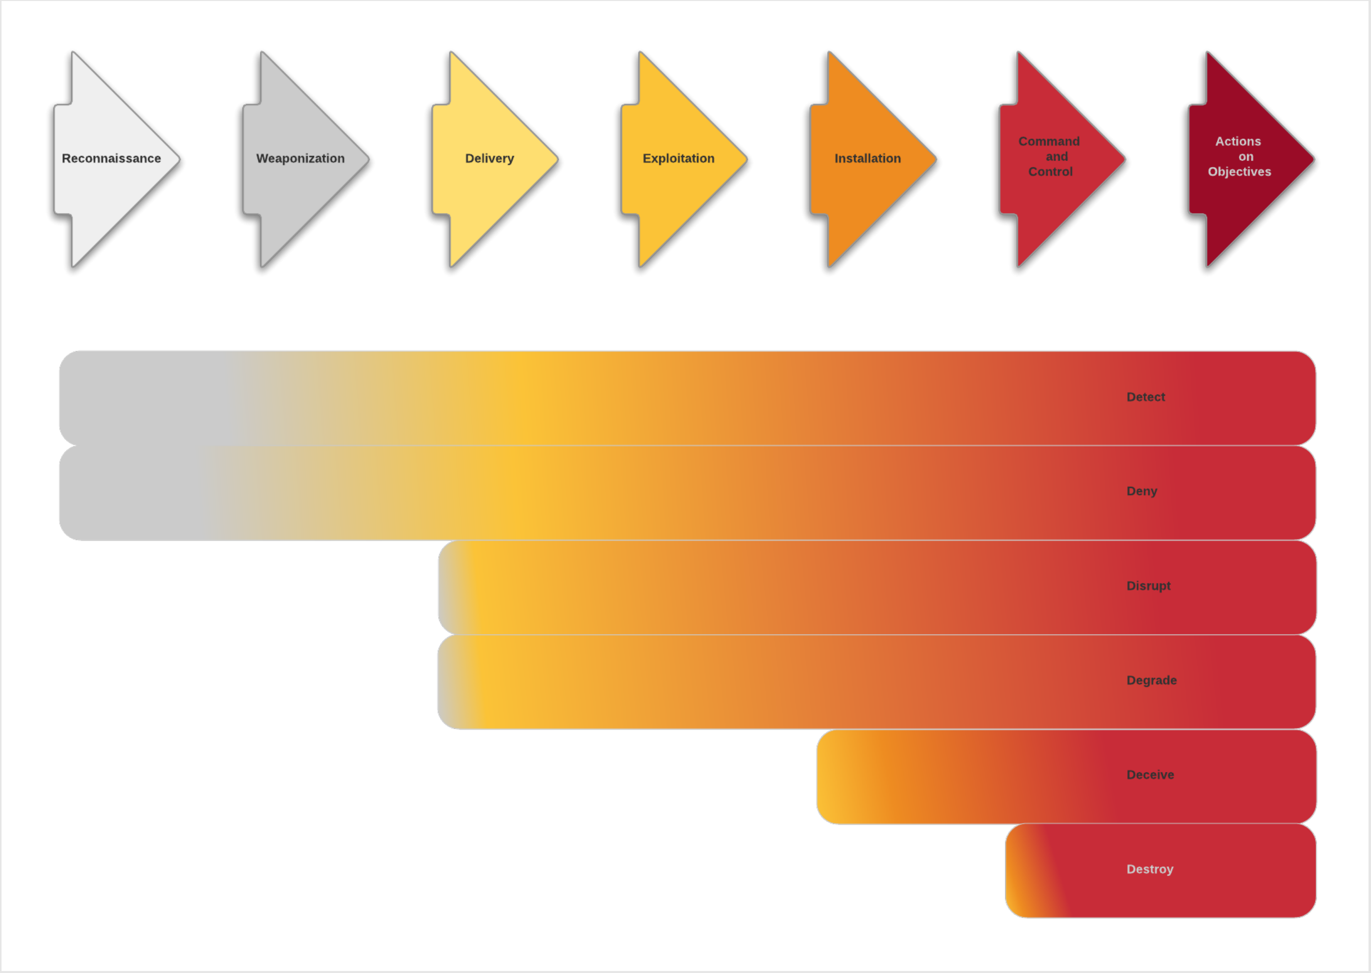

Create a taxonomy entry

✓ 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 left-hand navigation sidebar click **Taxonomy**.
  On the **Taxonomy** page, an overview of the existing entries is displayed in a table.
  You can sort the table view by column. 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 **+ Entry** button.

- On the **Taxonomy > Create new taxonomy** page, fill out the input fields to define the new taxonomy entry, and whether it is a *parent*, top-level entry, or a *child* entry, i.e. subordinate to a parent:
    - **Name**: define 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**: if you are creating a subordinate/child taxonomy entry, from the drop-down menu select the parent entry you want to relate the child to.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.
- The newly created taxonomy entry detail pane is displayed, where you can review the entry information. Child entry details on this pane include a clickable link to the corresponding parent.

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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

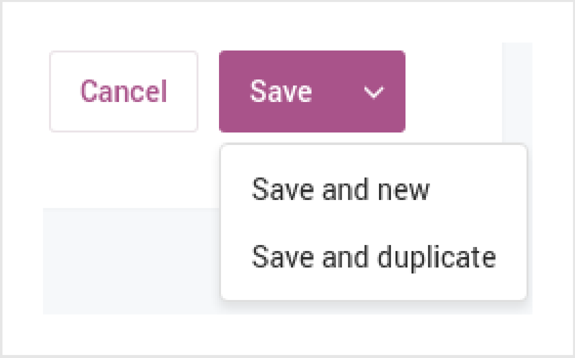

Edit a taxonomy entry

To edit an existing taxonomy entry, do the following:

- On the left-hand navigation sidebar click **Taxonomy**.
On the **Taxonomy** page, an overview of the existing entries is displayed in a table.
You can sort the table view by column. 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 dotted menu icon.
- From the drop-down menu select **Edit**.

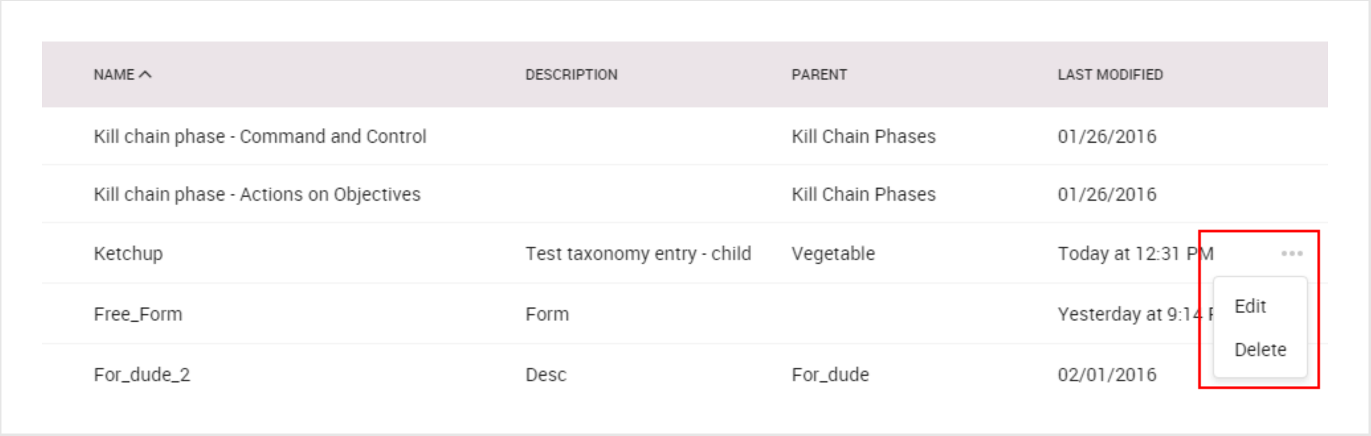

NAME ^	DESCRIPTION	PARENT	LAST MODIFIED
Kill chain phase - Command and Control		Kill Chain Phases	01/26/2016
Kill chain phase - Actions on Objectives		Kill Chain Phases	01/26/2016
Ketchup	Test taxonomy entry - child	Vegetable	Today at 12:31 PM
Free_Form	Form		Yesterday at 9:14 P
For_dude_2	Desc	For_dude	02/01/2016

- On the **Taxonomy > Edit Taxonomy** page, edit the name, the description, or the parent-child hierarchy relationship as needed.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

Delete a taxonomy entry

To delete an existing taxonomy entry, do the following:

- On the left-hand navigation sidebar click **Taxonomy**.
  On the **Taxonomy** page, an overview of the existing entries is displayed in a table.
  You can sort the table view by column. 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 dotted menu icon.
- From the drop-down menu select **Delete**.

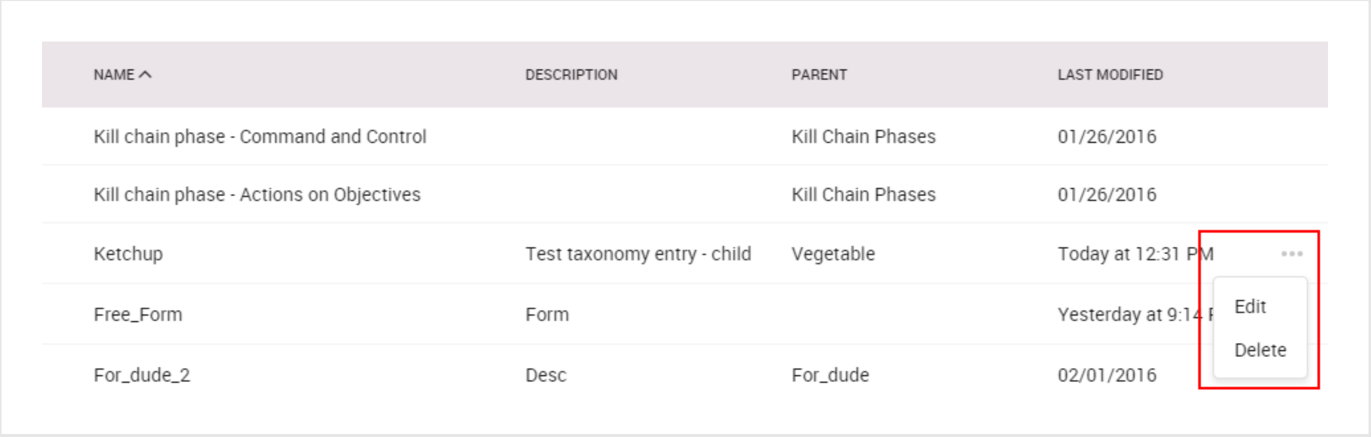

NAME ^	DESCRIPTION	PARENT	LAST MODIFIED
Kill chain phase - Command and Control		Kill Chain Phases	01/26/2016
Kill chain phase - Actions on Objectives		Kill Chain Phases	01/26/2016
Ketchup	Test taxonomy entry - child	Vegetable	Today at 12:31 PM
Free_Form	Form		Yesterday at 9:14 P
For_dude_2	Desc	For_dude	02/01/2016

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

If you delete a parent entry in the taxonomy section, any children related to the removed parent are kept in the taxonomy section, but they lose the parent-child relationship and they become top-level taxonomy entries.

Rules

Summary: Rules give you granular control over entity extracts 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 more 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.

Extract rules help you zero in on the valuable bits of information by filtering out the clutter from your data. By getting rid of the noise, you can see a clearer and sharper picture.

Entity rules allow you to automatically assign free tags or taxonomy tags to ingested entities. Besides adding semantic relevance, you can use tags within a workflow to group entities sharing similar characteristics.

Entity and extract rules are highly customizable to give you granular control over your data. For example, you can create rules to target specific entities or extracts 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.
- Extract rules act like filters: they filter entity extract as malicious, safe, or ignorable, based on a predefined set of criteria.

Entity rules

Add an entity rule

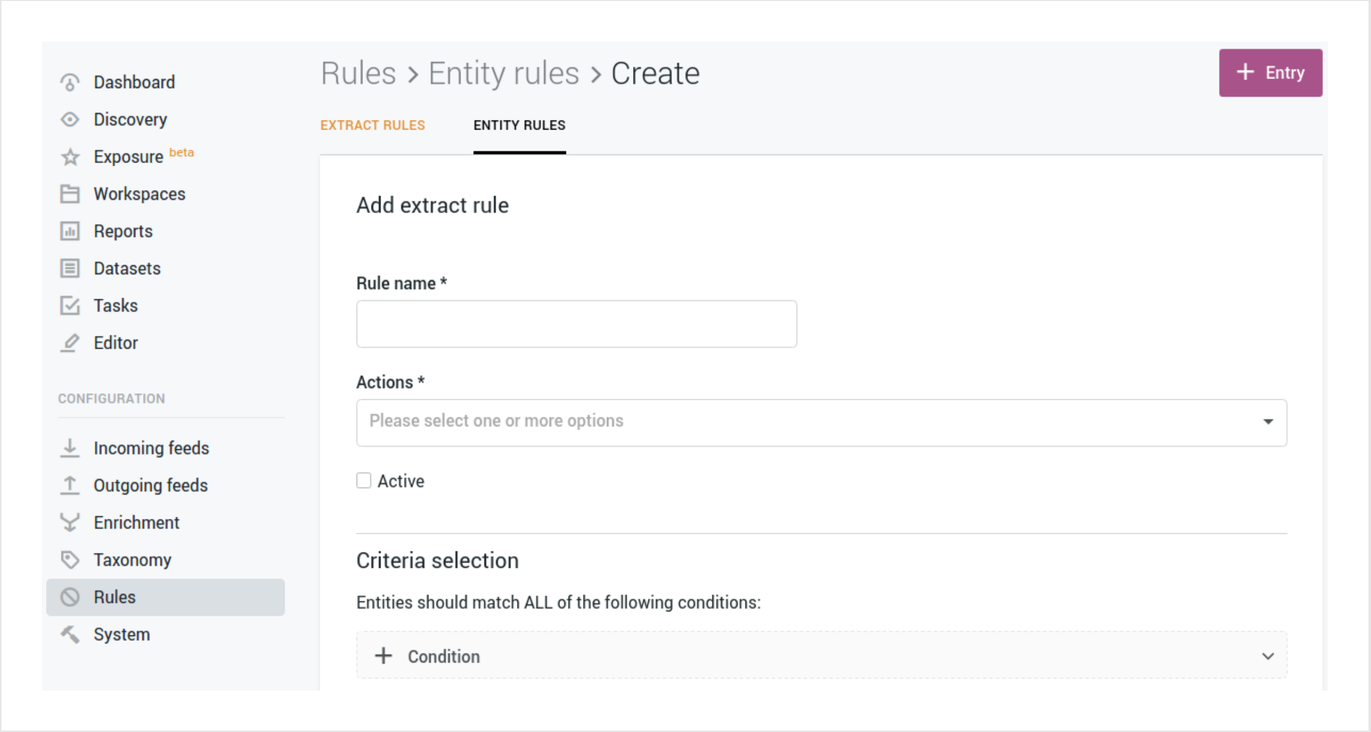

Rule name *

Actions * Please select one or more options

☐ Active

Criteria selection

Entities should match ALL of the following conditions:

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

To create a new entity rule, do the following:

- On the left-hand navigation sidebar click **Rules**.
- On the **Entity rules** page, click the ✚ **Entry** button.

- On the **Add entity rule** page, define the new rule criteria to automatically tag, add entities to datasets, or both:
  - **Rule name**: enter a name to identify the rule. It should be descriptive and easy to remember.
  - Select the **Active** 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 **Active** checkbox to enable the rule immediately after creating it.

Besides a name and an action, a valid rule needs to include 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:

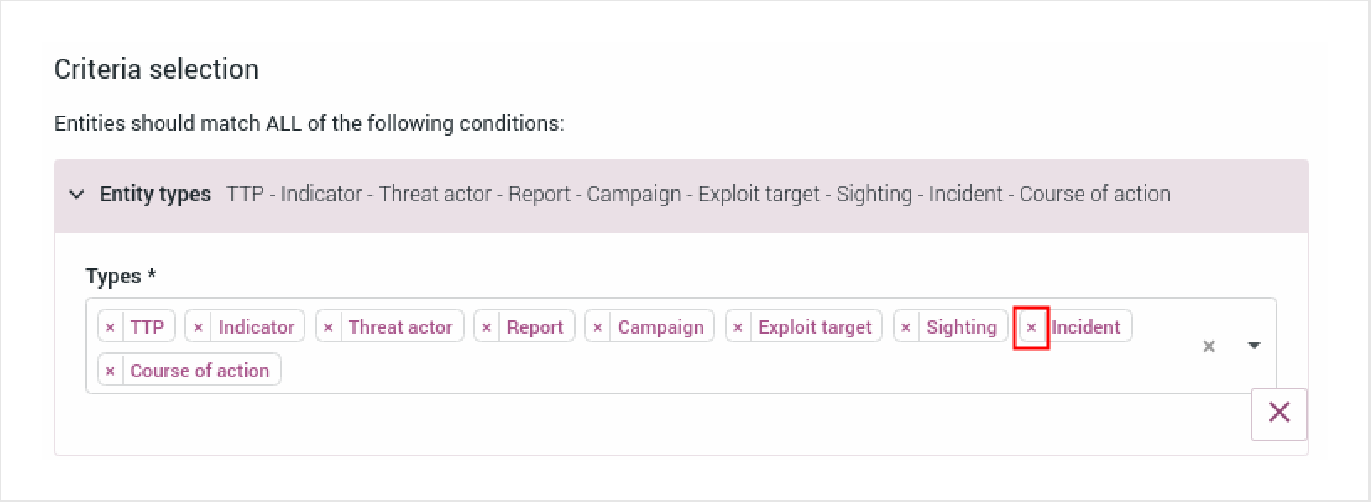

Criteria selection

Entities should match ALL of the following conditions:

Entity types TTP - Indicator - Threat actor - Report - Campaign - Exploit target - Sighting - Incident - Course of action

Types *

× TTP × Indicator × Threat actor × Report × Campaign × Exploit target × Sighting × Incident × Course of action

To remove a condition and its content from the criteria selection for the rule, click the deletion icons available on the right side of the input field, and immediately below the bottom-right corner of the input field:

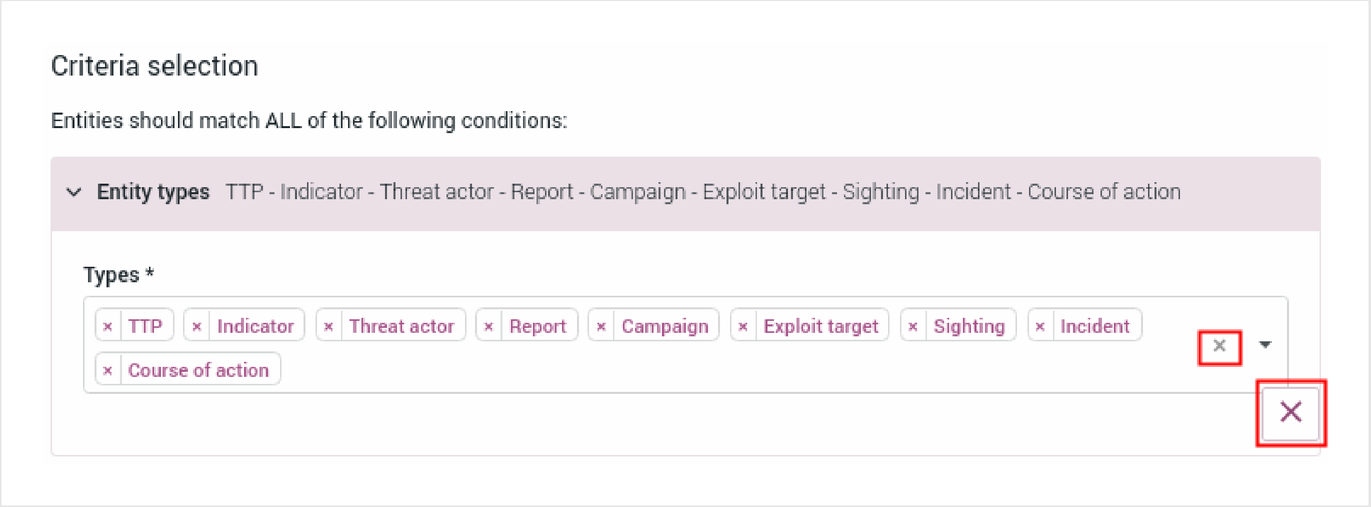

Criteria selection

Entities should match ALL of the following conditions:

Entity types TTP - Indicator - Threat actor - Report - Campaign - Exploit target - Sighting - Incident - Course of action
Types *
× TTP
× Indicator
× Threat actor
× Report
× Campaign
× Exploit target
× Sighting
× Incident
× Course of action

- **Content criteria**: this input field takes key/value pairs. The key is always a JSON path, and the value is always a regex.
  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.
  To write a JSON path pointing to the title or the description of an entity in JSON format, instead of `data.title` or `data.description` you only need to enter `title` or `description`. The `data` root is implied.
  - **Content > Path**: based on the specified JSON path you enter in this field, the rule searches for a corresponding match in the JSON data structures describing 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* inside the entity it should go look for a 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. Therefore, 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:

- On the left-hand navigation sidebar click **Discovery**.
- On the **Discovery > Entities** table format overview, click an entity.
- An overlay slides in from the side of the screen.
- On the selected entity detail pane, click **JSON** to view the entity JSON data structure.

Alternatively, do the following:

- On the left-hand navigation sidebar click **Discovery**.
- On the **Discovery > Entities** table format overview, click an entity.
- An overlay slides in from the side of the screen.
- On the selected entity detail pane, click **Actions > Download original > JSON** to view the JSON data structure for that entity type.

- **Content > Value**: define a regex to specify the data pattern you want to apply the rule to.
  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 this 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 the conditional criteria for the current rule.
- **Source**: from the drop-down menu select an incoming feed or an enricher to use as 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.
  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.
- When you are done, 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

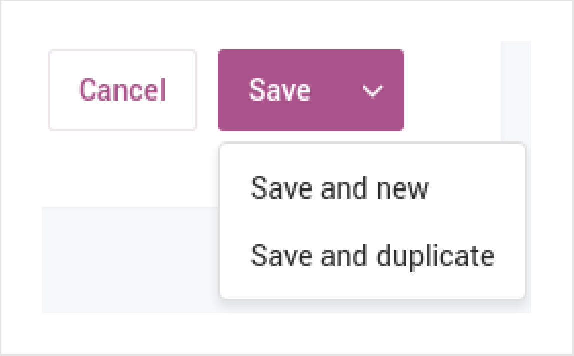

Extract rules

Add an extract rule

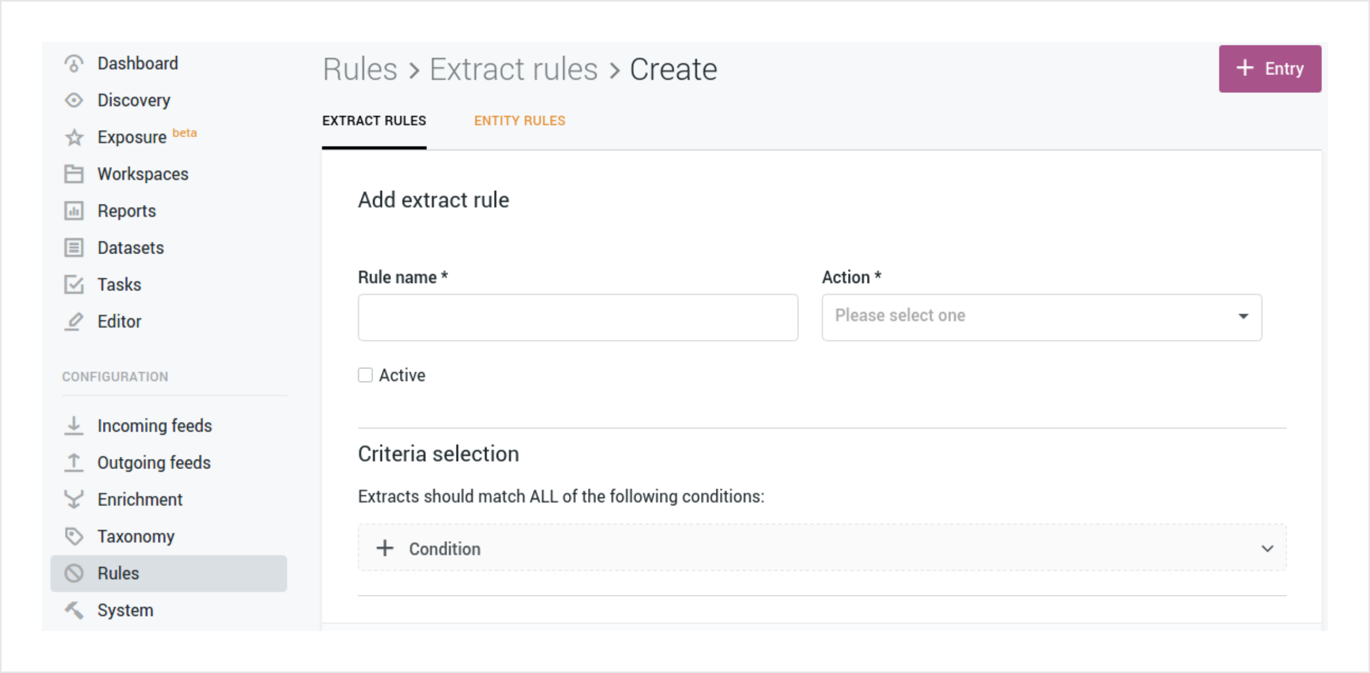

✓ On the forms, input fields marked with an asterisk are required.

To create a new extract rule, do the following:

- On the left-hand navigation sidebar click **Rules**.
- On the **Extract rules** page, click the **+ Entry** button.

- On the **Add extract rule** page, define the new rule criteria to filter entity extracts as needed:
  - **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 extracts matching *all* the conditions defined under **Criteria selection** are ignored. If any extracts are found that can be ignored, you can delete them in bulk from the platform by selecting **Delete all matching extracts**.
      It is a good idea to review the extracts before deleting them.
    - **Mark as safe**: *all* entity extracts matching *all* the conditions defined under **Criteria selection** are flagged as safe, and therefore non-threatening.
    - **Mark as malicious** *all* entity extracts matching *all* the conditions defined under **Criteria selection** are flagged as malicious. These are the ones analysts probably want to follow up on close. The maliciousness confidence level makes it easier to triage and prioritize threat severity.

      When an extract is marked as malicious, it cannot transition to *safe* or *ignore*. It can only become more malicious, that is, it can only transition from a low to a high confidence level. The rationale behind it is that once an extract is flagged as harmful, it cannot become 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 extract(s) may or may not be malicious.
      - **Malicious - Medium confidence**: based on the available intelligence, the threat represented by the entity extract(s) is likely to be malicious.
      - **Malicious - High confidence**: based on the available intelligence, the threat represented by the entity extract(s) is malicious.
  - Select the **Active** checkbox to enable the rule immediately after creating it.

Besides a name and an action, a valid rule needs to include 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:

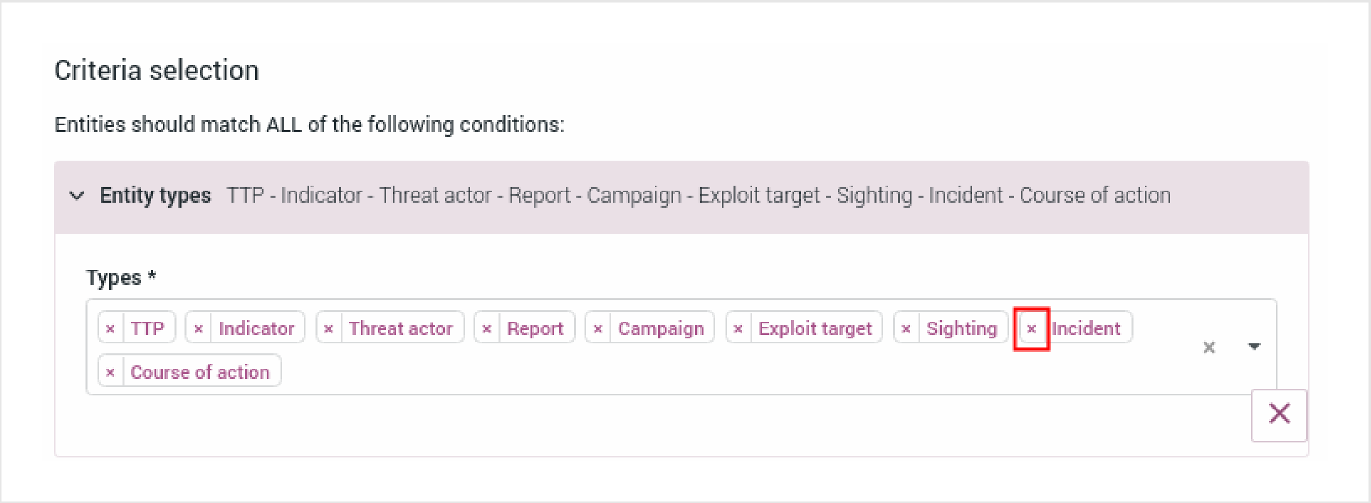

Criteria selection

Entities should match ALL of the following conditions:

Entity types TTP - Indicator - Threat actor - Report - Campaign - Exploit target - Sighting - Incident - Course of action

Types *

× TTP × Indicator × Threat actor × Report × Campaign × Exploit target × Sighting × Incident × Course of action

To remove a condition and its content from the criteria selection for the rule, click the deletion icons available on the right side of the input field, and immediately below the bottom-right corner of the input field:

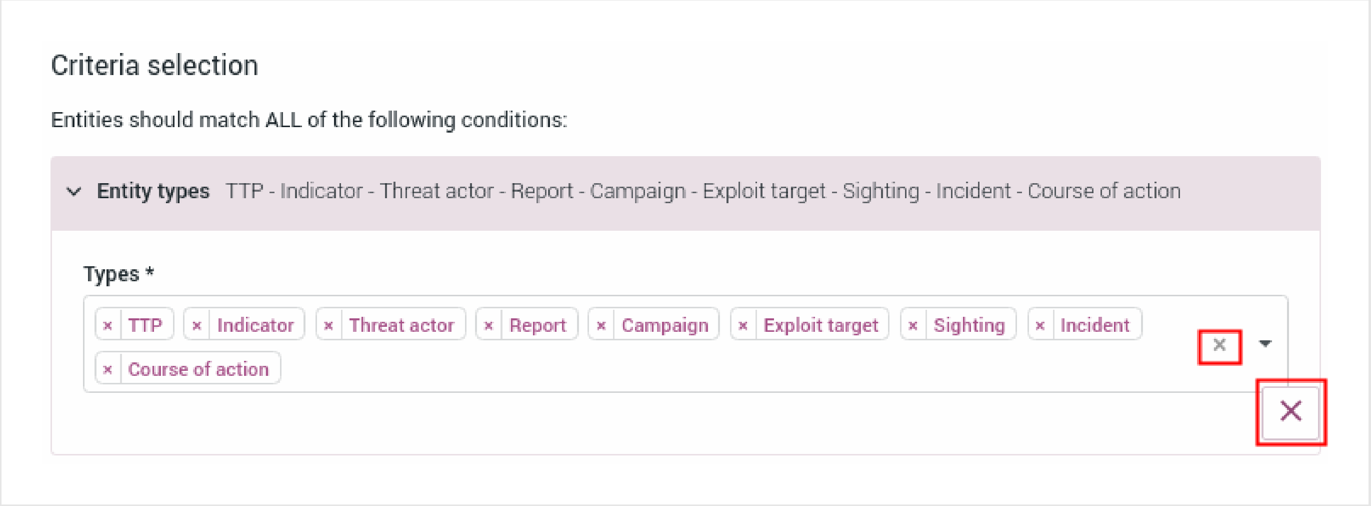

Criteria selection

Entities should match ALL of the following conditions:

Entity types TTP - Indicator - Threat actor - Report - Campaign - Exploit target - Sighting - Incident - Course of action

Types *

- **Extract types**: from the drop-down menu select one or more extract types, i.e. the data types describing the corresponding entity extract data. For example you may want to include in the rule processing a specific city name, an actor, a range of IP addresses or telephone numbers.
- **Paths**: based on the specified JSON path you enter in this field, the rule searches for a corresponding match in the JSON data structures describing 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* inside the entity it should go look for a 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. Therefore, 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:

- On the left-hand navigation sidebar click **Discovery**.
- On the **Discovery > Entities** table format overview, click an entity.
- An overlay slides in from the side of the screen.
- On the selected entity detail pane, click **JSON** to view the entity JSON data structure.

Alternatively, do the following:

- On the left-hand navigation sidebar click **Discovery**.
- On the **Discovery > Entities** table format overview, click an entity.
- An overlay slides in from the side of the screen.
- On the selected entity detail pane, click **Actions > Download original > JSON** to view the JSON data structure for that entity type.

- **Value matches**: define a regex to specify the data pattern you want to apply the rule to.
  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 this 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 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**.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

Save extract 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

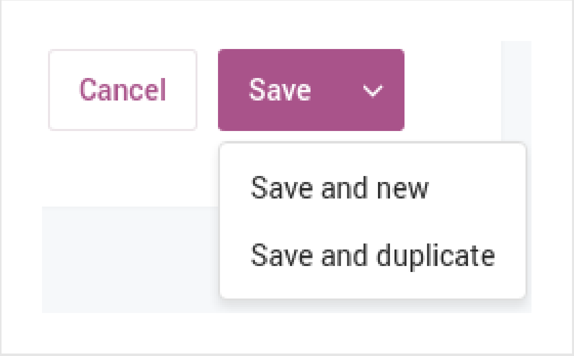

Derivation and levels

Derivation — **Original** vs **Derived** extracts — and levels — level **1** and level **2** extracts — work together to make it easier to identify and analyze the extracts you need to act on by including them in your prevention and/or detection toolchains.

You can filter extracts by derivation and levels to discard unwanted noise. You can zero in on specific bits of information to examine them, and to investigate any relationships they may have with other entities or extracts.

Level 1	The data origin is a value extracted from a field inside a *CybOX* object. Example: a URI in a CybOX `URIObj` field.
Level 2	The data origin is a value extracted from a field inside a *STIX* object. Example: a URI in a STIX `Reference` field.
Original	The original data is extracted and stored in the resulting extract as is. Example: a URI.
Derived	The original data is processed, and then a subset of that data is extracted and stored in the resulting extract. Example: a domain that is part of a URI.

Extracts are bits of information that contribute with additional context to an entity description. The platform can flag extracts to automate processes such as:

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

A set of rules can handle the flags and route extracts to a prevention and/or a detection system, or mark 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 extract 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 extract as **Original/1**, it handles it as follows:

- It assigns the extract an initially *low maliciousness* level;
- It flags it as a *level 1* extracted data to indicate that it originates from a CybOX object. Therefore, it is directly related to its source, 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 extract 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</stixComm
on: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 extract as **Original/2**, it handles it as follows:

- It does not assign the extract any maliciousness level;
- It flags it as a *level 2* extracted data to indicate that it does not originate from a CybOX object, but from a STIX field. Therefore, it is more distant from its source, and possibly not 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 extract value, in this example a URI -->
<URIObj:Value condition="Equals">http://x4z9arb.cn/4712</URIObj:Value>

<!-- The derived extract 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 extract as **Derived/1**, it handles it as follows:

- It does not assign the extract any maliciousness level;
- It flags it as a *level 1* extracted data to indicate that it originates from a CybOX object. Therefore, it is directly related to its source, 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 extract value, in this example a URI -->
<stixCommon:Reference>https://technet.microsoft.com/library/security/2887505</stixCommon:Reference>

<!-- The derived extract 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 extract as **Derived/2**, it handles it as follows:

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

Edit rules

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

- On the left-hand navigation sidebar click **Rules**.
- On the **Rules** page, go to **Extract rules** or to **Entity rules**, 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.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

Delete rules

ℹ It is a good idea to review the extracts before deleting them.

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

- On the left-hand navigation sidebar click **Rules**.
- On the **Rules** page, go to **Extract rules** or to **Entity rules**, and then click the row corresponding to the rule you want to delete.
- On the row corresponding to the rule you want to delete, click the dotted menu icon, and then from the context menu select **Delete**.
- On the pop-up confirmation dialog, confirm your choice.
- The rule is removed from the list.

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

Filter rules

On the **Rules** page you can see table format overviews of the existing extract 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 sources for the rules
- **Active**: select if you want to display only **Active** rules, only **Inactive** rules, or both sets
- **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 extract rules.

Example

For example let’s assume we want to apply an extract rule that zeroes in on ipv4 IP address extracts. We want the rule to target IP address extracts only when they are included in a sighting.
The criteria we set for the rule are:

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

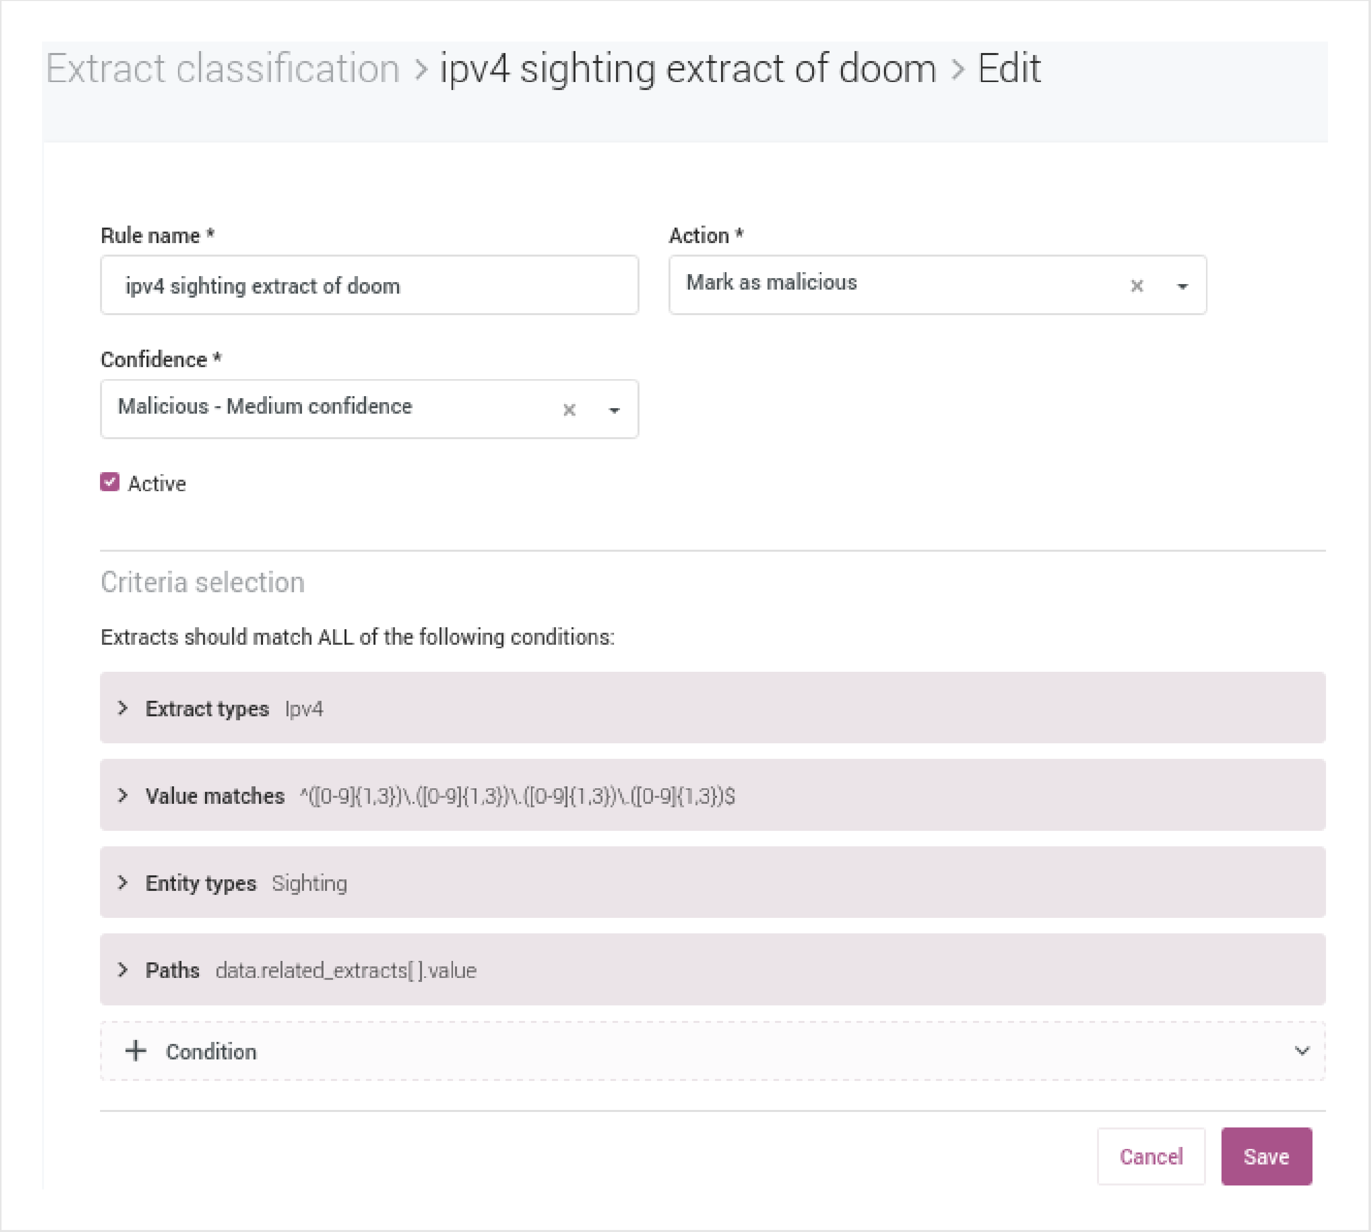
Extract classification > ipv4 sighting extract of doom > Edit

Rule name *
ipv4 sighting extract of doom

Action *
Mark as malicious

Confidence *
Malicious - Medium confidence

☑ Active

Criteria selection

Extracts should match ALL of the following conditions:

Extract types Ipv4

Value matches `^([0-9]{1,3})\.([0-9]{1,3})\.([0-9]{1,3})\.([0-9]{1,3})$`

Entity types Sighting

Paths data.related_extracts[ ].value

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 extracts of 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%2C6AS2Me
0aD0dEag3aS0hI7h42&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
  },

  "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": []
}
```

Extracts

View matching extracts

When the **Action** for an extract rule is **Ignore**, no specific action is automatically executed on the matching extracts. If you want to delete them, you need to manually initiate the action.

However, you may wish to inspect the ignored extracts before deleting them. You can do so on the **Matches** tab on the rule detail pane, which becomes available when the specified rule action is set to **Ignore**.

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

- On the left-hand navigation sidebar click **Rules**, and then choose **Extract rules**.
- On the **Rules** page, go to **Extract rules** or to **Entity rules**, and then click the row corresponding to the rule you want to view to display the rule detail pane..
- The **Matches** tab is available between the **Details** and the **History** tabs.

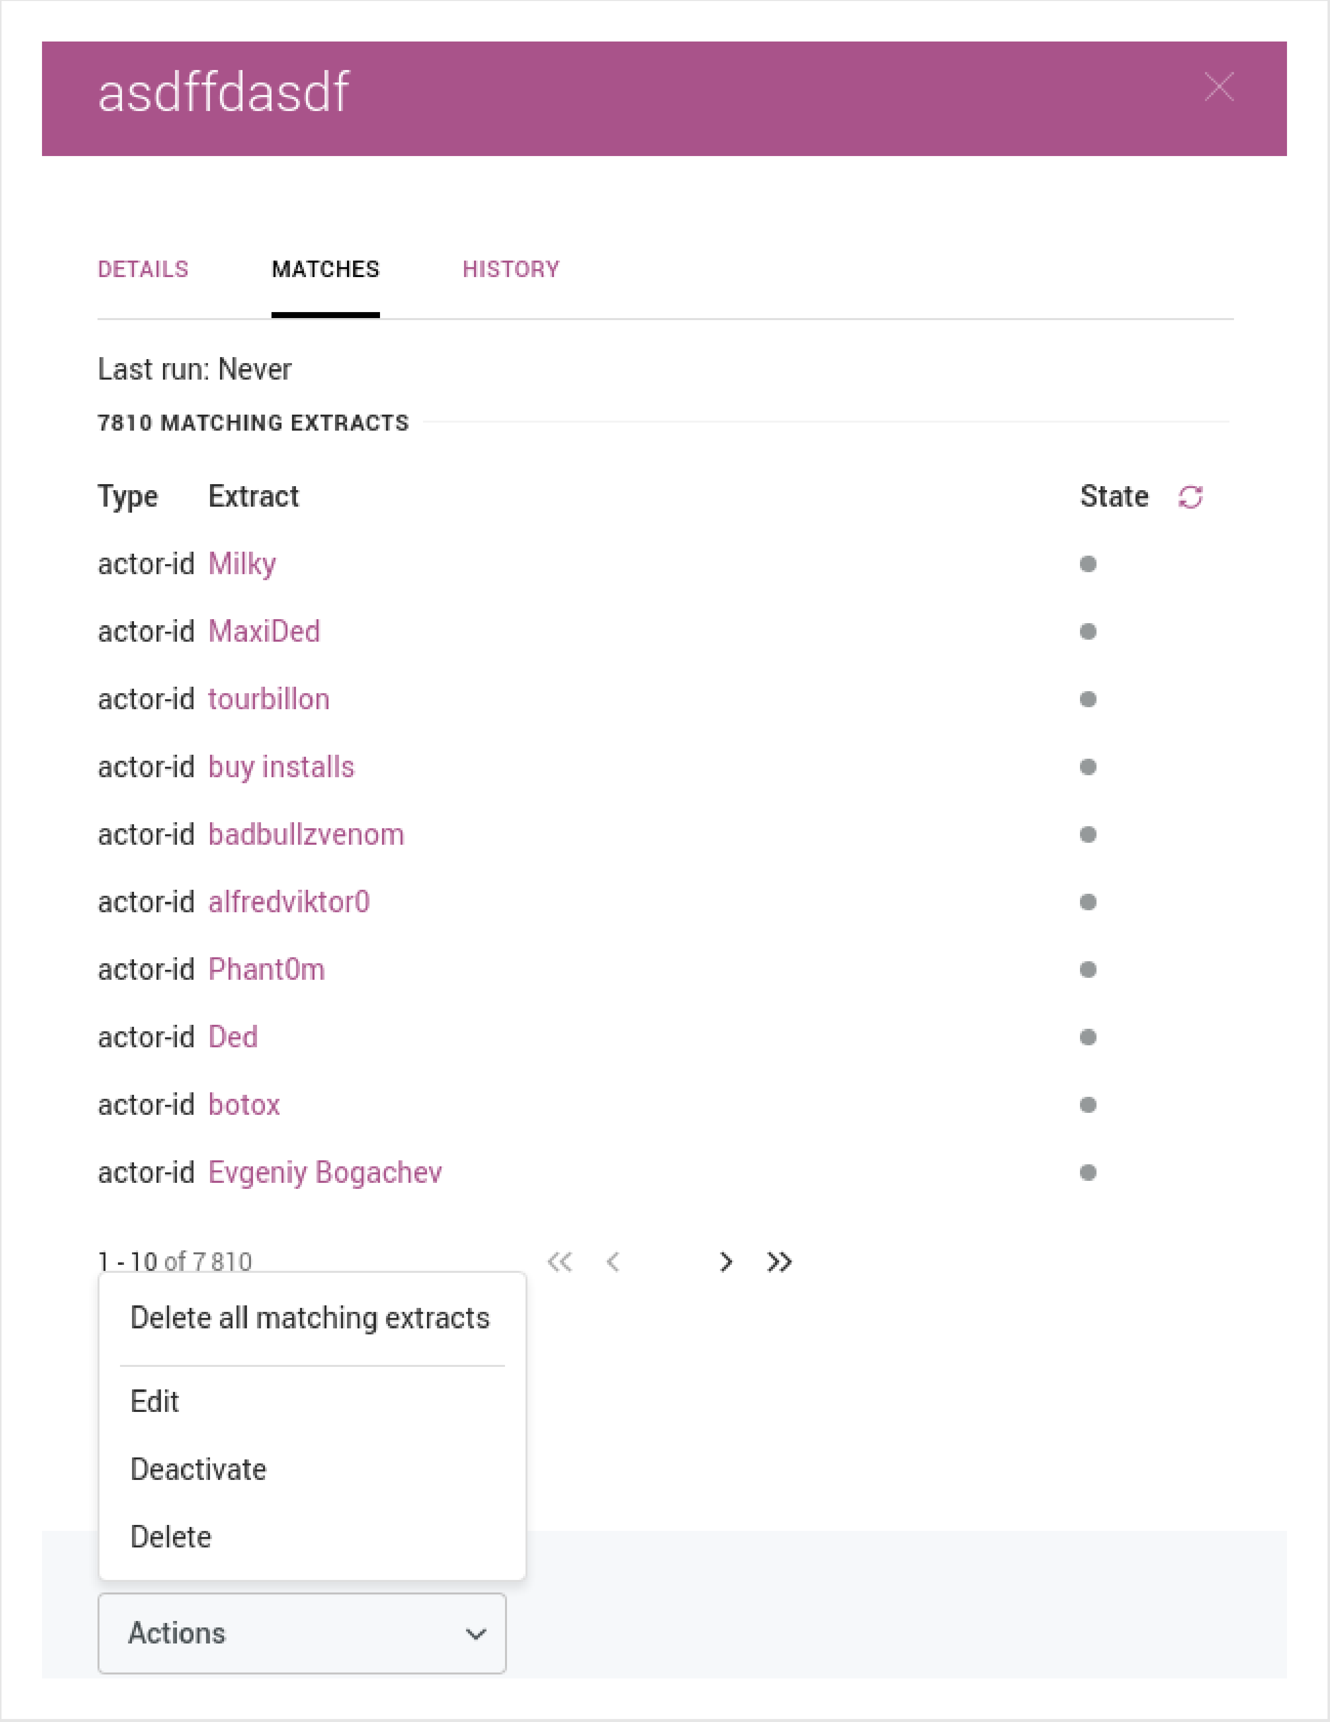

Last run: Never

7810 MATCHING EXTRACTS

Type	Extract	State
actor-id	Milky	●
actor-id	MaxiDed	●
actor-id	tourbillon	●
actor-id	buy installs	●
actor-id	badbullzvenom	●
actor-id	alfredviktor0	●
actor-id	Phant0m	●
actor-id	Ded	●
actor-id	botox	●
actor-id	Evgeniy Bogachev	●

1 - 10 of 7 810

The **Matches** tab shows the matching extract type(s) the rule has identified and the corresponding state:

- **Kind**: the matching extract type; for example, *domain*
- **Value**: the domain name value; for example, *www.iphishyourdata.biz*.

On this tab, you can further explore, or carry out actions like these:

- To view a list of all the entities that share a given extract, click the desired extract name on the pane.
- To refresh the view, click the ⟳ refresh icon on the upper-right portion of the pane.
- To edit, deactivate or delete the rule, or to delete all matching extracts when the **Action** of the rule is set to **Ignore**, select an option from the **Actions** pop-up menu.

Delete matching extracts

When the **Action** configuration option of an extract rule is set to **Ignore**, any extracts matching the rule criteria can be disregarded, and they can be deleted.
To delete all extracts matching an ignore action rule, do the following:

- On the left-hand navigation sidebar click **Rules**, and then choose **Extract rules**.
- On the row corresponding to the rule whose matching extracts you want to delete, click the dotted menu icon, and then from the context menu select **Delete all matching extracts**.
- The extracts matching the rule are deleted from the platform database, as well as from the platform history.

Get extract types via API

You can retrieve a JSON response with all supported extract 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/
```

- 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"
    }

  ]
}
```

Extract types

The available extract types are:

actor-id
asn
bank-account
card
card-owner
company
cve
domain
email
email-subject
file
forum-name

forum-thread
rum-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
eui-64
malware
mutex
name
nationality

netname
organization
person
port
process
product
raw-artifact
registrar
telephone
uri
uri-hash-sha256
winregistry

Dashboard

Summary: 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.

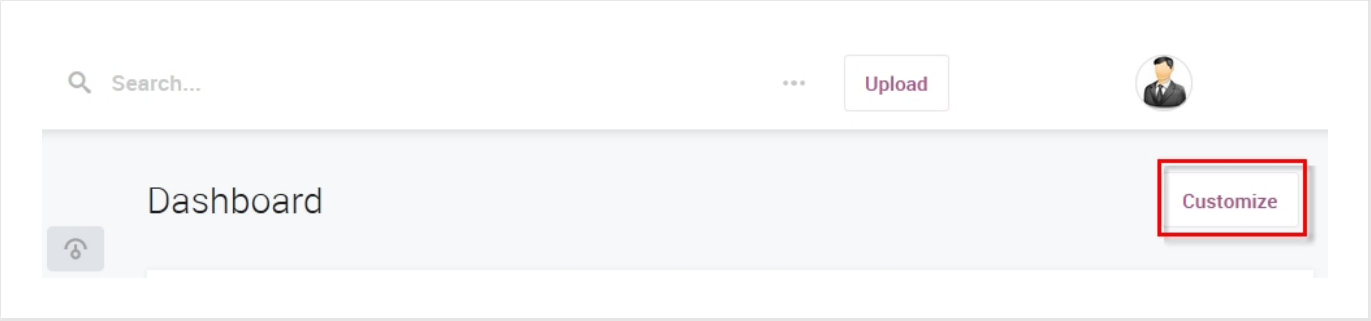

- To add a gauge from the list to the dashboard view, select the corresponding checkbox.
- To remove a gauge from the dashboard view, deselect the corresponding checkbox.
- When you are done, click **Return to the dashboard** to apply the changes and go back to the dashboard.

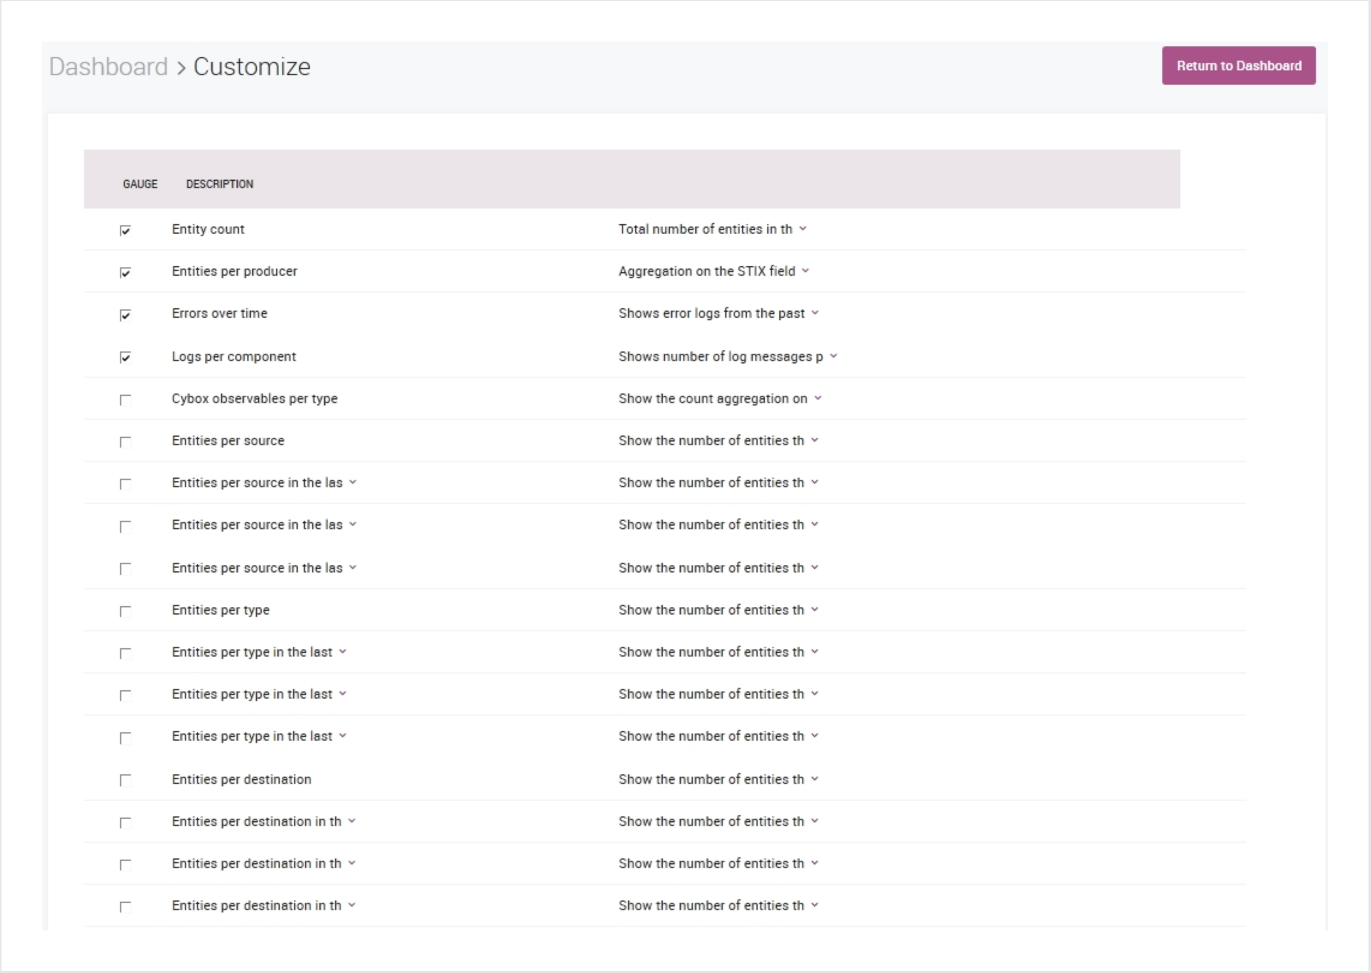

GAUGE	DESCRIPTION	
☑	Entity count	Total number of entities in th ˅
☑	Entities per producer	Aggregation on the STIX field ˅
☑	Errors over time	Shows error logs from the past ˅
☑	Logs per component	Shows number of log messages p ˅
☐	Cybox observables per type	Show the count aggregation on ˅
☐	Entities per source	Show the number of entities th ˅
☐	Entities per source in the las ˅	Show the number of entities th ˅
☐	Entities per source in the las ˅	Show the number of entities th ˅
☐	Entities per source in the las ˅	Show the number of entities th ˅
☐	Entities per type	Show the number of entities th ˅
☐	Entities per type in the last ˅	Show the number of entities th ˅
☐	Entities per type in the last ˅	Show the number of entities th ˅
☐	Entities per type in the last ˅	Show the number of entities th ˅
☐	Entities per destination	Show the number of entities th ˅
☐	Entities per destination in th ˅	Show the number of entities th ˅
☐	Entities per destination in th ˅	Show the number of entities th ˅
☐	Entities per destination in th ˅	Show the number of entities th ˅

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

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

Search

You can find the search box on the top bar:

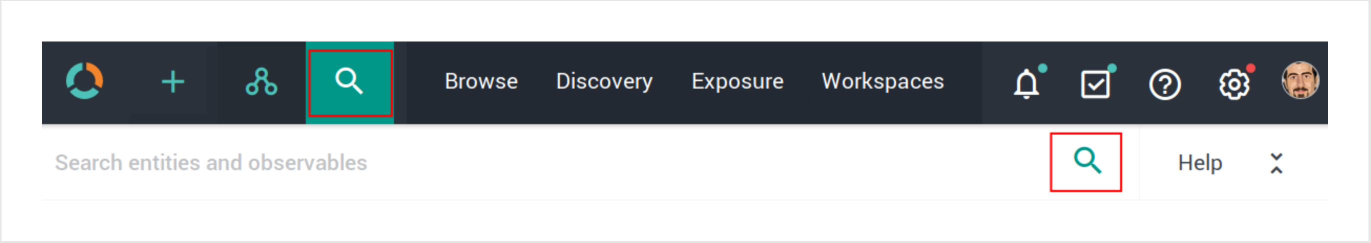

Enter search terms and search queries, and then press **ENTER** or click the search icon to run the search. The searches you run from this search box are platform-wide.

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.

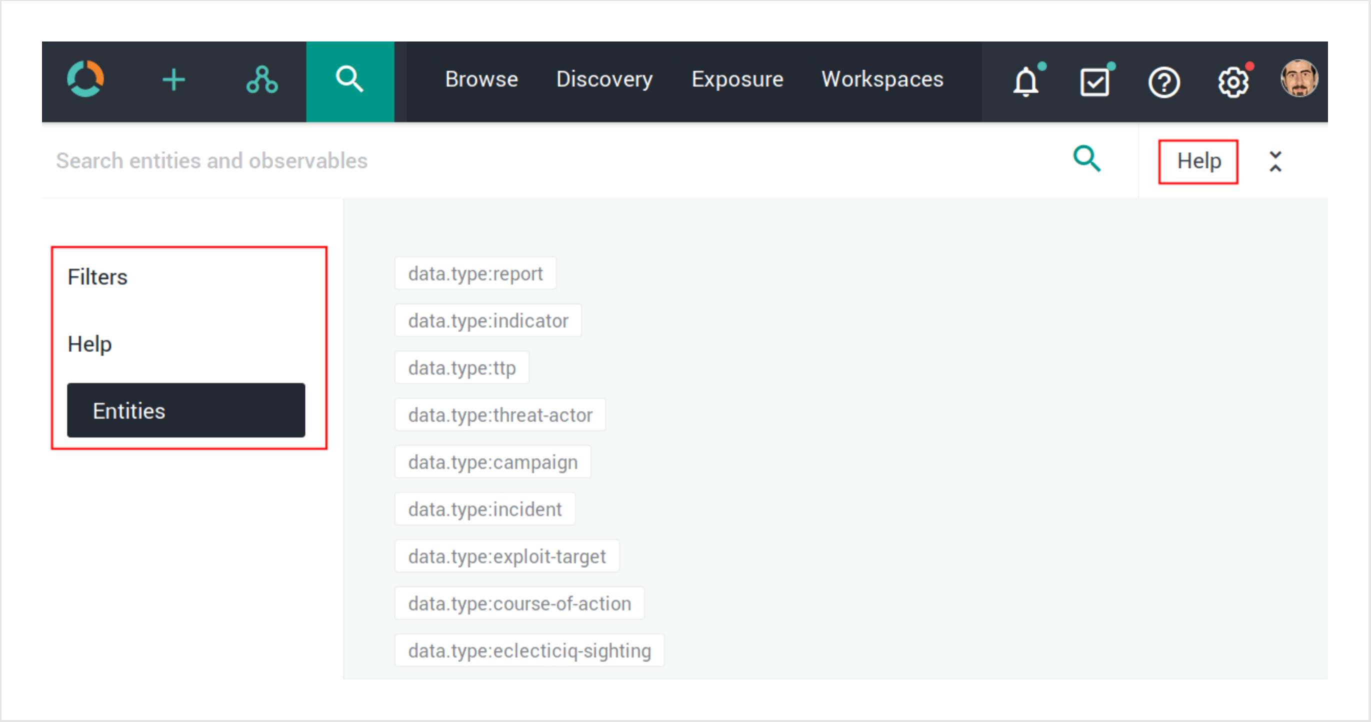

data.type:report
data.type:indicator
data.type:ttp
data.type:threat-actor
data.type:campaign
data.type:incident
data.type:exploit-target
data.type:course-of-action
data.type:eclecticiq-sighting

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

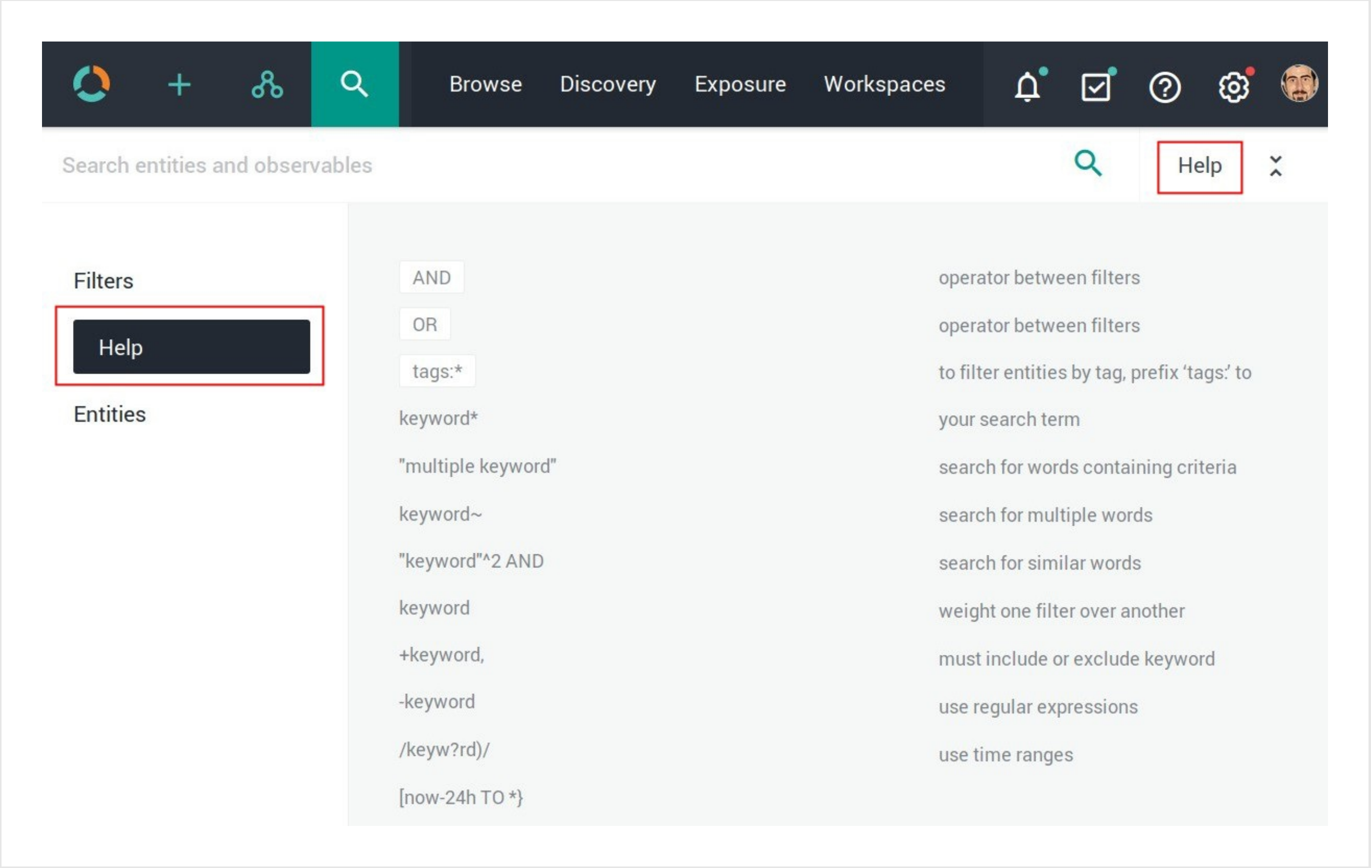

AND	operator between filters
OR	operator between filters
tags:*	to filter entities by tag, prefix 'tags:' to
keyword*	your search term
"multiple keyword"	search for words containing criteria
keyword~	search for multiple words
"keyword"^2 AND	search for similar words
keyword	weight one filter over another
+keyword,	must include or exclude keyword
-keyword	use regular expressions
/keyw?rd)/	use time ranges
[now-24h TO *}	

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

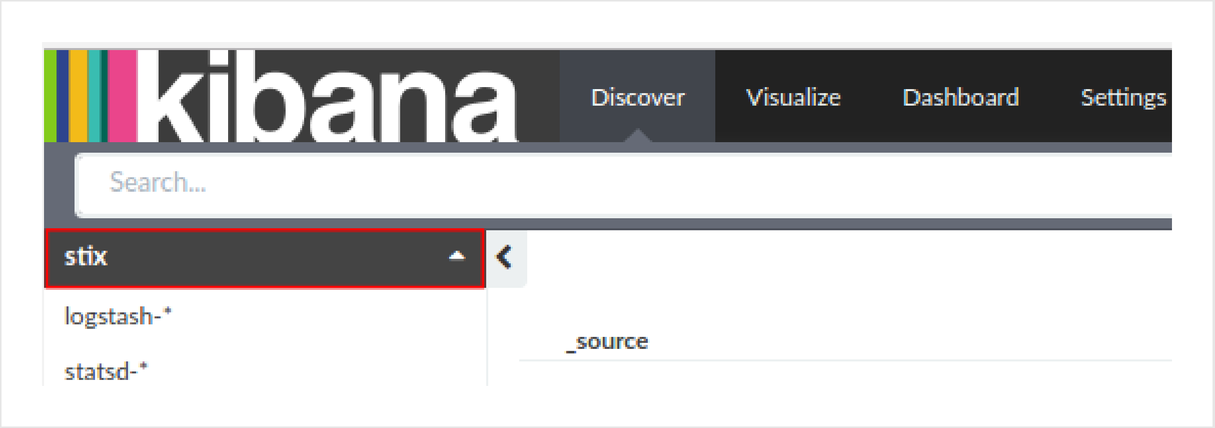

- On the main menu bar, select **Settings**:

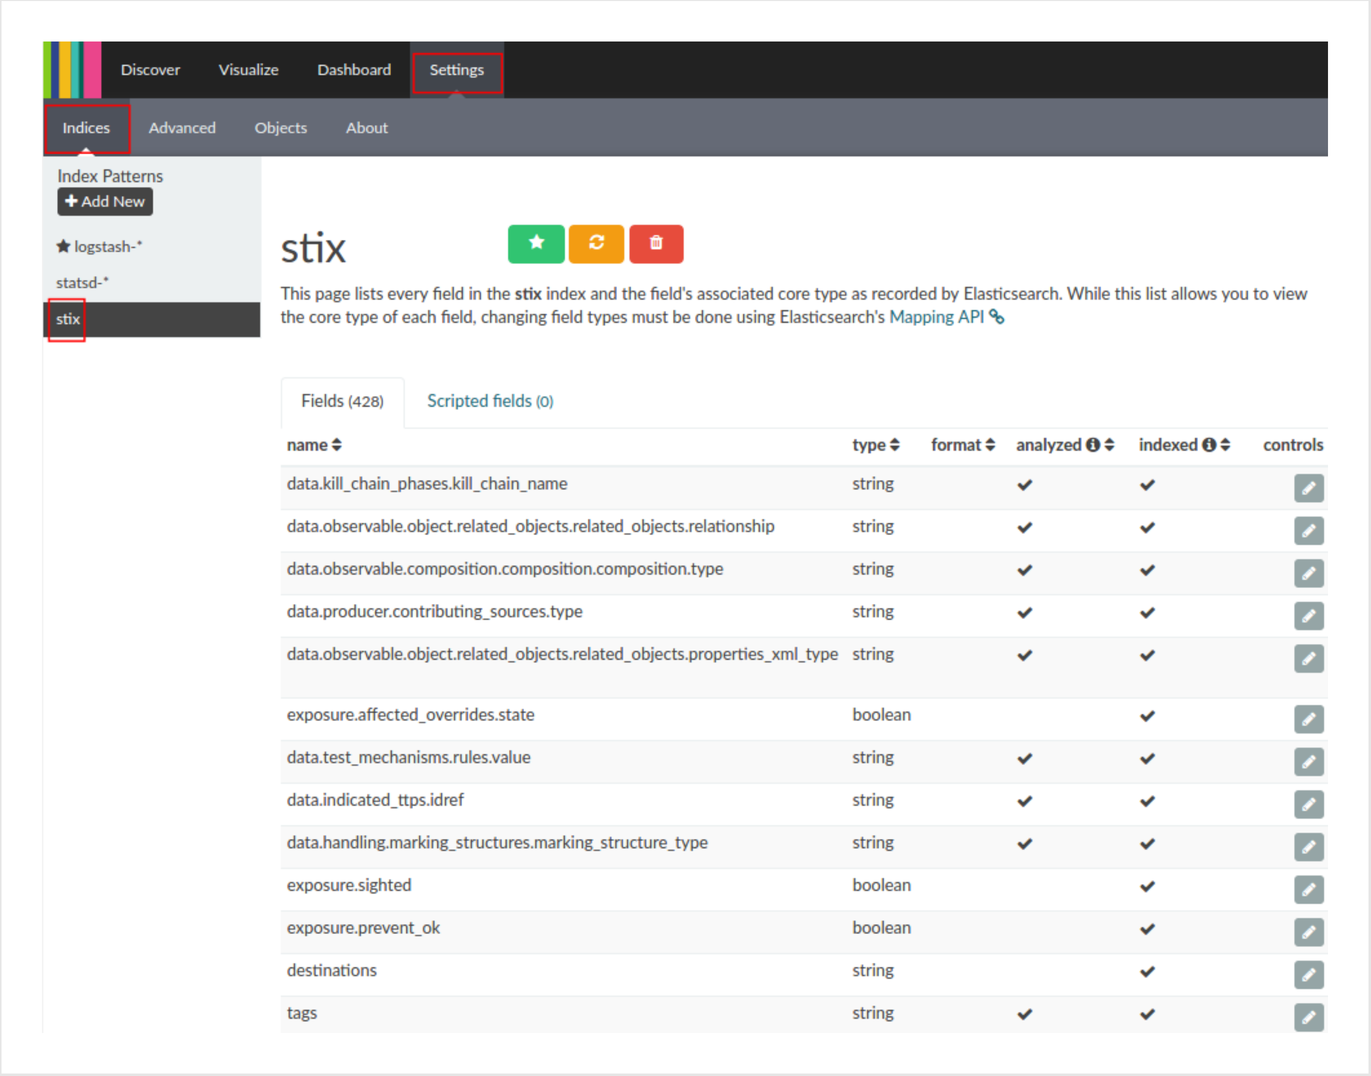

stix

This page lists every field in the **stix** index and the field's associated core type as recorded by Elasticsearch. While this list allows you to view the core type of each field, changing field types must be done using Elasticsearch's Mapping API

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

Search timeout

By default, Elasticsearch search queries that do not resolve time out after 20 seconds.
To set a different search timeout value, do the following:

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

Upload data files

Summary: 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.

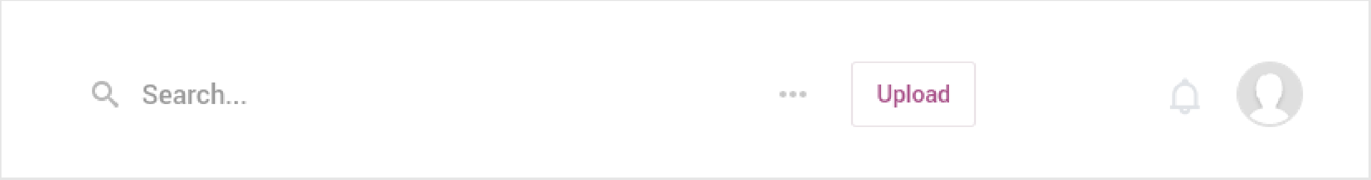

Upload data

- **Content type**: from the drop-down menu select a content type corresponding to the file format you are about to upload. The availanle 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

Summary: 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.

Within the platform, discovery rules work like configurable, specialized intel providers:

- 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 all matching data *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 the 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 all matching data for the rule *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.

ℹ

- When a rule is active, it is automatically executed every 15 minutes.
- 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**.

Alternatively:

- On the top navigation bar click ⚙ **> Rules > Discovery**.
- **Rules > Discovery** shows an overview of the existing discovery rules.
  You can sort the table view by column. 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:

- Go to **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*
  - **Active**: select or deselect this checkbox to enable or disable the rule.
- When you are done, 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

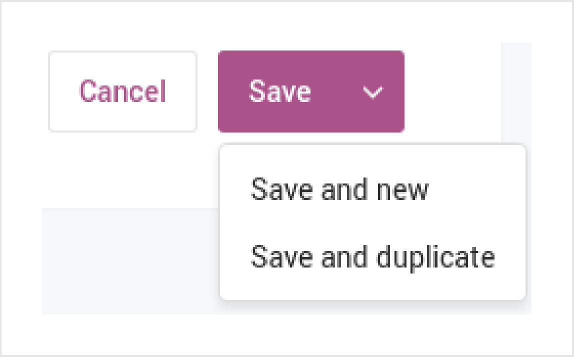

✔ On the forms, input fields marked with an asterisk are required.

To create a new discovery rule, do the following:

- Go to **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*
    - **Active**: select or deselect this checkbox to enable or disable the rule.

- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

Edit discovery rules

To edit a rule, do the following:

- Go to **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.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

Alternatively:

- Go to **Rules > Discovery**.
- On the rule overview, click the dotted menu 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.
- When you are done, click **Save** to store your changes, or **Cancel** to discard them.

> You can also edit the 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 all matching data for the rule *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:

- Go to **Rules > Discovery**.
- On the rule overview, click the dotted menu 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:

- Go to **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.

⟳ sent	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

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

Workspaces help you structure and organize your workload to keep it manageable and efficient.

You can use workspaces as collaborative environments where a group of analysts can zero in on selected potential threats to assess them; alternatively, you can use workspaces like containers to sort and manage related threats.

In a workspace the tools you need are ready at hand, so you can concentrate on your tasks.

Workspaces are organized in tabs to easily give you access to:

- Quick overviews
- Datasets
- Graphs to visualize threat relationships
- Any relevant files you may want to check
- Comments and feedback from other colleagues.

Workspace types

You can assign types to workspaces to clarify their purpose.

Type	Description
Generic	A generic workspace to collect structured, semi-structured or unstructured information.
Team	A team workspace to stress collaboration and knowledge sharing. You can use it to manage tasks and workload at team level. For example a team workspace allows you to organize and share information at team level, assign tasks to team members and keep track of progress.
Topic	A topic workspace helps you categorize, order and structure information. Use it to to gather threat information related to a specific research or investigation area. For example you can use it to focus on intelligence that is related to prevention, detection, or threat assessment operations.
Case	A case workspace is a structured container to organize intelligence on a case basis. For example a specific cyber attack, or suspicious activity originating from a limited region.

Access workspaces

To access workspaces, do the following:

- Sign in to the platform.
- On the left-hand navigation sidebar, click **Workspaces**.
- The **Workspaces** page displays an overview of the workspaces available to the currently signed in user.

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

Create a workspace

- In the **Workspaces** page, click the **Workspace** button.
- The **Workspaces > Create new** form page includes several input fields you can populate to define the new workspace.

✓ 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.	*B-R5RB Bloodbath*	
Type	Single choice drop-down menu	Defines the workspace type, the main purpose you're creating it for. Make an appropriate selection among the available options: **Generic**, **Team**, **Topic**, and **Case**.	*Team*	
Contact info	String, alphanum. [A-Z a-z] [0-9]	Enter here the details of a contact person.	*Mr. Smith*	
Collaborators	Multiple choice drop-down menu	Share the workspace with one or more collaborators. To do so, select the relevant people from the drop-down list. To remove a selection from the input field, click the ✖ icon corresponding to the item(s) you want to remove.	*Mr. White, Mr. Pink, Mr. Brown, Mr. Blonde*	

Field	**Type**	**Description**	**Example**	
Description	Text input field	An internal description for the workspace. It provides some high-level details about the purpose of the workspace. The content of this field is visible to workspace members only in the workspace **Overview** tab, under the **Short description** header.	*Only the workspace collaborators can see this description: we're doomed!*	
Public description	Text input field	Same as **Description**. The difference is that any content in this field is visible to all signed in users of the platform. The content is displayed in the workspace **Public page** tab, under the **Short description** header.	*All signed in platform users can see this description: everything is under control.*	
Analysis	Text input field	Work notes and analysis findings to provide more context about and insight into the workspace content and/or its purpose.	*Workspaces promote team collaboration and knowledge sharing.*	
Is Public	Checkbox	Default setting: deselected. Leave it deselected to keep it private and accessible only to the workspace collaborators. Select this checkbox to open up the workspace and make it public. An additional **Public page** tab becomes available, where all signed in platform users can view the workspace public content.	-	

After populating these fields, click **Add** to create the workspace.

Workspace types

Summary: 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.

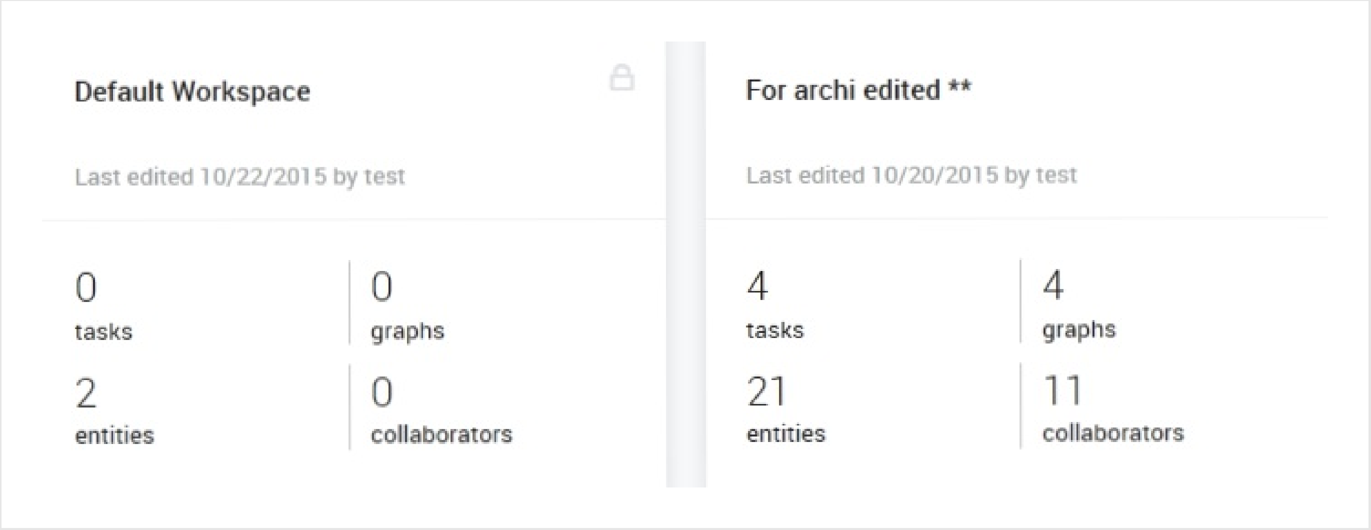

Default Workspace

Last edited 10/22/2015 by test

0 tasks
0 graphs
2 entities
0 collaborators

For archi edited **

Last edited 10/20/2015 by test

4 tasks
4 graphs
21 entities
11 collaborators

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

Summary: 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

Summary: 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

You can filter tasks to view only a sub-set of the available ones. You can choose to view only tasks created by and/or assigned to the current user, or only tasks in a specific status.

View tasks created by or assigned to the current user

- Click the **My Tasks** link.
- From the drop-down menu select one or more checkboxes to display the following task sub-sets:
  - **Created By Me**: the filter returns only the tasks created by the current user.
  - **Assigned To Me**: the filter returns only the tasks that are assigned to the current user.
  - Select both checkboxes to display all tasks created by *and* assigned to the current user.

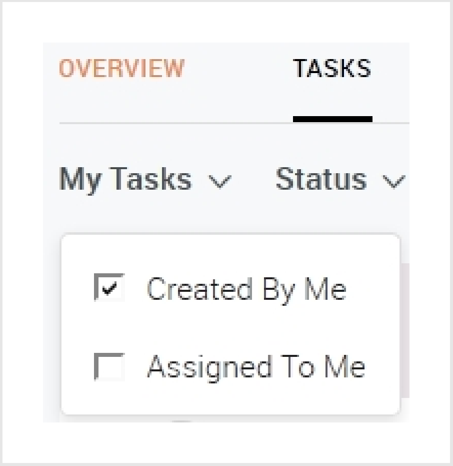

View tasks with a specific status

- Click the **Status** link.
- From the drop-down menu select one or more checkboxes to display only the tasks that satisfy the checked status criteria.

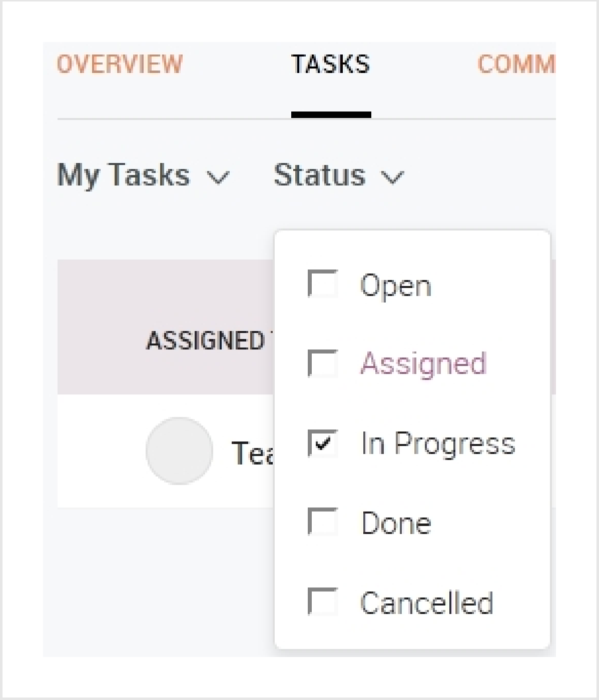

View task details

To view specific details about a task, click the corresponding row in the task overview. An overlay slides in from the side of the screen. It displays detailed task information in a flash-card format. You can review task details like:

- Task status
- Task owner/Assignee
- Creation date
- Any entities it references

You can read existing comments, or add new ones to provide additional information by clicking the **Add comment** link.

You can also act on a task. To do so, click **Action**. From the pop-up menu, choose the action you want to carry out:

Action	Description
Edit	Takes you to the **Tasks > Edit** section. Here you can modify any content about the task, for example following a change in scope or focus.
Change Due Date	Allows you to edit the scheduled deadline for the task.
Re-assign	Allows you to assign the task to a different owner than the current one.
Cancel Task	Allows you to cancel the task and remove it from the task list.

When you complete a task, click **Complete Task!** to confirm task completion.

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 decide whether to take action, for example when work on a task does not evolve as planned.

Status	Description
Open	The default status a task takes upon its creation.
Assigned	The task has been assigned to a workspace collaborator who owns it, but who has not started it, yet.
In Progress	The task owner has started working on the assigned task.
Done	The task has been completed.
Cancelled	The task has been canceled.

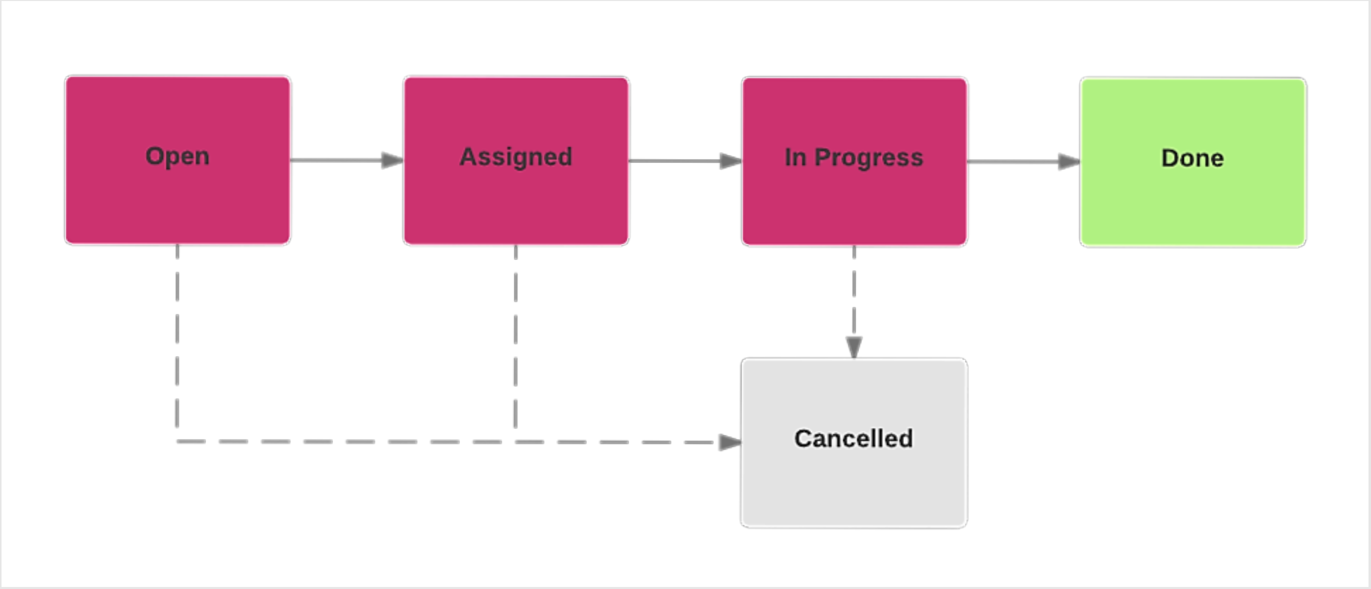

A standard task status flow

Edit tasks

To bypass the pop-up flash-card view and jump directly to the edit or delete options for a task, click the dotted menu icon on the row corresponding to the task you want to modify:

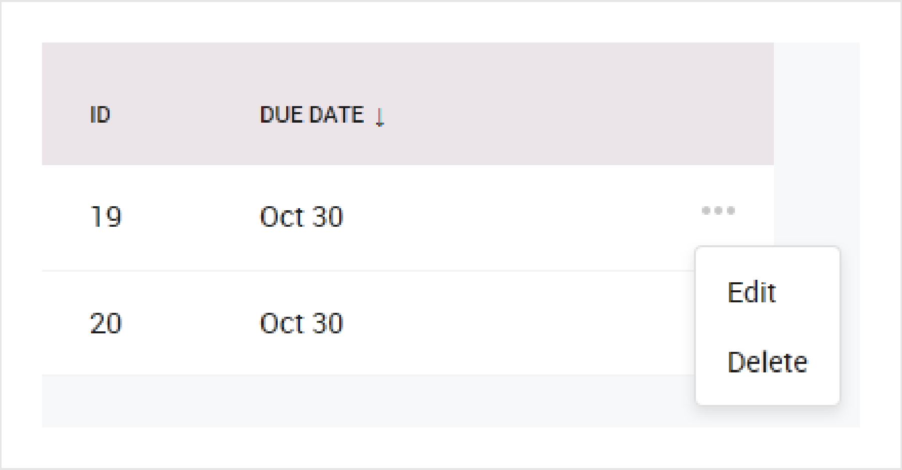

ID	DUE DATE ↓
19	Oct 30
20	Oct 30

Action	Description
Edit	Takes you to the task input form, where you can modify existing information and add new details about the task.
Change Due Date	Allows you to remove the task from the overview. A confirmation dialog prompt is displayed before deleting the task.

Workspace Comments tab

Summary: 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.

Last generated on Mar 8, 2017

Workspace Saved graphs tab

Summary: 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 dotted menu icon on the relevant graph tile:

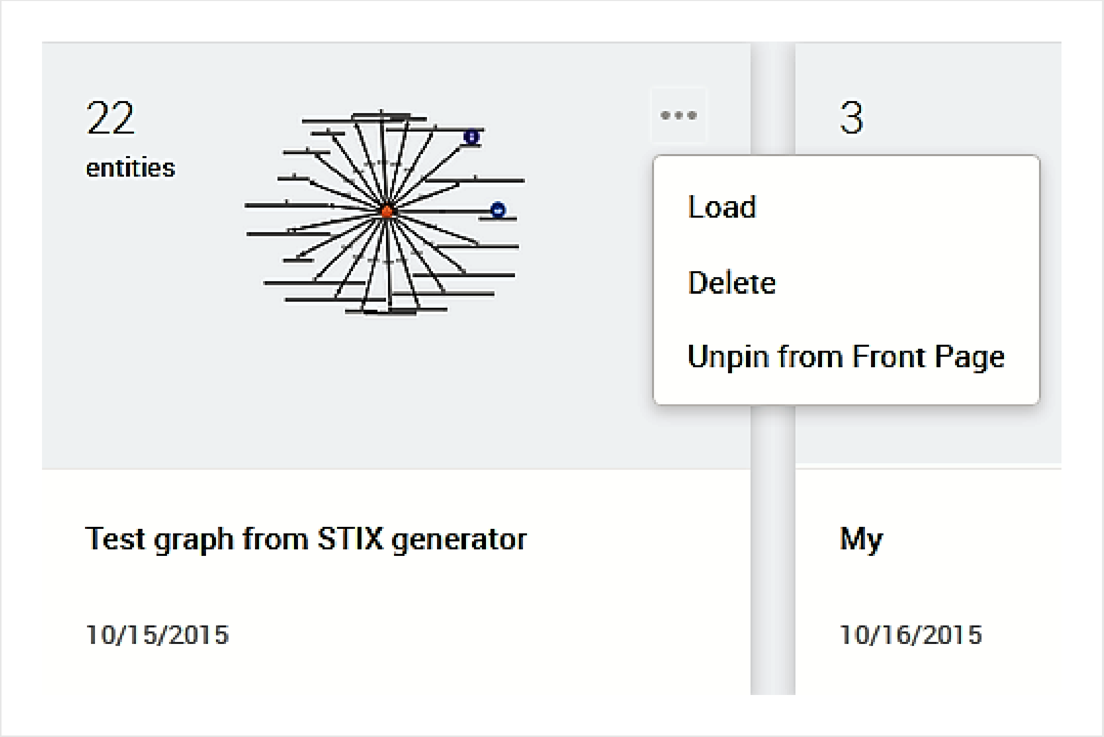

- **Load** allows you to load the selected graph to analyze entity relationships and references in an intuitively visual way.
- **Delete** purges and discards the graph and its content.
- **Unpin from Front Page** removes the graph quick access reference from the **Overview** tab, but keeps the graph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

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

Workspace Entities tab

Summary: 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

You can edit and modify entity data by choosing the desired option from the **Actions** pop-up menu on the bottom half of each entity detail pane tab. For example you can load it onto the graph to further examine it in a visual environment, you can create a task to take action and respond to the entity, or you can download it.

Grayed-out menu options are disabled in the current context, and are not available.

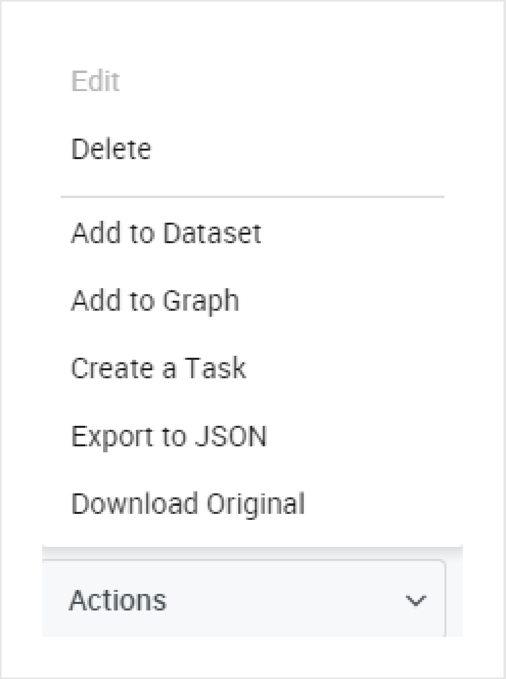

Select this menu option…	…to carry out this action
Edit	Makes the entity detail pane fields writable, so that you can edit and modify the entity.
Delete	Permanently removes the entity from the platform. If you confirm the entity deletion in the pop-up confirmation dialog, the operation is completed, and it cannot be undone.
Add to dataset	Adds the entity to the specified dataset(s).
Add to graph	Adds the selected entity to the graph to further examine it in a visual environment.
Create a Task	Makes the entity data actionable by creating a task, and by assigning it to a user.
Export to JSON	Exports the entity as a JSON object that can be saved, for example locally.
Download Original	Downloads the selected entity in its original data format.

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

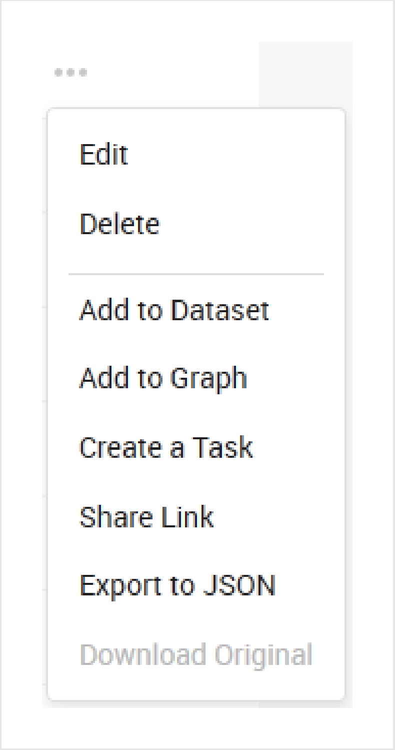

View entity details

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

- Go to the entity editor published entity overview.
- Click anywhere on a 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).

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

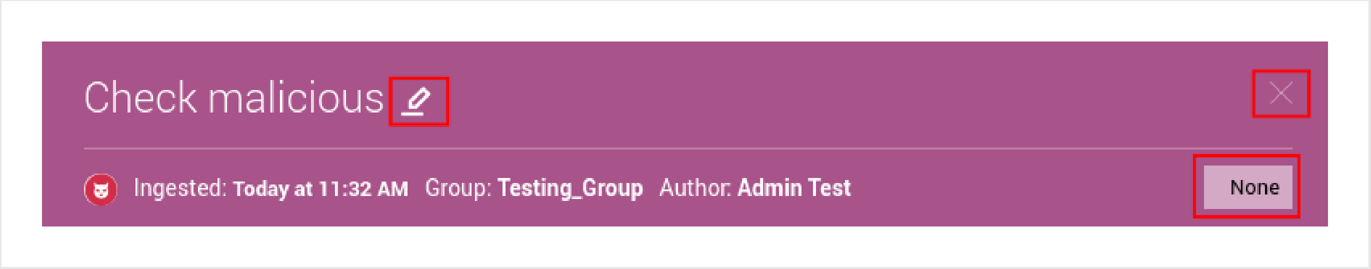

Check malicious

Ingested: Today at 11:32 AM Group: Testing_Group Author: Admin Test

None

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.

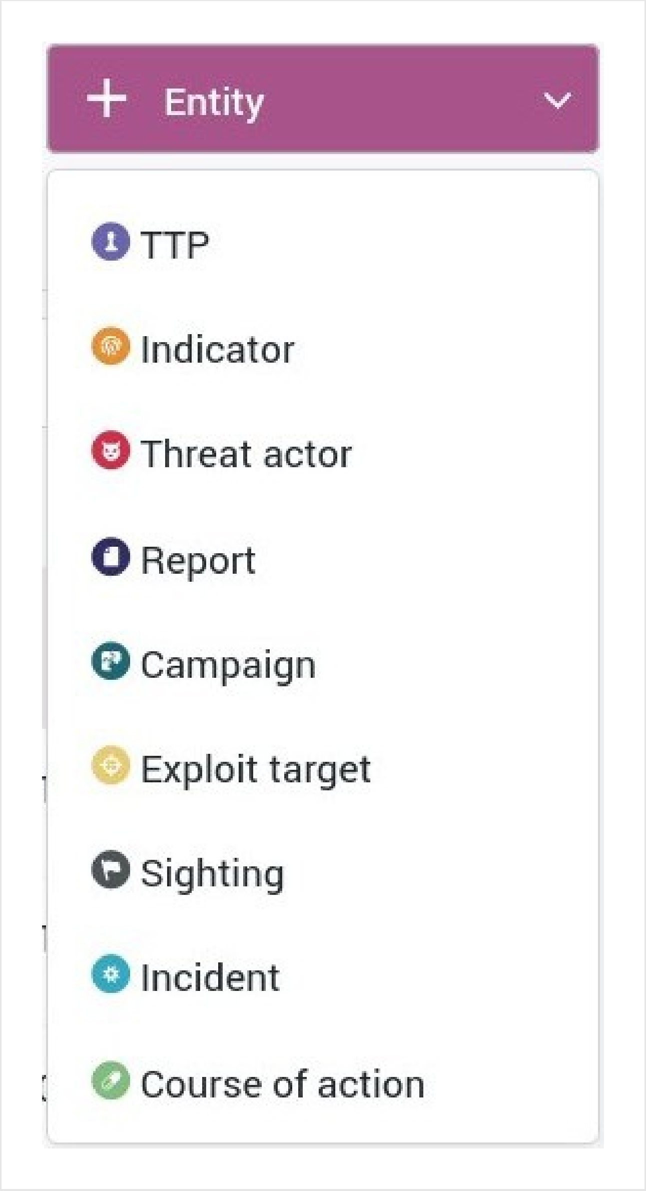

The drop-down menu discloses the available entity types you can create

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.

Entity type	Description
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**: when you are done, 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**. Draft entities associated with a workspace are available in the corresponding workspace under **Entities > Draft entities**.
  You can also click the drop-down arrow to save the current entry as a draft and create a new draft one right away — **Save as draft and new** — or to save the current entry and duplicate it — **Save as draft and duplicate** — for example to use it as a quick template for a new entry.
- **Publish**: when you are done, 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 entities**. Published entities associated with a workspace are available in the corresponding workspace under **Entities > Entities**.
  You can also click the drop-down arrow to save the current entry, publish it to the platform, and create a new entry one right away — **Publish and new** — or to publish the current entry and duplicate it — **Publish and duplicate** — for example to use it as a quick template for a new entry.

Filter entities

You can filter entities to zero in on specific sub-sets:

- 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.	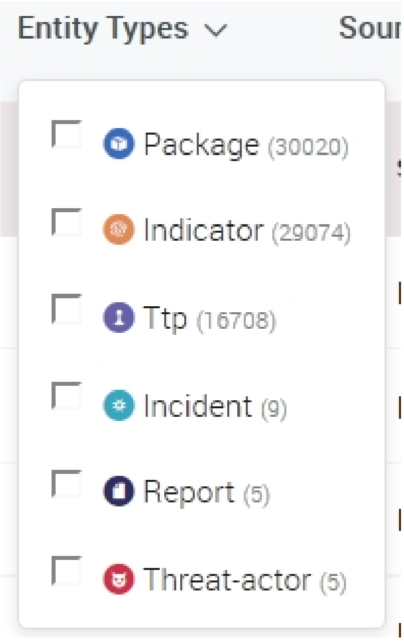 Entity Types Sour Package (30020) Indicator (29074) Ttp (16708) Incident (9) Report (5) Threat-actor (5)
Source Types	Filter entities by source/origin.	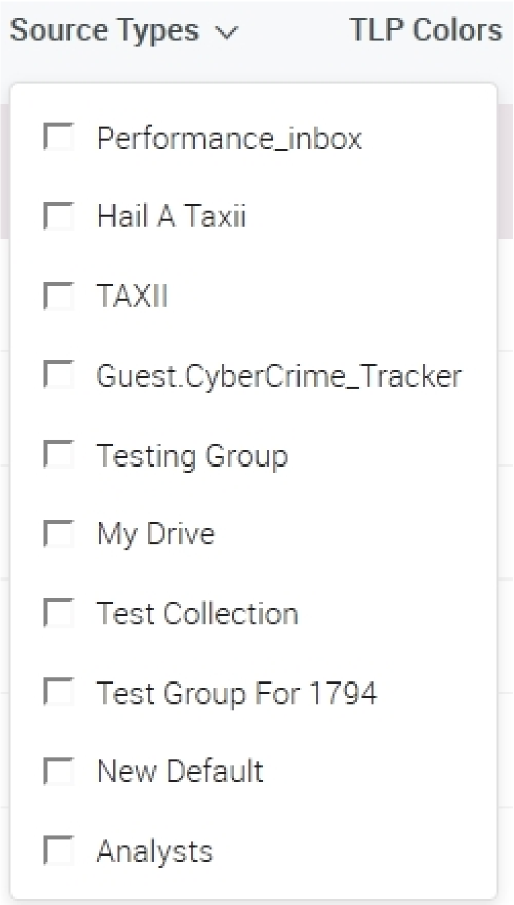 Source Types TLP Colors Performance_inbox Hail A Taxii TAXII Guest.CyberCrime_Tracker Testing Group My Drive Test Collection Test Group For 1794 New Default Analysts

Filter menu	Description	
TLP Colors	Filter entities by Traffic Light Protocol color code.	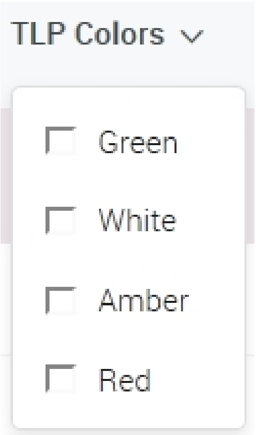
Date	Filter entities included in a date range.	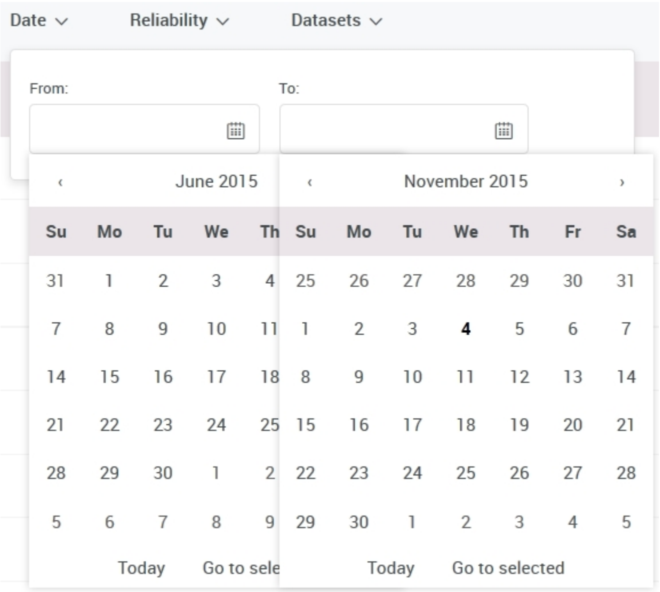 Date Reliability Datasets From: To: June 2015 November 2015 Su Mo Tu We Th Fr Sa Today Go to selected
Reliability	Filter entities based on their reliability index.	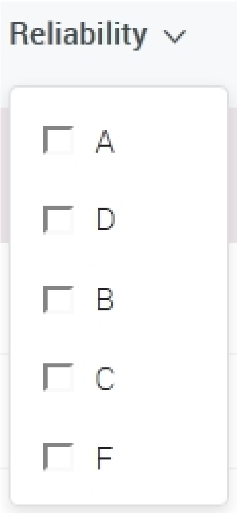

Filter menu	Description	
Datasets	Filter entities based on the dataset they belong to.	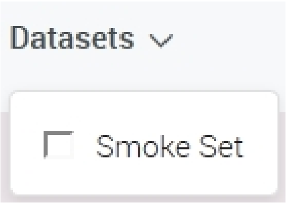

Entity reliability

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

Summary: 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.

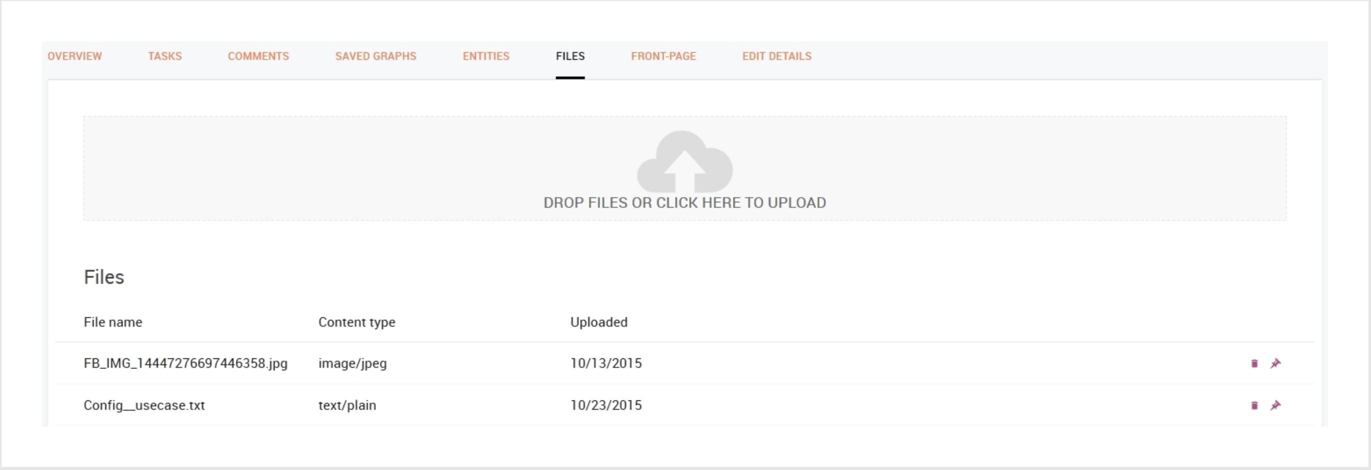

Files

File name	Content type	Uploaded
FB_IMG_14447276697446358.jpg	image/jpeg	10/13/2015
Config__usecase.txt	text/plain	10/23/2015

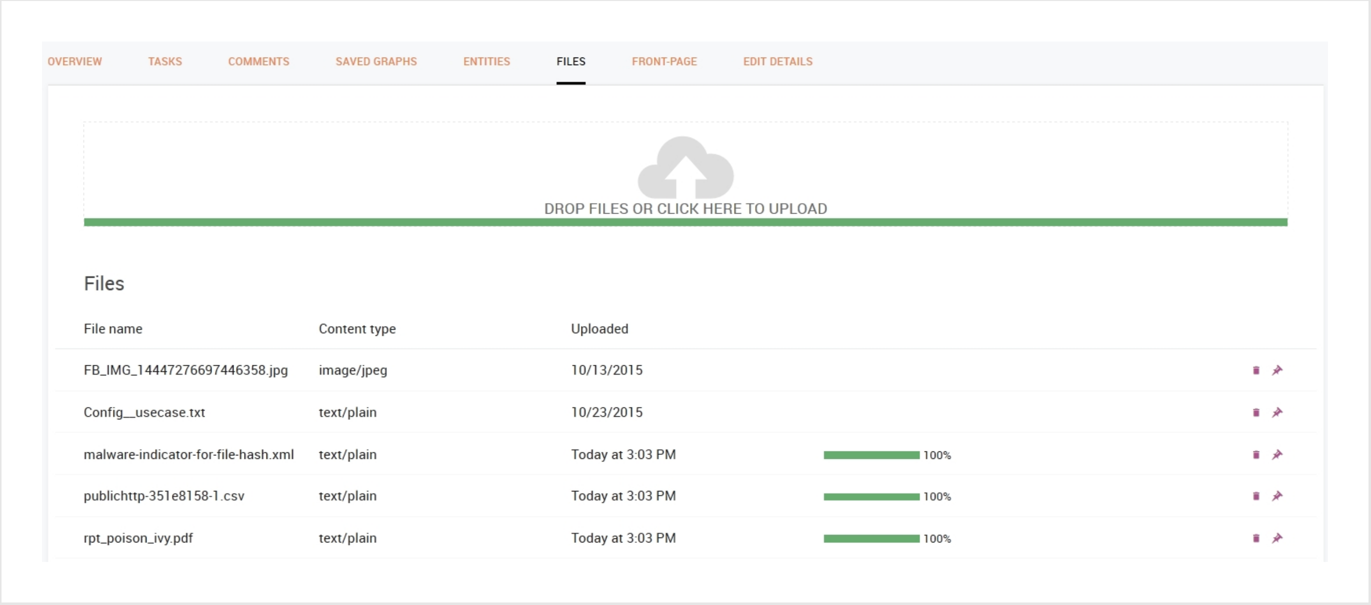

Files

File name	Content type	Uploaded	
FB_IMG_14447276697446358.jpg	image/jpeg	10/13/2015	
Config__usecase.txt	text/plain	10/23/2015	
malware-indicator-for-file-hash.xml	text/plain	Today at 3:03 PM	100%
publichttp-351e8158-1.csv	text/plain	Today at 3:03 PM	100%
rpt_poison_ivy.pdf	text/plain	Today at 3:03 PM	100%

Workspace Edit details tab

Summary: 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

Summary: 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. The graph is a canvas where you can add and remove entities, examine and manipulate them to gain insights, look for relationships, and extract bits of information from entities for further investigation.

The graph is a quick and intuitive way to investigate complex relationships using different layouts, the histogram filtering options, and the timebar.

Add entities to the graph

You can quickly load entities onto the graph from almost anywhere in the platform.
For example:

- On the top navigation bar click **Browse**, **Discovery**, **Exposure**, or **Workspaces > <workspace_name> > Entities**.
  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
  By default, the lower half of the entity detail pane includes an **Actions** menu.
- On the pop-up menu select **Actions > Add to graph** to load the entity onto the graph.

> ℹ The **Actions > Add to graph** menu option is available for published entities and observables.
> It is not available for draft entities.

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***.
  The 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**.

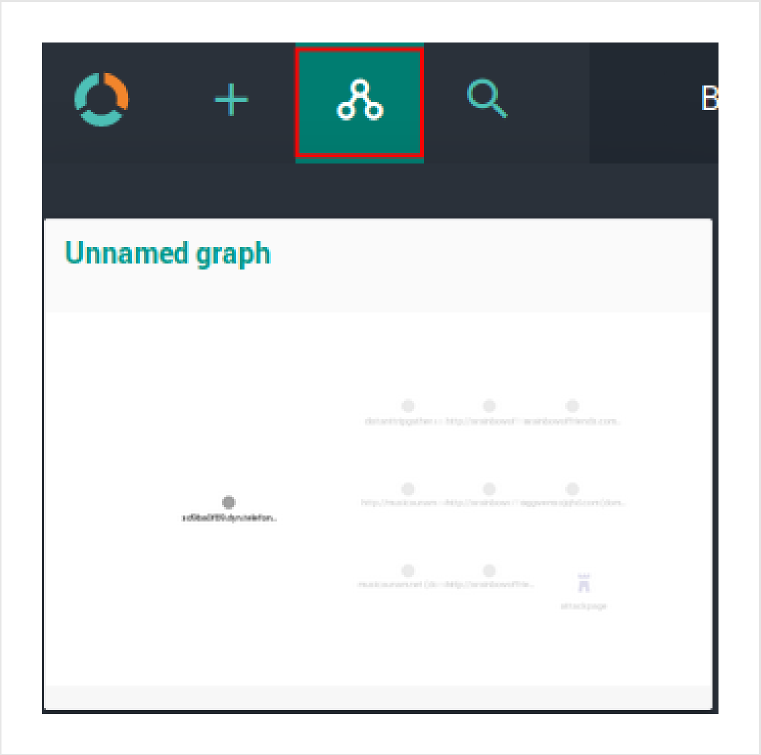

Open the graph from the Saved graphs section in a workspace

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

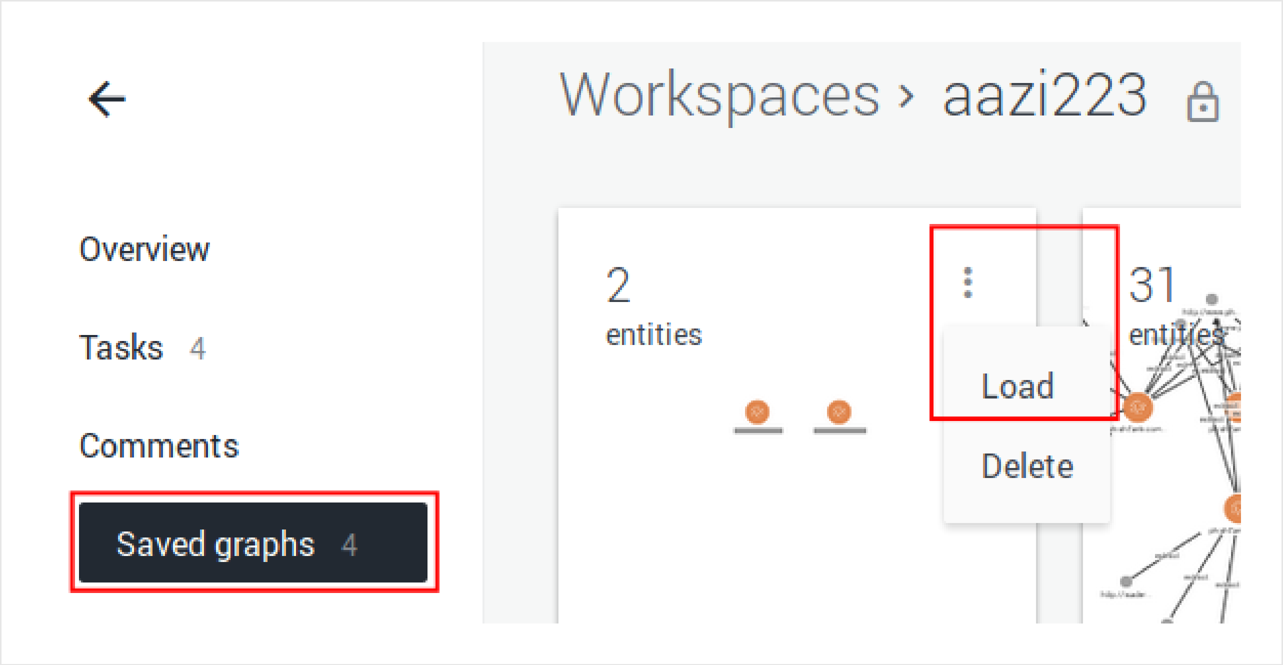

Open the graph from the Neighborhood tab on the entity detail pane

- 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.

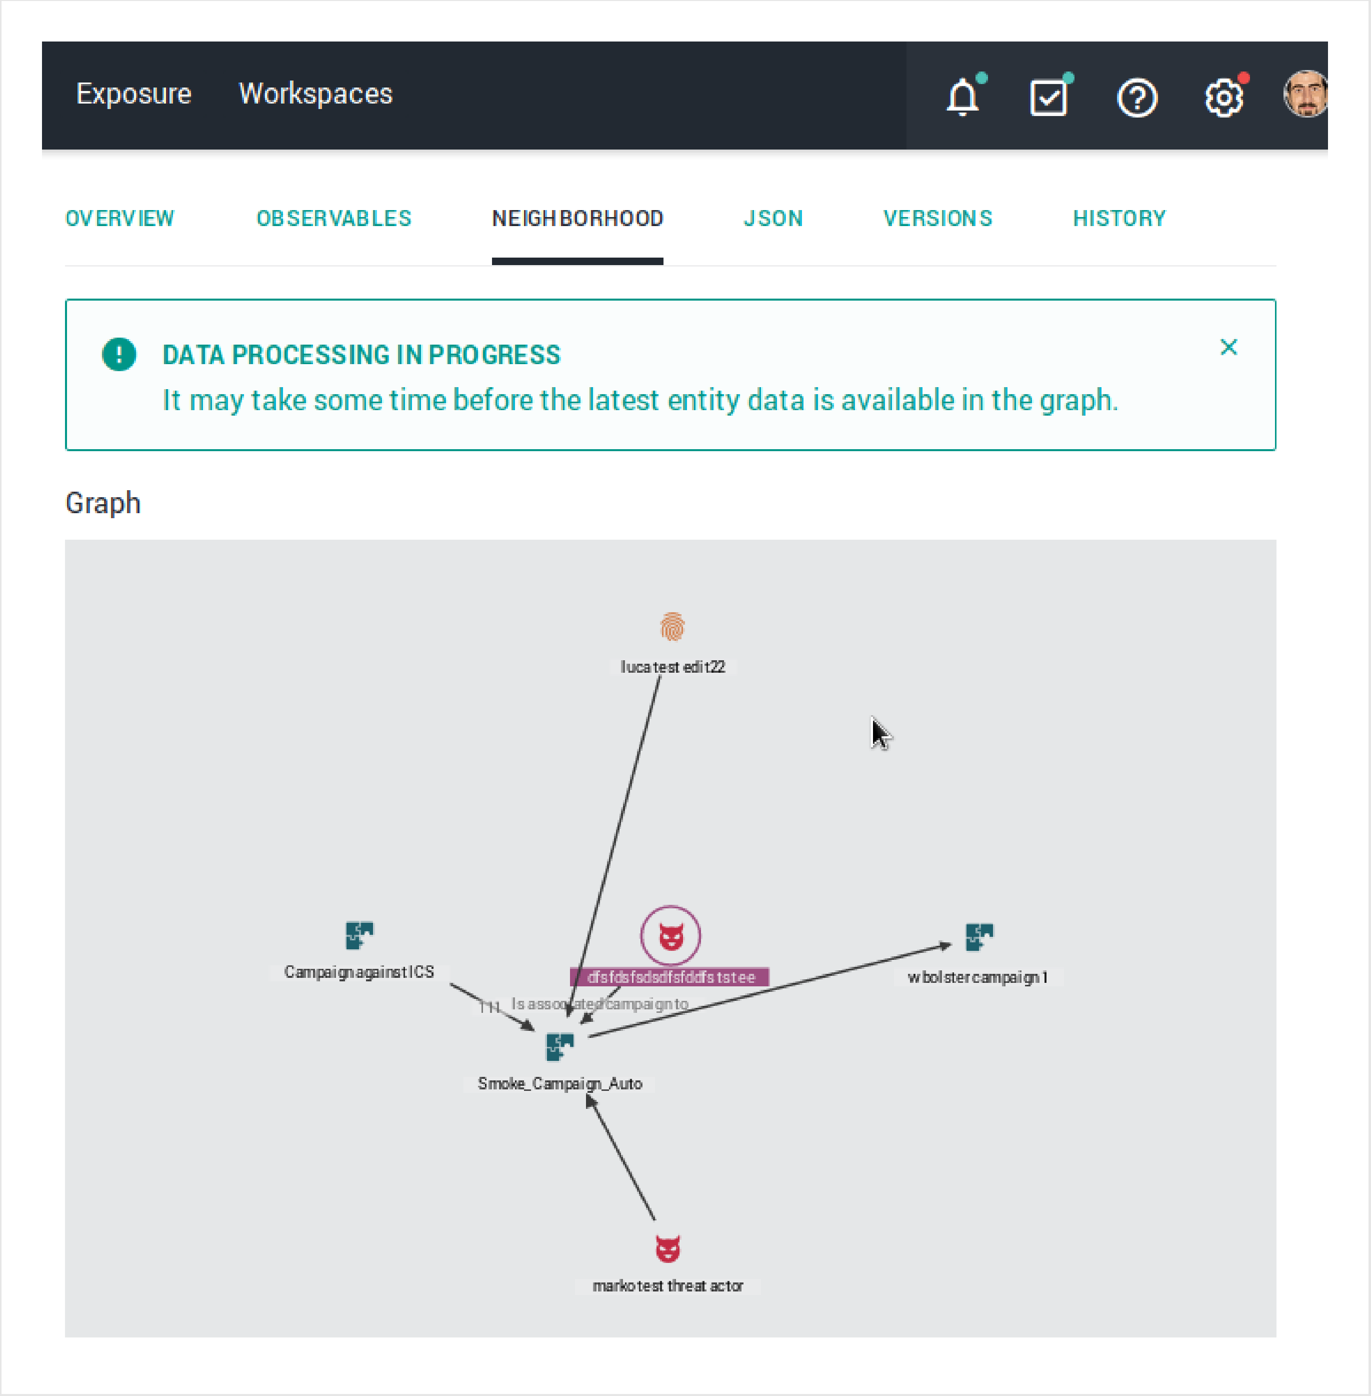

DATA PROCESSING IN PROGRESS

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

Graph

Examine entities on the graph

To start working with an entity on the graph, right-click it to display the available options in the context menu. Greyed 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.

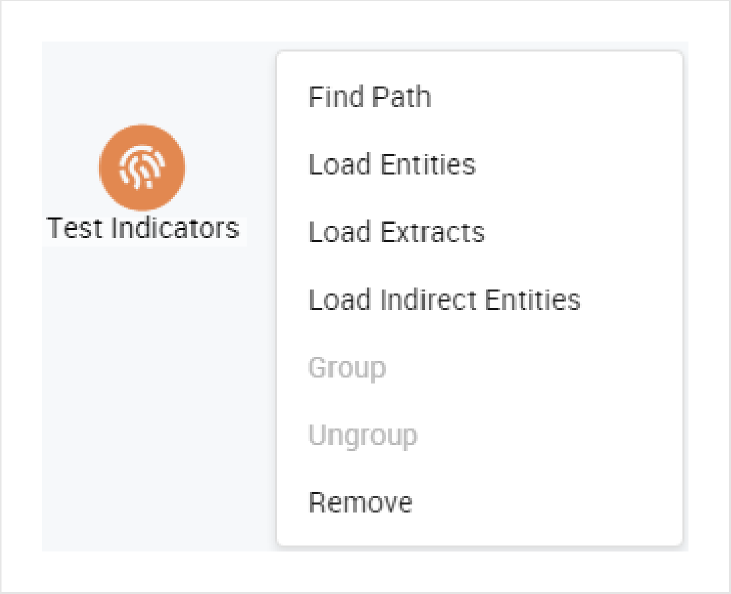

- **Load Entities**: click this option to to display on the graph any entities that are linked to the selected one, for example because they share one or more IP addresses, or target email addresses, and so on.

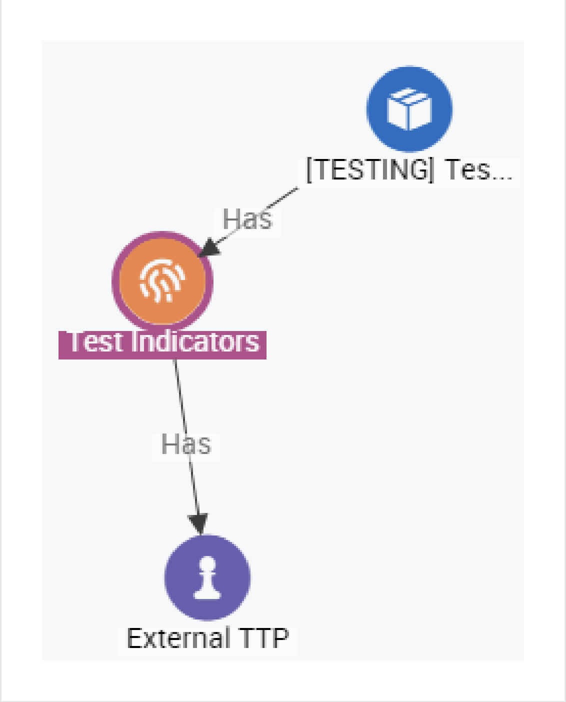

- **Load Extracts**: click this option to to display on the graph any extracts associated with the selected entity, so bits of information like an IP address, or an email address related to the entity.

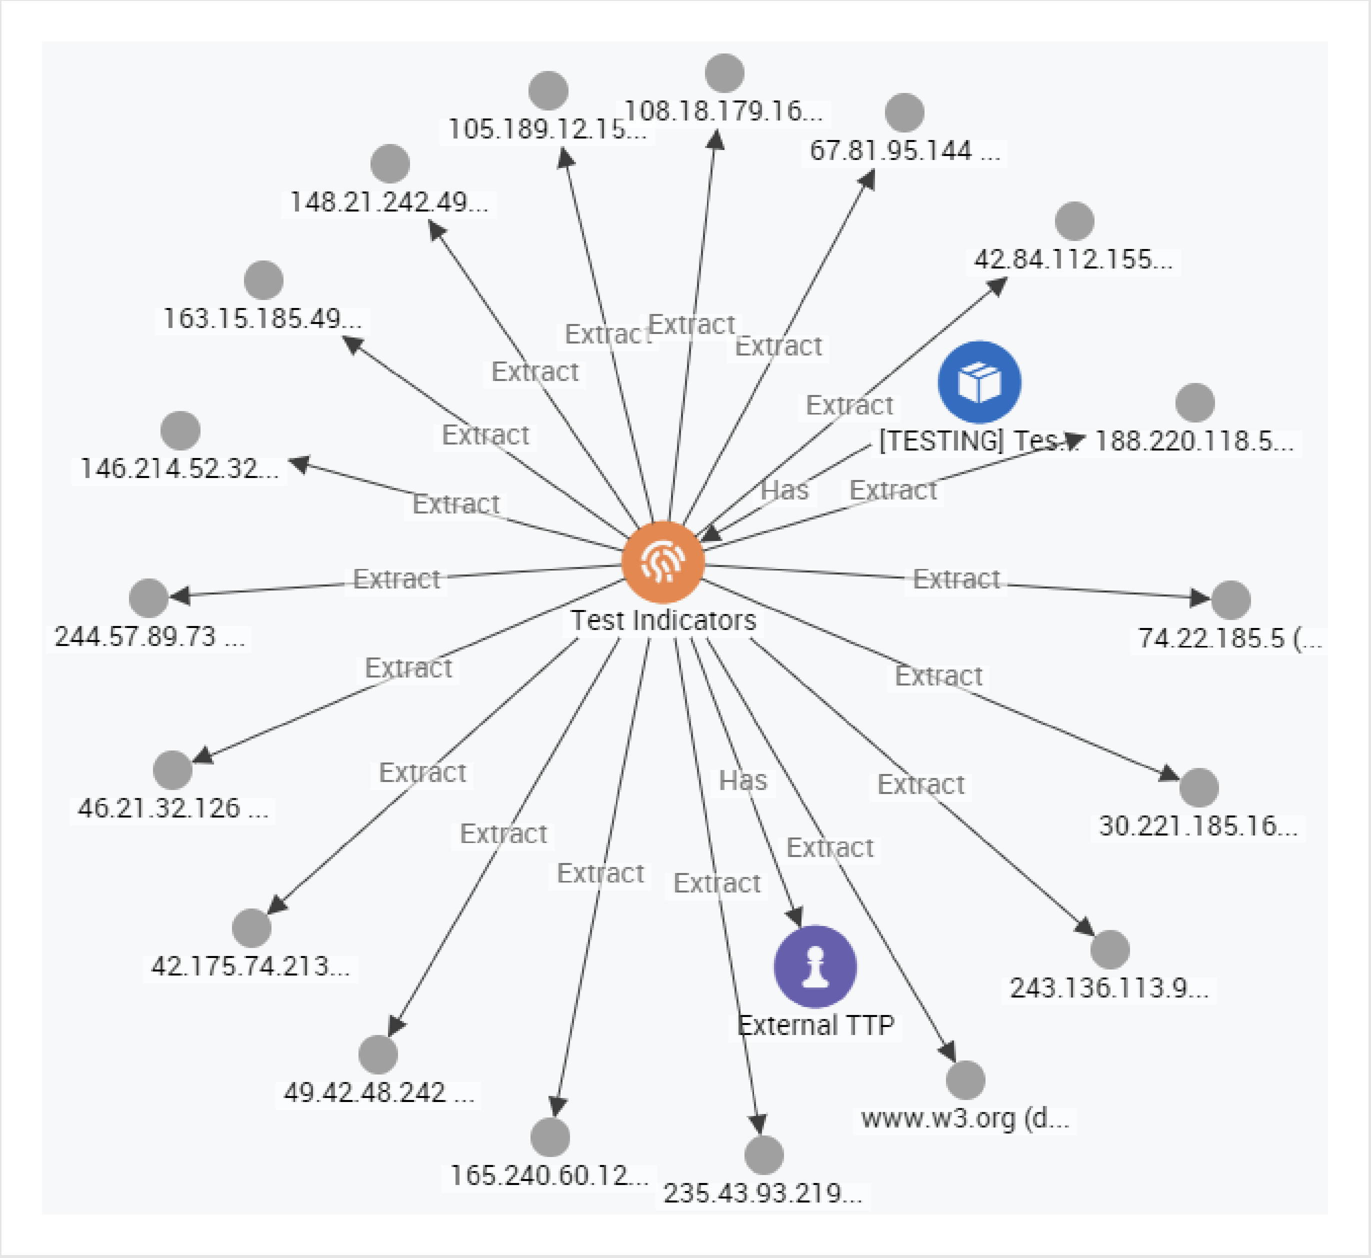

- **Load Indirect Entities**: click this option to to display on the graph any entities that are indirectly linked to the selected one, so through one or more intermediate nodes.

You can also select, group, ungroup, and remove loaded entities:

- **Group** allows you to group together entities that you want to handle together: select the entities by dragging the mouse, right-click one of the entities included in the selection, and then click **Group**.
- **Ungroup** undoes entity grouping by restoring them as separate entities on the graph.
- **Remove** allows you to remove the selected entity from the graph. It works also on 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 bar, click the icon corresponding to the desired layout to automatically rearrange the view accordingly.

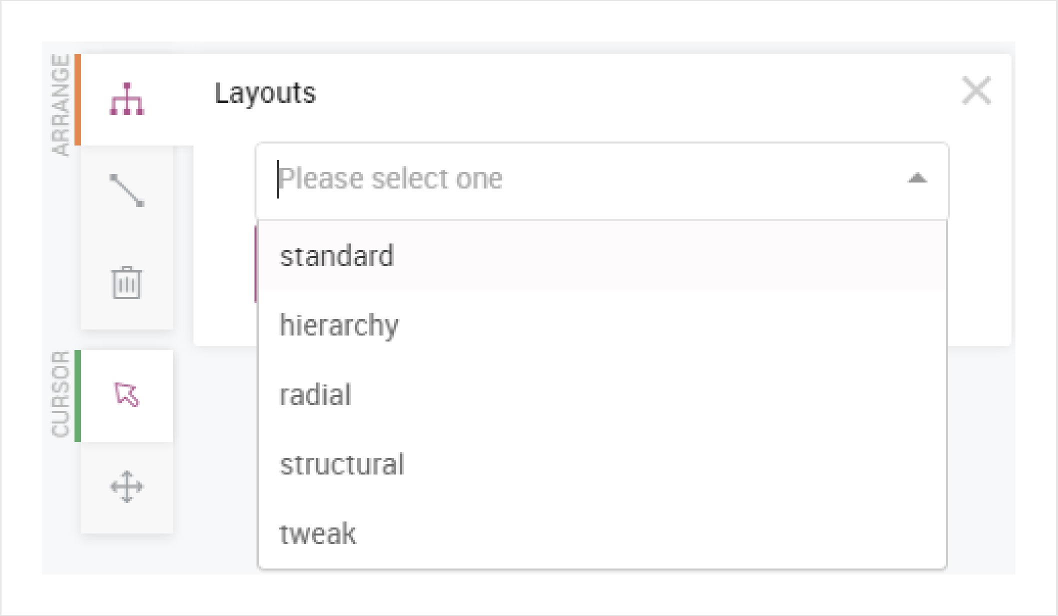

Layout type	Description
Standard	In the standard layout, links on the graph are a consistent length. Nodes and edges overlap as little as possible, and they are evenly distributed on the graph surface. Its goal is consistency and simplicity; it is a catch-all for any kind of data and any dataset size, especially when you are looking for patterns and symmetries.
Hierarchy	It is a family tree 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 as the dependency chain grows. This layout is the best option when dealing 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 adapts itself as links are created and destroyed, allowing the viewer to see clearly 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, 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:

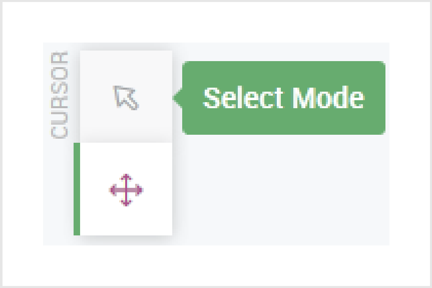

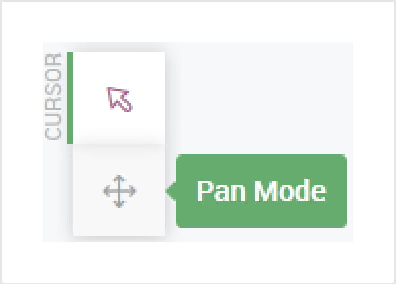

- Click **Select mode** to select and deselect one or multiple entities on the graph. You can group, ungroup, or remove the selected entities, as well as further examine them using the context menu options.
- Click **Pan mode** to move up and down, as well as left and right on the graph canvas. This is helpful when working on a graph containing a large amount of entities.
- 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.

Datasets

Summary: 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 left-hand navigation sidebar, click **Datasets**.
- Click the ✚ **Dataset** button.

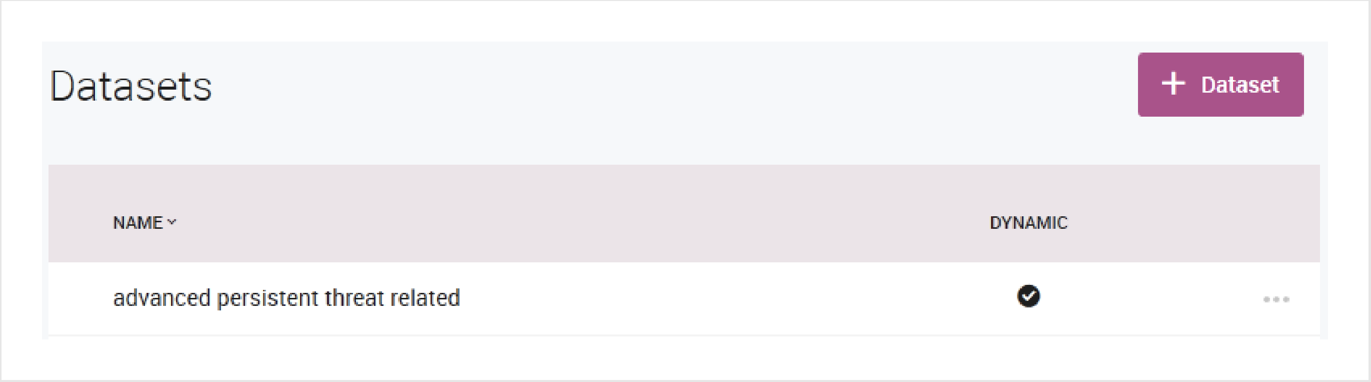

Datasets

NAME	DYNAMIC
advanced persistent threat related	✓

- In **Datasets > Create Dataset**, specify a name for the new dataset.
- If you want to create a dynamic dataset, click the **Is Dynamic** checkbox, and then specify a valid query string in **Search Query**. You can build your query using **Elasticsearch query syntax** `(https://www.elastic.co/guide/en/elasticsearch/reference/current/full-text-queries.html)`.
- 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

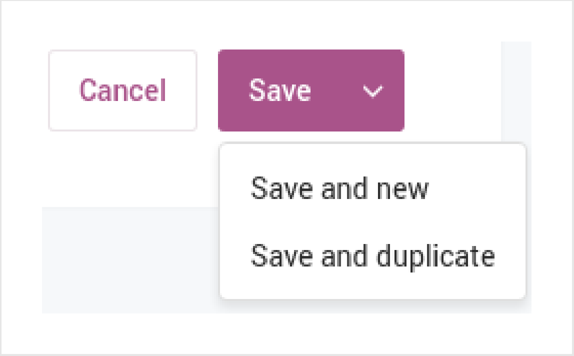

Edit or delete a dataset

To edit or delete an existing dataset, do the following:

- On the left-hand navigation sidebar, click **Datasets**.
- On the **Datasets** page, click the click the dotted menu icon on the row corresponding to the dataset you want to modify.

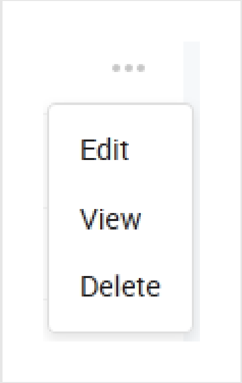

  - Choose **Edit** to rename an existing dataset, or to enable/disable an ad-hoc search query by toggling the **Is Dynamic** option.
  - Choose **View** to inspect the content of the dataset. It shows an overview where you can review the dataset entities, and you can manipulate them through option menus and feature sets comparable to the ones available in the workspaces and in the graph view.
  - Choose **Delete** to remove a dataset from the platform.

Tasks

Summary: 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, do the following:

- On the left-hand navigation sidebar, click **Datasets**.
- Click the ✚ **Task** button.

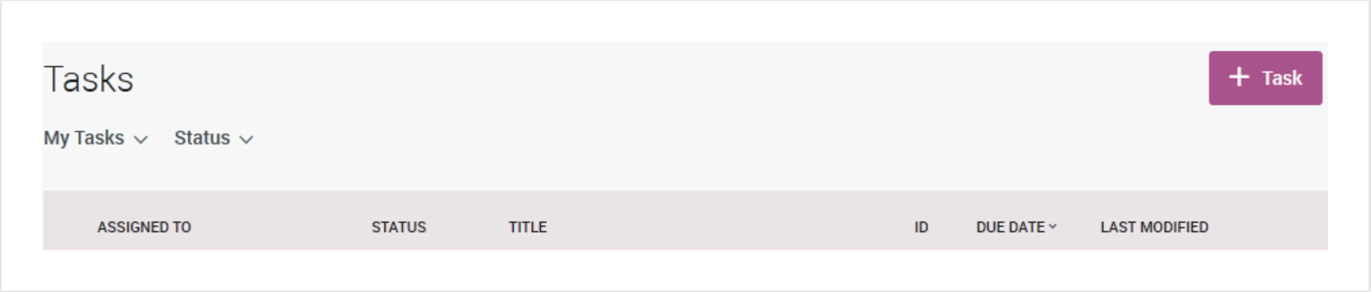

- In **Tasks > Create**, specify a name for the new dataset.
  On the forms, input fields marked with an asterisk are required.

Field name	Description
Name	The owner of the task. It should be descriptive and easy to remember.
Short description	A short description of the task to provide more context and.or details.
Guidance angle	Guidelines and recommendations to instruct the assignee, pointers to any relevant reference or context, so that the assignee is aware of any requirements or expectations related to the task.
Assigned to	The owner of the task. From the drop-doen menu, select the collaborator you want to assign the task to.
Due date (UTC)	Deadline: the scheduled completion date for the task.
Workspaces	From the drop-down menu select one or more workspaces to associate the task to.
Stakeholders	From the drop-down menu select one or more stakeholders sponsoring the task, for example a team leader or a project manager.

Field name	Description
Entities	Click the ✚ **more** link to add one or more entities to the task. The selected entities are the object of the task activites. Type a search string into the search field, click the **Search** button, and then select one or more entries from the search result list.

- 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 right away a new item of the same type; for example, a dataset, a feed, a rule, 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.

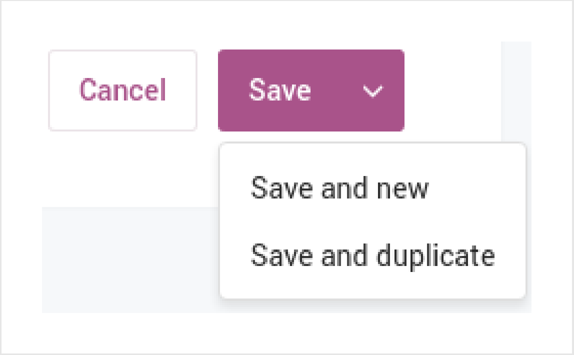

View tasks

You can filter tasks to view only a sub-set of the available ones. You can choose to view only tasks created by and/or assigned to the current user, or only tasks in a specific status.

View tasks created by or assigned to the current user

- Click the **My Tasks** link.
- From the drop-down menu select one or more checkboxes to display the following task sub-sets:
  - **Created By Me**: the filter returns only the tasks created by the current user.
  - **Assigned To Me**: the filter returns only the tasks that are assigned to the current user.
  - Select both checkboxes to display all tasks created by *and* assigned to the current user.

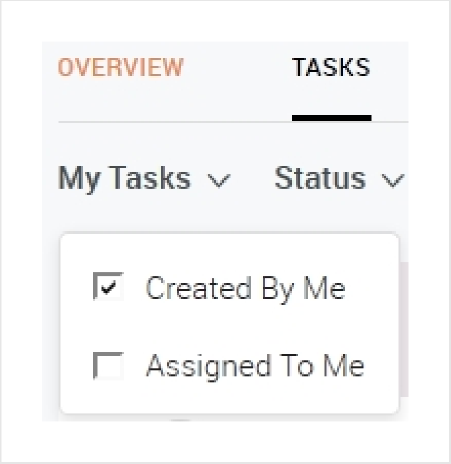

View tasks with a specific status

- Click the **Status** link.
- From the drop-down menu select one or more checkboxes to display only the tasks that satisfy the checked status criteria.

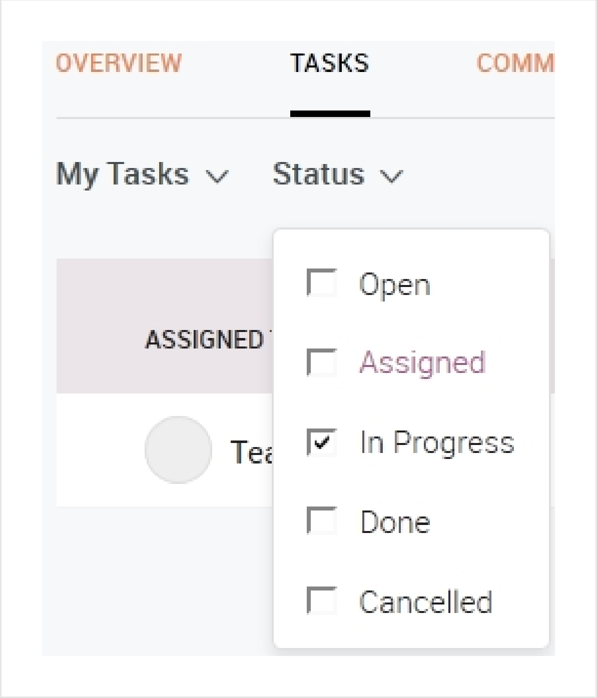

View task details

To view specific details about a task, click the corresponding row in the task overview. An overlay slides in from the side of the screen. It displays detailed task information in a flash-card format. You can review task details like:

- Task status
- Task owner/Assignee
- Creation date
- Any entities it references

You can read existing comments, or add new ones to provide additional information by clicking the **Add comment** link.

You can also act on a task. To do so, click **Action**. From the pop-up menu, choose the action you want to carry out:

Action	**Description**
Edit	Takes you to the **Tasks > Edit** section. Here you can modify any content about the task, for example following a change in scope or focus.
Change Due Date	Allows you to edit the scheduled deadline for the task.
Re-assign	Allows you to assign the task to a different owner than the current one.
Cancel Task	Allows you to cancel the task and remove it from the task list.

When you complete a task, click **Complete Task!** to confirm task completion.

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 decide whether to take action, for example when work on a task does not evolve as planned.

Status	**Description**
Open	The default status a task takes upon its creation.
Assigned	The task has been assigned to a workspace collaborator who owns it, but who has not started it, yet.
In Progress	The task owner has started working on the assigned task.
Done	The task has been completed.
Cancelled	The task has been canceled.

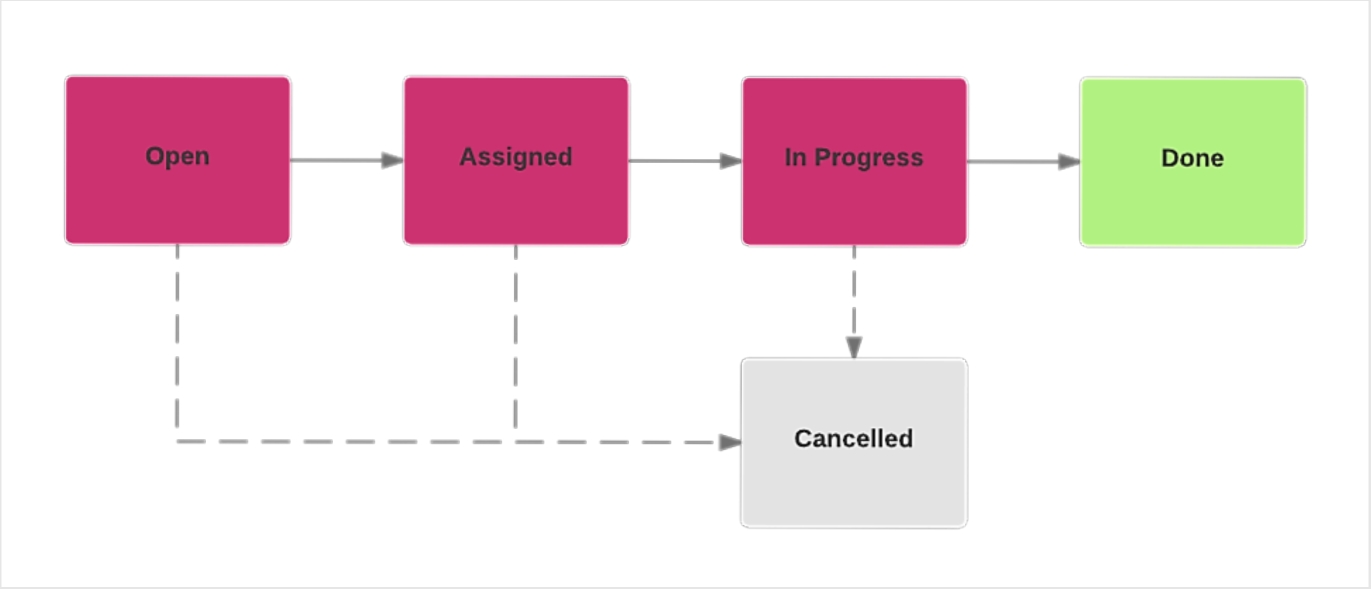

A standard task status flow

Editor

Summary: 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 > **.

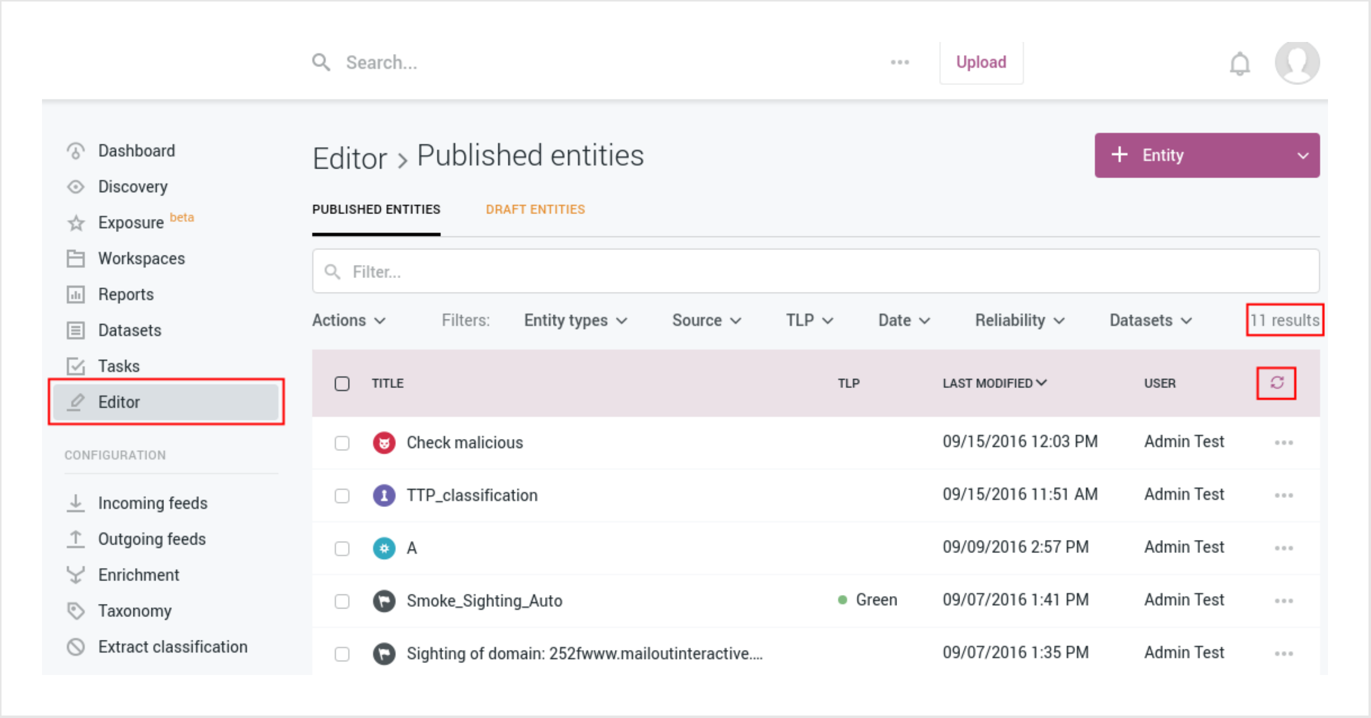

☐	TITLE	TLP	LAST MODIFIED	USER
☐	Check malicious		09/15/2016 12:03 PM	Admin Test
☐	TTP_classification		09/15/2016 11:51 AM	Admin Test
☐	A		09/09/2016 2:57 PM	Admin Test
☐	Smoke_Sighting_Auto	● Green	09/07/2016 1:41 PM	Admin Test
☐	Sighting of domain: 252fwww.mailoutinteractive....		09/07/2016 1:35 PM	Admin Test

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 a 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).

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

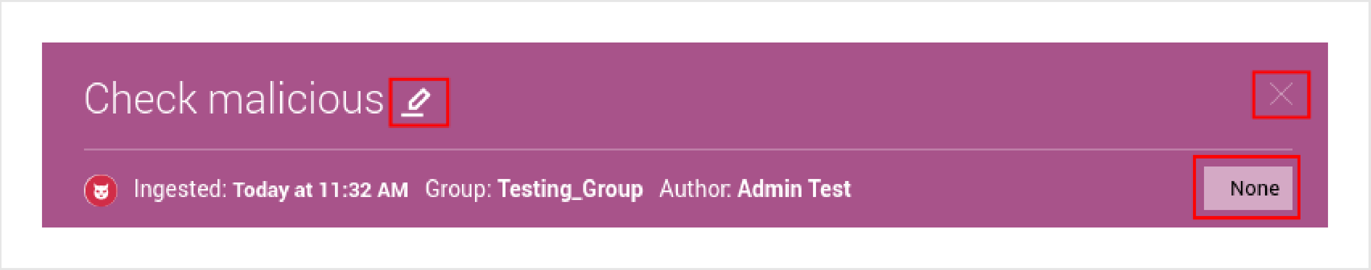
Check malicious

Ingested: Today at 11:32 AM Group: Testing_Group Author: Admin Test

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:

Title

The name of the entity as shown also on the detail pane header section.

Confidence

it flags the estimated level of confidence concerning the maliciousness of the (potential) threat the entity represents.
Allowed values:

- **Unknown**

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

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.

To manage tags and their parent-child relationships, go to **Taxonomy**.

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

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 reports 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 intel 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 **System > Groups**.

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.
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).`

Exposure

Exposed

Exposed entities are ingested and processed. However, their intelligence value is not used to produce any 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. In other words, nobody is doing anything with this information.

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.

Extracts

Type

The type of extract associated with the entity.

Extract

The value of the extract, for example an IP address, a domain, an email address, a city name, and so on.

State

Flags the estimated likeness of the extract to pose a threat. An extract threat potential can be malicious, safe or unknown. Extract rules help automate extract flagging.
To override the current extract state, click the dotted menu icon, and then select an action from the drop-down menu:

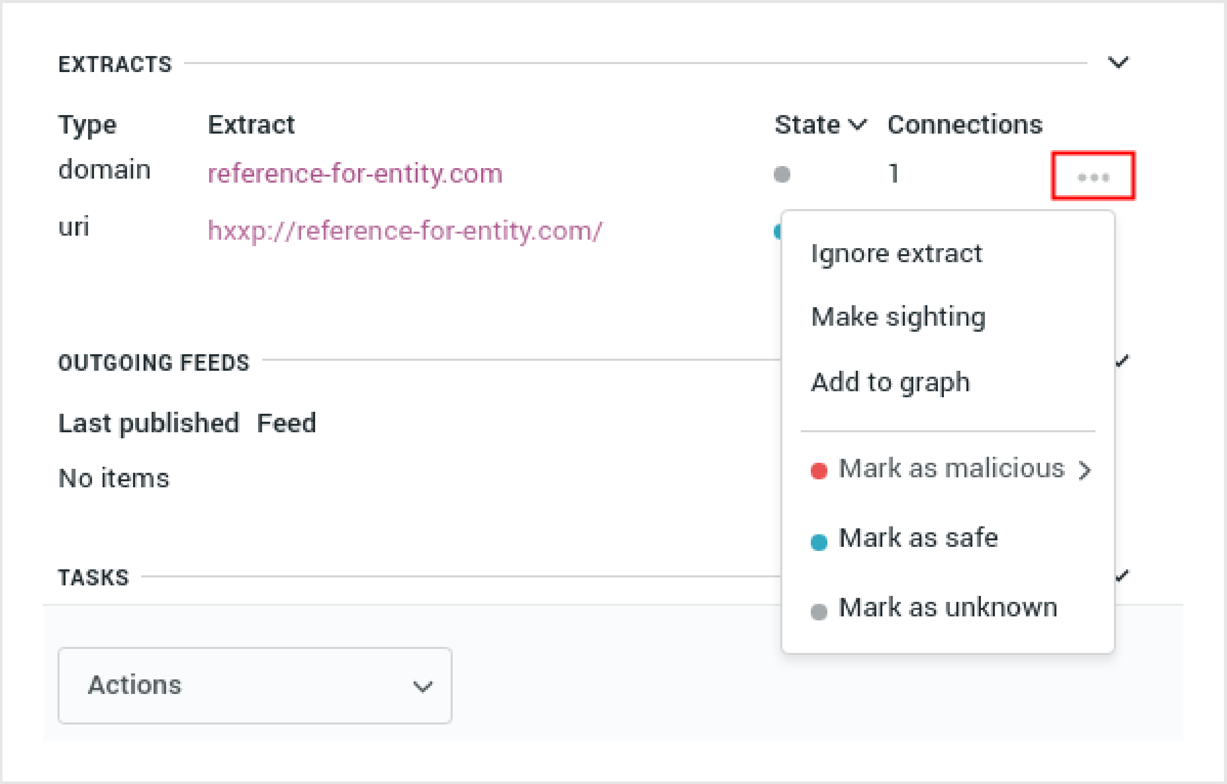

Connections

The number of connections the extract has or may have with other extracts or other entities in the platform.

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.

Tasks

Actionable user tasks associated with the entity are listed here You can create tasks and assign them to yourself or to other users to request follow-ups, for example further investigation or a call for action.

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 extracts related to the entity. Enrichers poll external data sources; enricher rules process the data and create the relayionships between entities and retrieved extracts.

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

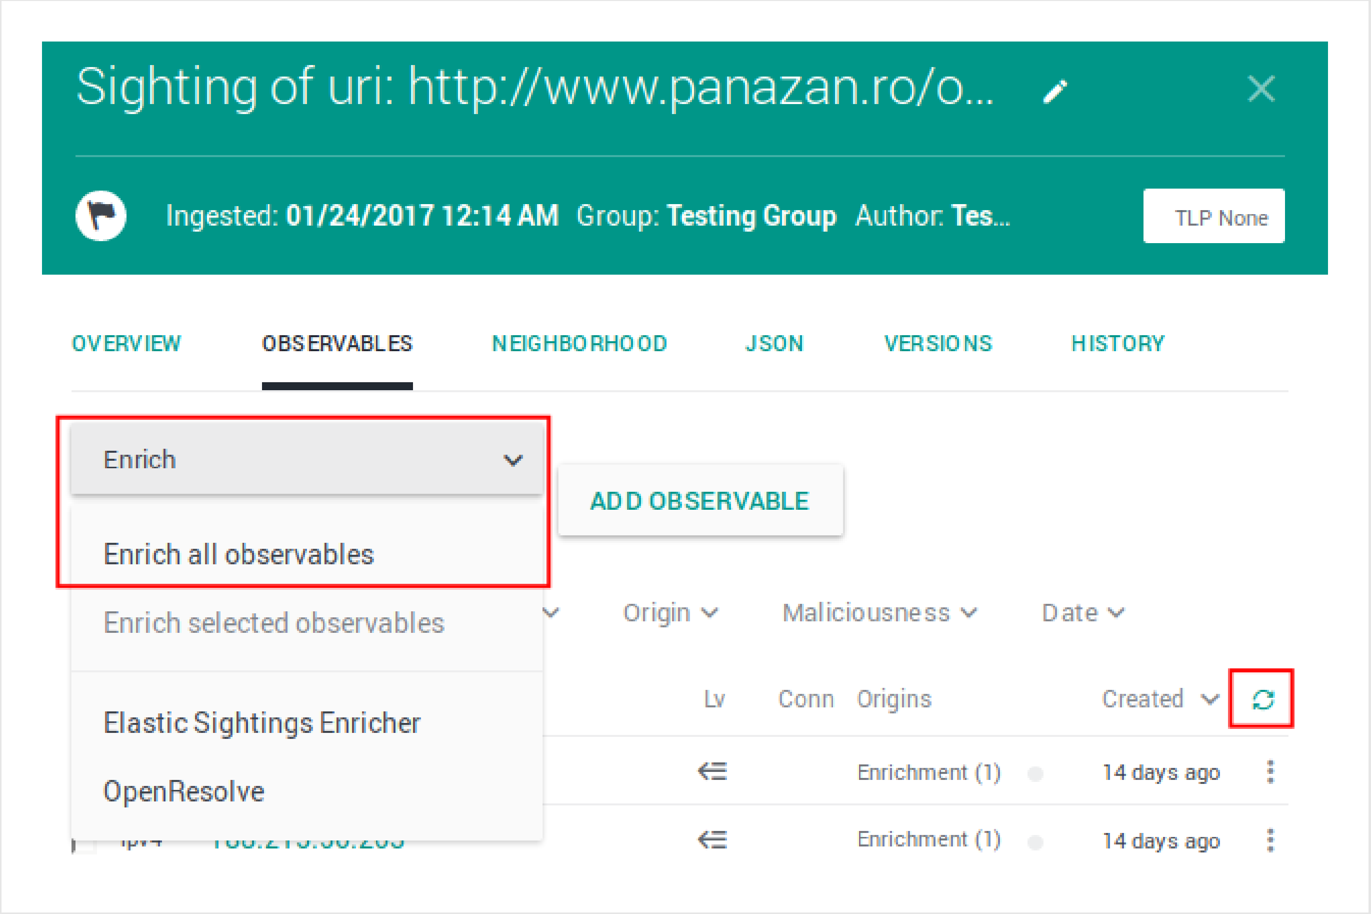

Neighborhood

During an analysis you may want to quickly inspect an entity to check relationships to other entities and extracts. 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 is a faster alternative: click it to see a small graph displaying close-range relationships the entity has with nearby entities and extracts.

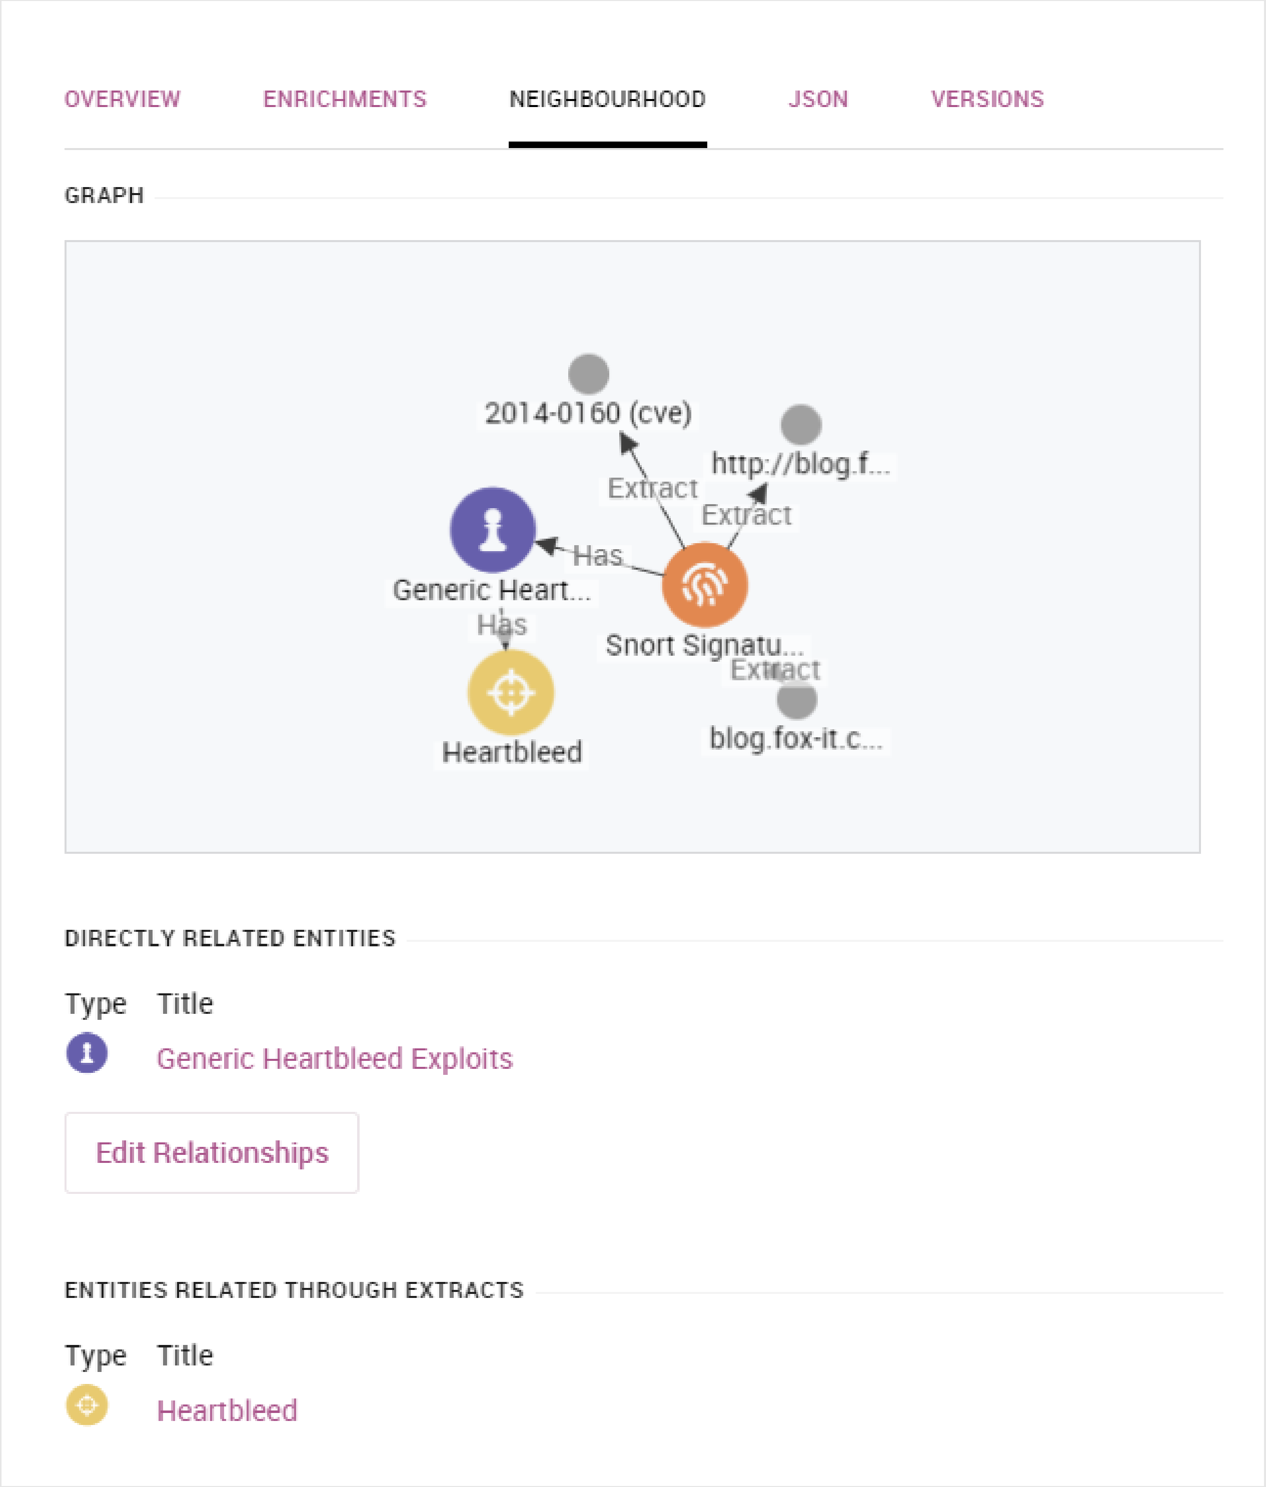

DIRECTLY RELATED ENTITIES

Type	Title
ℹ	Generic Heartbleed Exploits

ENTITIES RELATED THROUGH EXTRACTS

Type	Title
⌖	Heartbleed

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

If more than 100 relationships are found for the entity, only the 30 most recently created relationships are displayed on the **Neighborhood** graph. In this case, a notification message is displayed to inform the user:

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

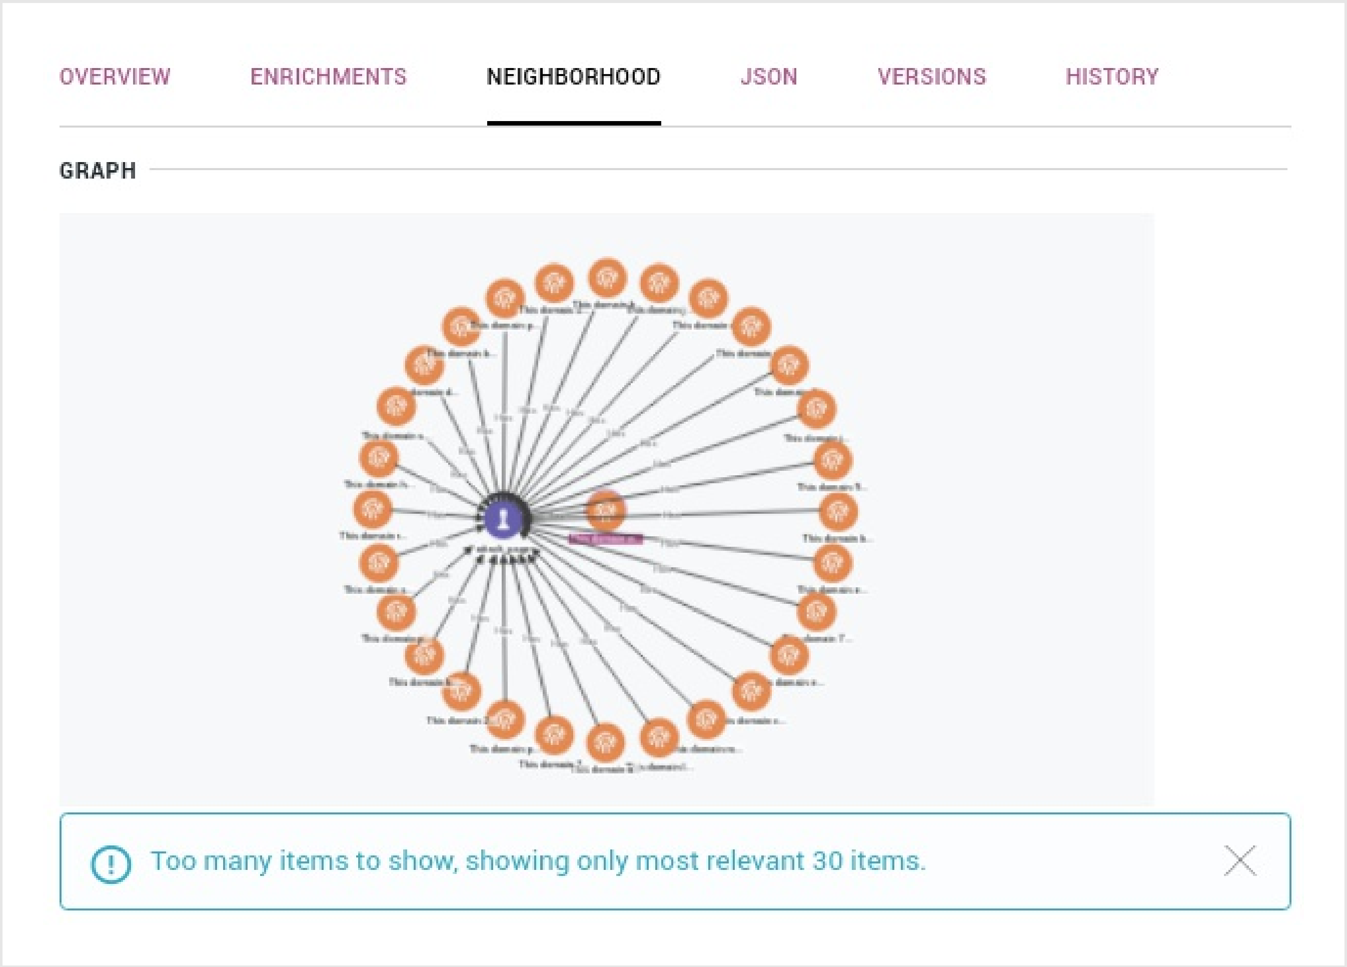

Too many items to show, showing only most relevant 30 items.

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 goes out of sync. 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.

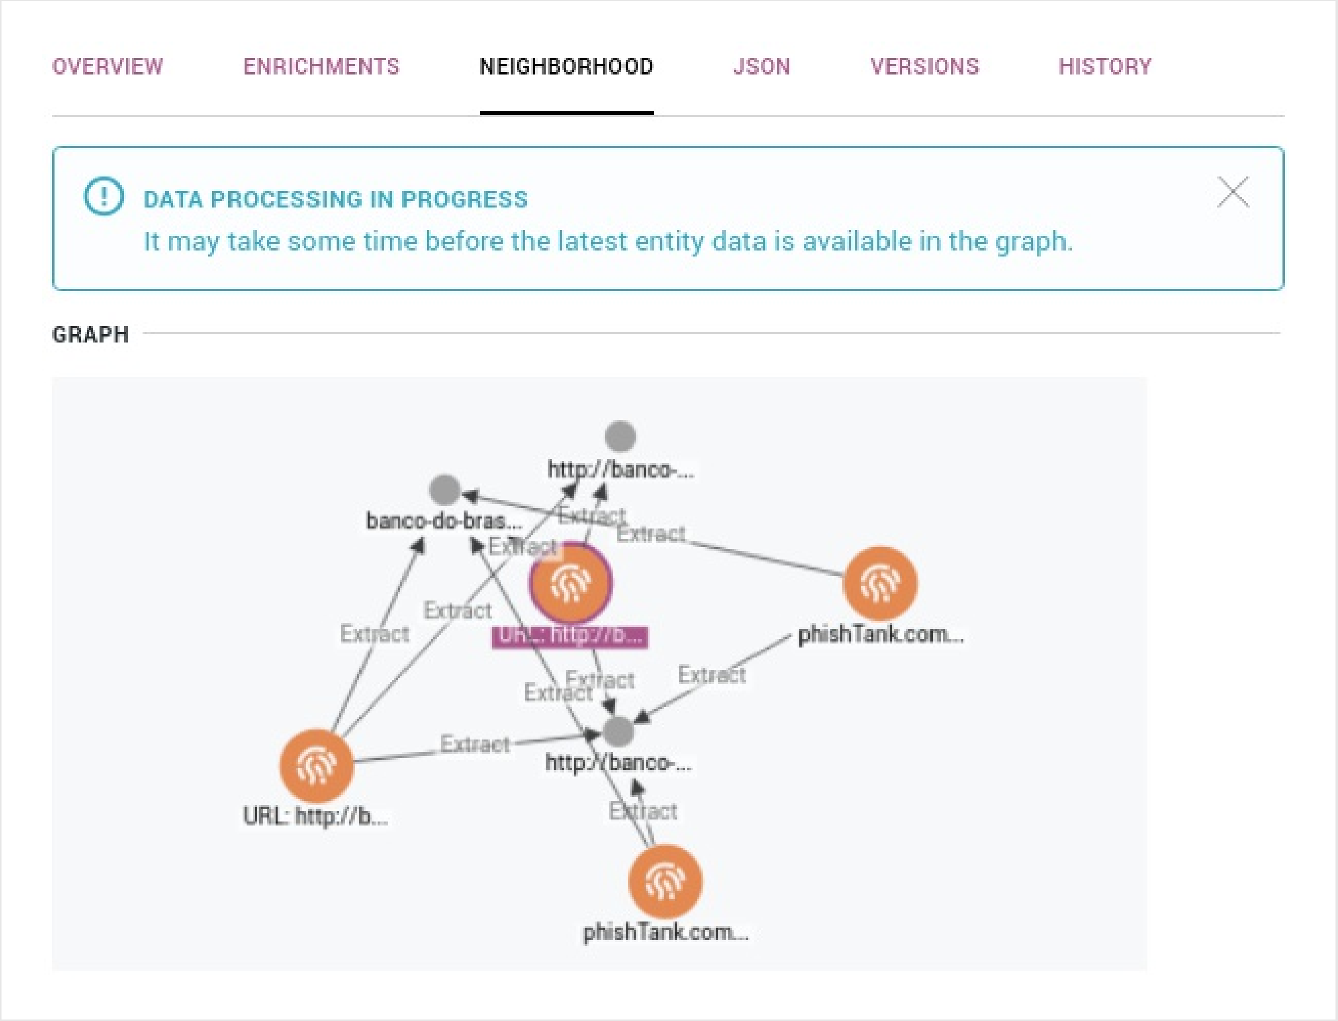

DATA PROCESSING IN PROGRESS

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

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

In the visual graph representation you can inspect any relationships the entity may have with other entities in the platform. Relationships can be *direct* — the entities are immediately related to each other — or *indirect* — the entities are related through a shared entity or a shared extract.

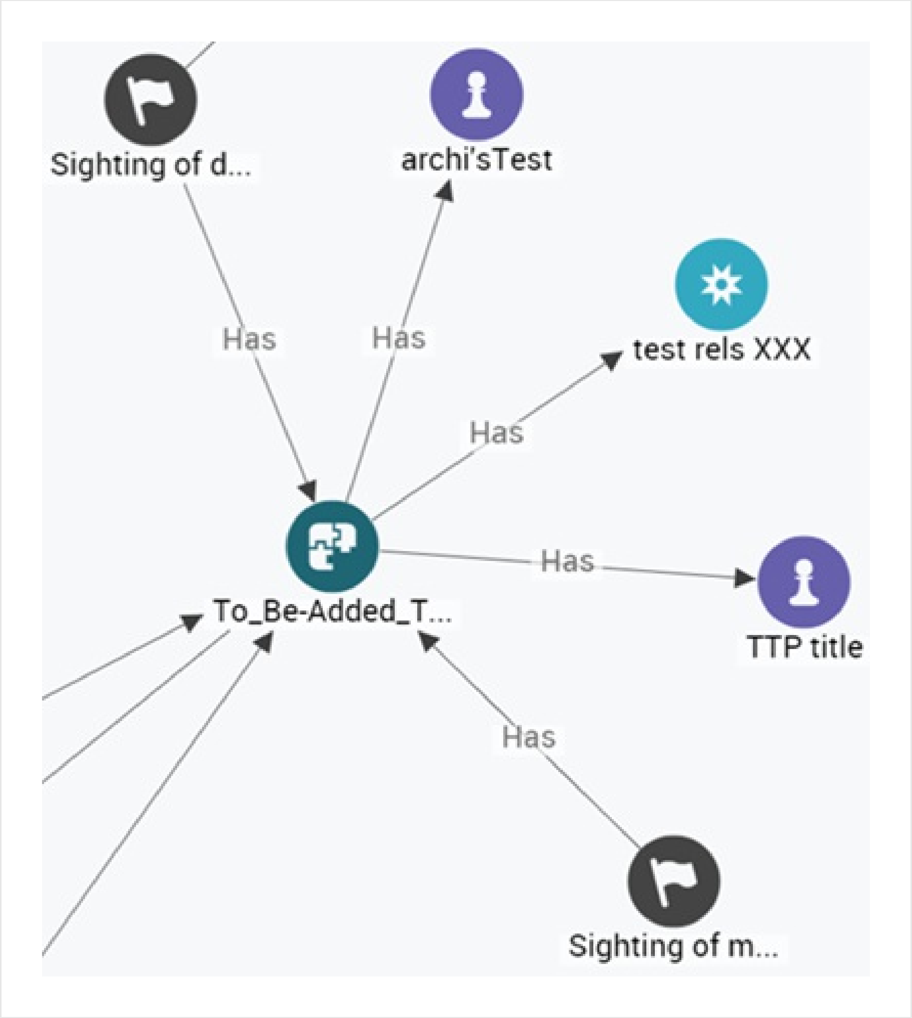

Entities with direct relationships

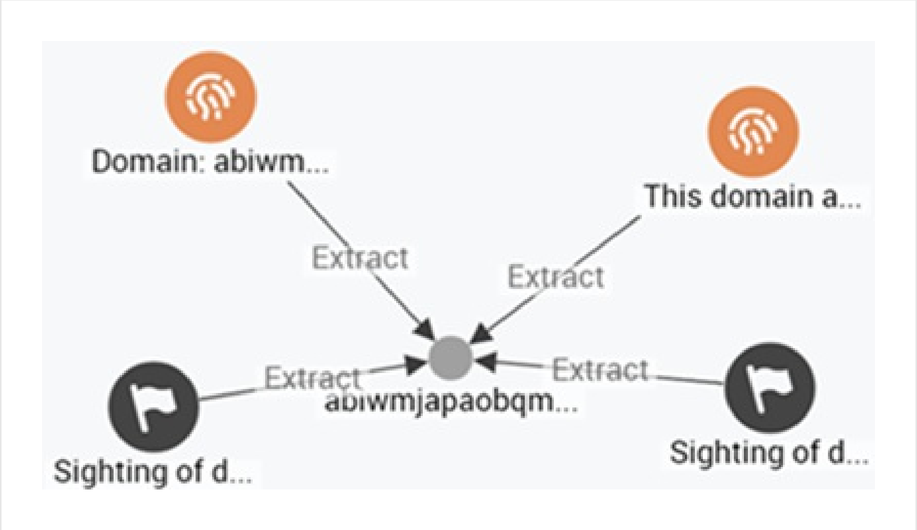

Entities with indirect relationships

To visually examine the entity more closely, click the small graph to launch the larger and more powerful platform graph.

To edit entity relationships, click **Edit relationships**.

Directly related entities

Entities that are directly related to the active entity displayed on the entity pane are listed under **Directly related entities**, and they are sorted by entity type.

Besides the entity name, you can see the TLP color code assigned to the entity, if available, and the entity ingestion time.
To refresh the view, if necessary, click the refresh icon: ↻.

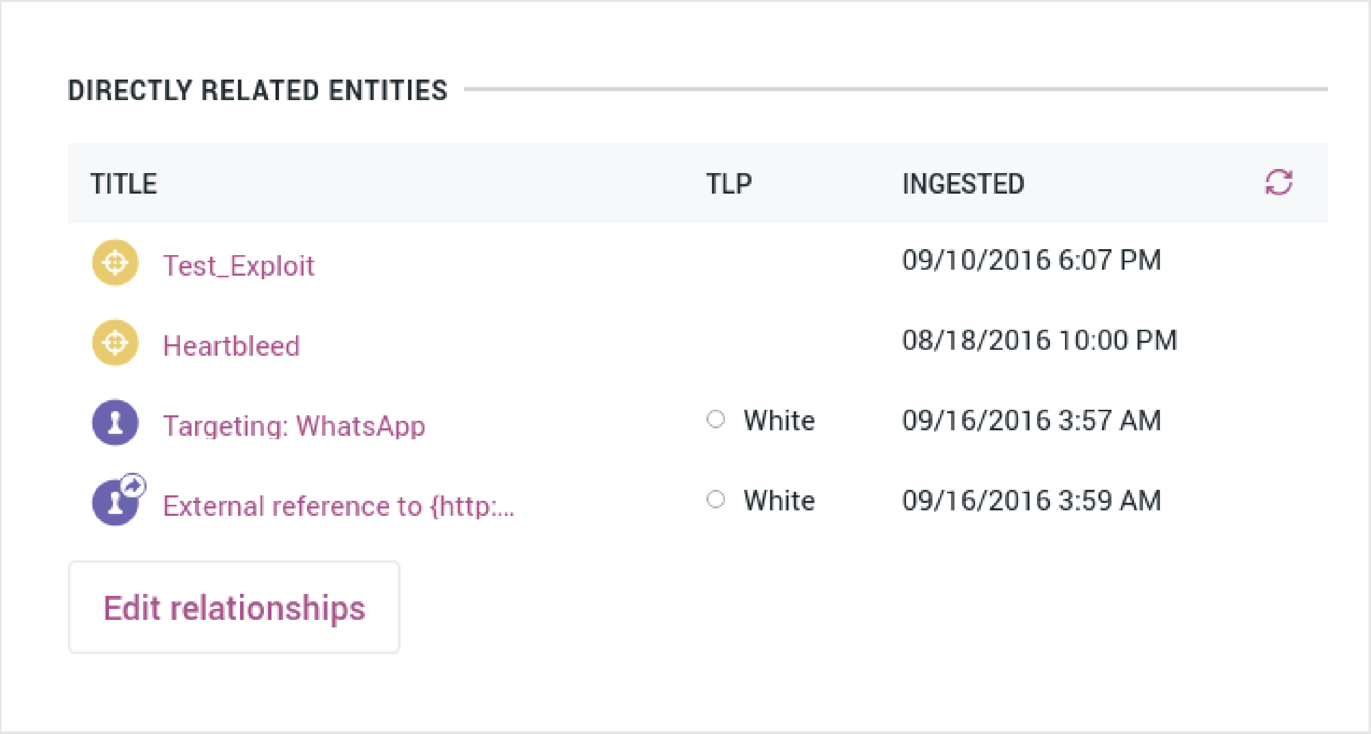

DIRECTLY RELATED ENTITIES

TITLE	TLP	INGESTED
Test_Exploit		09/10/2016 6:07 PM
Heartbleed		08/18/2016 10:00 PM
Targeting: WhatsApp	○ White	09/16/2016 3:57 AM
External reference to {http:...	○ White	09/16/2016 3:59 AM

Entities related through extracts

Entities that are indirectly related through extracts to the active entity displayed on the entity pane are listed under **Entities related through extracts**, and they are sorted by entity type.
Besides the entity name, you can see the TLP color code assigned to the entity, if available, and the entity ingestion time.
To refresh the view, if necessary, click the refresh icon: ↻.

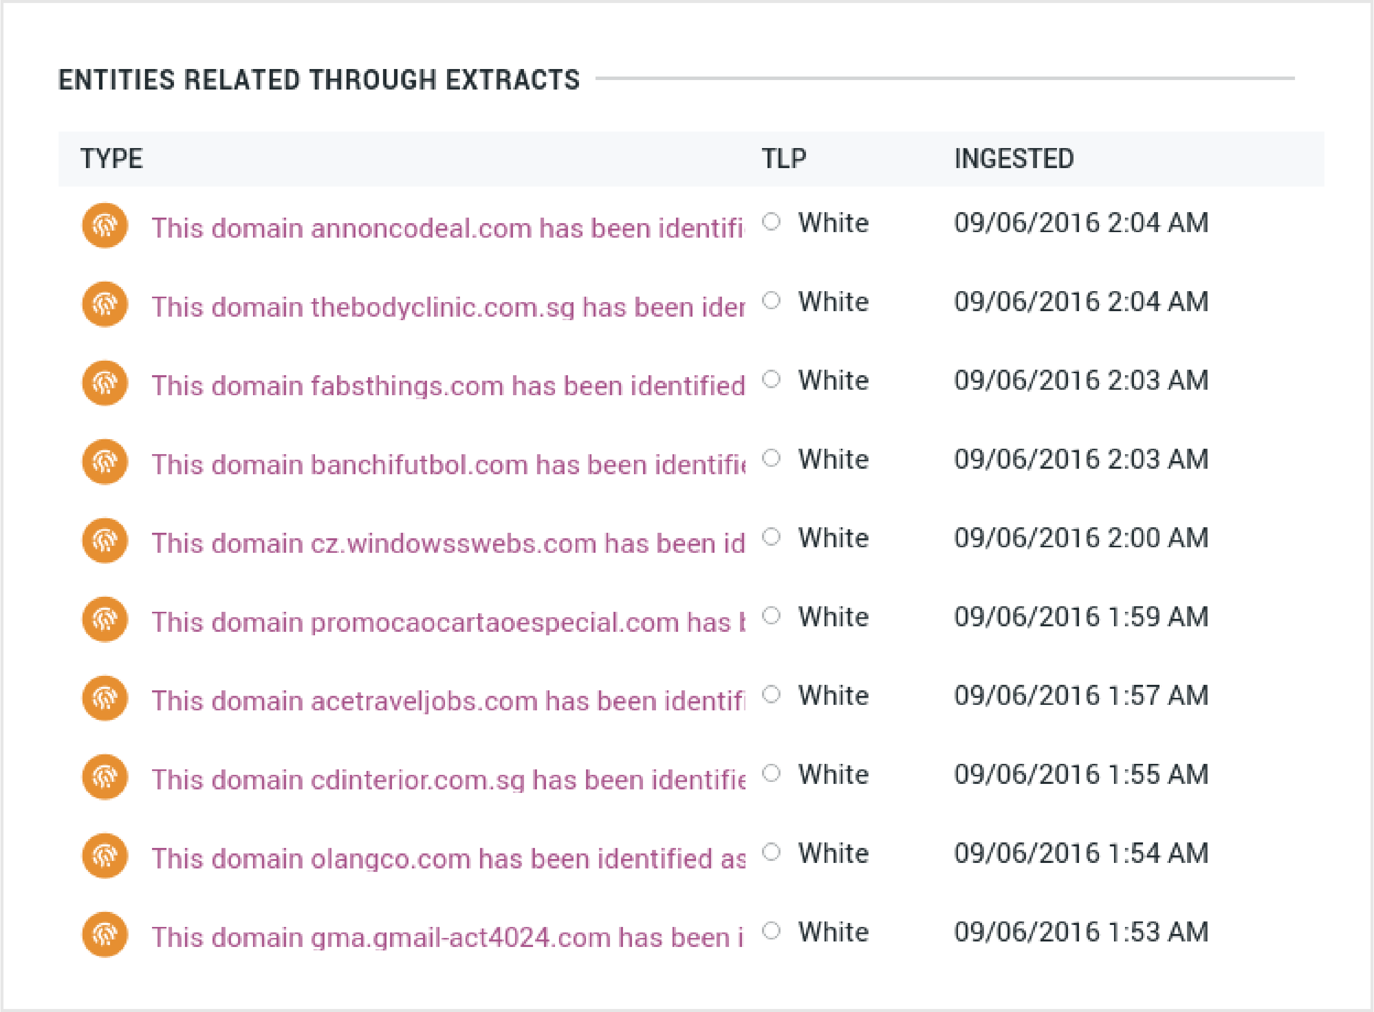

ENTITIES RELATED THROUGH EXTRACTS

TYPE	TLP	INGESTED
This domain annoncodeal.com has been identifi	○ White	09/06/2016 2:04 AM
This domain thebodyclinic.com.sg has been ider	○ White	09/06/2016 2:04 AM
This domain fabsthings.com has been identified	○ White	09/06/2016 2:03 AM
This domain banchifutbol.com has been identifi	○ White	09/06/2016 2:03 AM
This domain cz.windowsswebs.com has been id	○ White	09/06/2016 2:00 AM
This domain promocaocartaoespecial.com has	○ White	09/06/2016 1:59 AM
This domain acetraveljobs.com has been identifi	○ White	09/06/2016 1:57 AM
This domain cdinterior.com.sg has been identifi	○ White	09/06/2016 1:55 AM
This domain olangco.com has been identified as	○ White	09/06/2016 1:54 AM
This domain gma.gmail-act4024.com has been i	○ White	09/06/2016 1:53 AM

Related datasets

Entities can belong to one, more, or no datasets. Any datasets the entity is part of are listed here.

Related workspaces

Entities can belong to one, more, or no workspaces. Any workspaces the entity is part of are listed here.

Related tasks

Actionable user tasks associated with the entity are listed here You can create tasks and assign them to yourself or to other users to request follow-ups, for example further investigation or a call for action.

JSON

This tab displays the entity as a JSON object. You can expand or compress the nodes to show or hide the corresponding data.

Versions

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

History

The **History** tab shows a reverse chronological order of the actions performed on the entity, like creation or modification, along with the user names of the corresponding actors.

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.
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)`.

Entity type	Description
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.

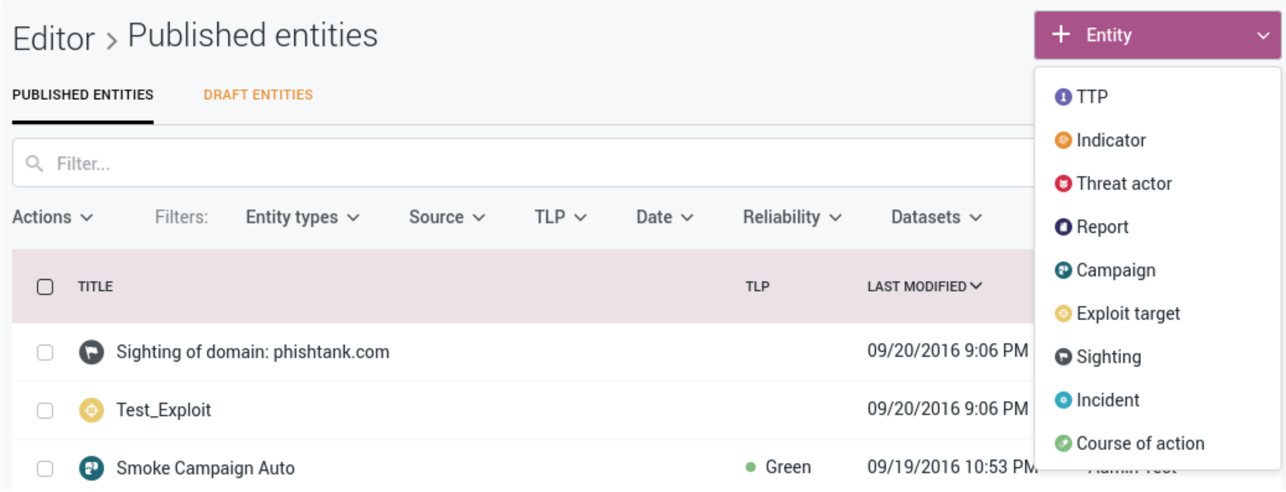

☐	TITLE	TLP	LAST MODIFIED
☐	Sighting of domain: phishtank.com		09/20/2016 9:06 PM
☐	Test_Exploit		09/20/2016 9:06 PM
☐	Smoke Campaign Auto	● Green	09/19/2016 10:53 PM

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 concerning the maliciousness of the (potential) threat the entity represents.
Allowed values:

- **Unknown**
- **None**
- **Low**

- **Medium**
- **High**

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.
Press **ENTER** on your keyboard to confirm the current input, and to display a new input field.
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*
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 reports 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.
To manage tags and their parent-child relationships, go to **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.
Press **ENTER** on your keyboard to confirm the current input, and to display a new input field.
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 existingworkspace, 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**: when you are done, 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**. Draft entities associated with a workspace are available in the corresponding workspace under **Entities > Draft entities**.
  You can also click the drop-down arrow to save the current entry as a draft and create a new draft one right away — **Save as draft and new** — or to save the current entry and duplicate it —**Save as draft and duplicate** — for example to use it as a quick template for a new entry.
- **Publish**: when you are done, 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 entities**. Published entities associated with a workspace are available in the corresponding workspace under **Entities > Entities**.
  You can also click the drop-down arrow to save the current entry, publish it to the platform, and create a new entry one right away — **Publish and new** — or to publish the current entry and duplicate it —**Publish and duplicate** — for example to use it as a quick template for a new entry.

Edit an entity

Delete an entity

Filter entities

You can filter entities to zero in on specific sub-sets:

- 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	
Entity Types	Filter entities by type.	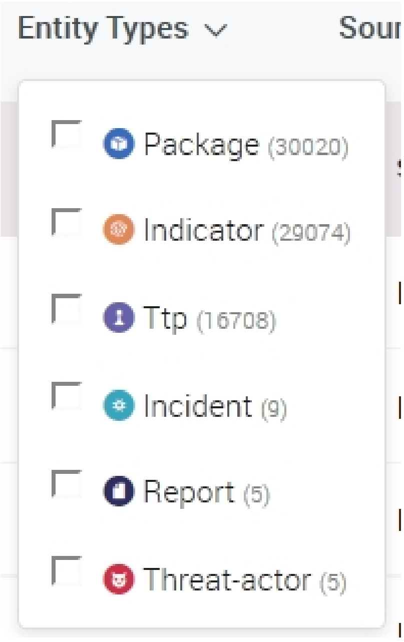 Entity Types Sour ☐ Package (30020) ☐ Indicator (29074) ☐ Ttp (16708) ☐ Incident (9) ☐ Report (5) ☐ Threat-actor (5)
Source Types	Filter entities by source/origin.	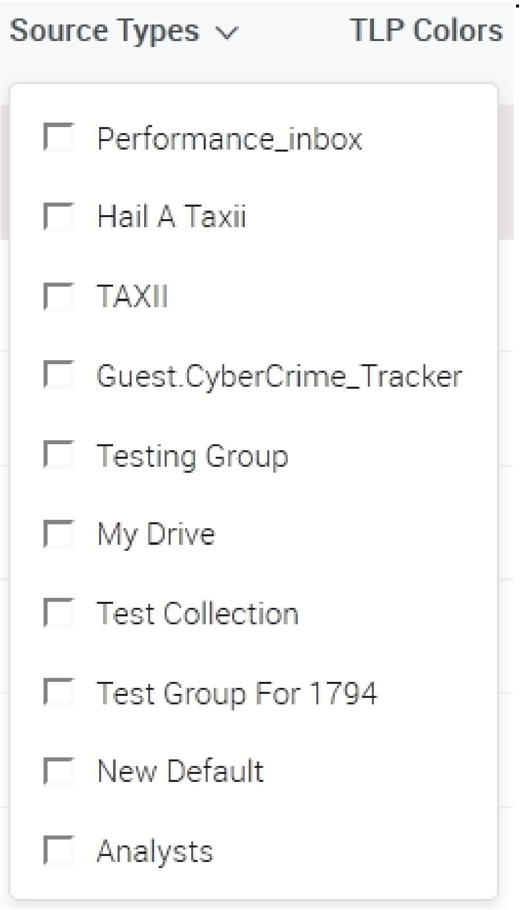 Source Types TLP Colors ☐ Performance_inbox ☐ Hail A Taxii ☐ TAXII ☐ Guest.CyberCrime_Tracker ☐ Testing Group ☐ My Drive ☐ Test Collection ☐ Test Group For 1794 ☐ New Default ☐ Analysts

Filter menu	Description	
TLP Colors	Filter entities by Traffic Light Protocol color code.	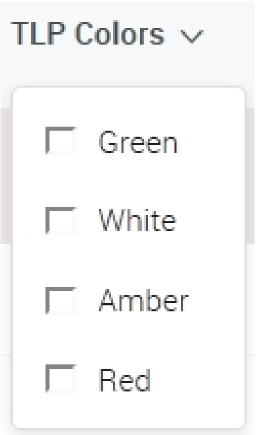
Date	Filter entities included in a date range.	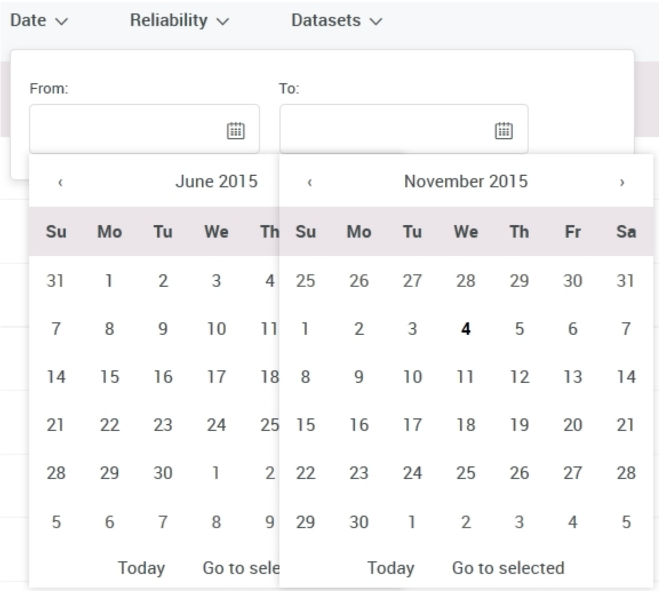 Date Reliability Datasets From: To: June 2015 Su Mo Tu We Th 31 1 2 3 4 7 8 9 10 11 14 15 16 17 18 21 22 23 24 25 28 29 30 1 2 5 6 7 8 9 Today Go to sele November 2015 Su Mo Tu We Th Fr Sa 25 26 27 28 29 30 31 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 1 2 3 4 5 Today Go to selected
Reliability	Filter entities based on their reliability index.	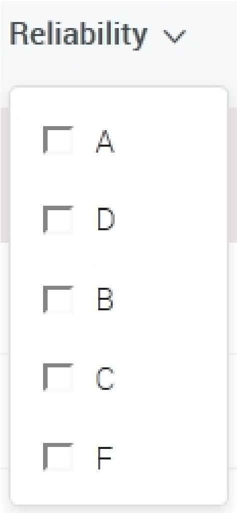

Filter menu	Description	
Datasets	Filter entities based on the dataset they belong to.	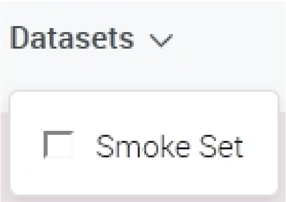

Search entities

Work with entities

You can edit and modify entity data by choosing the desired option from the **Actions** pop-up menu on the bottom half of each entity detail pane tab. For example you can load it onto the graph to further examine it in a visual environment, you can create a task to take action and respond to the entity, or you can download it.

Grayed-out menu options are disabled in the current context, and are not available.

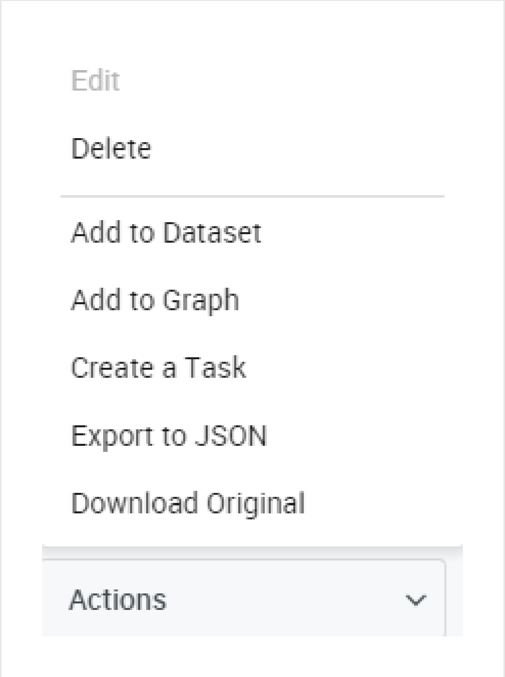

Select this menu option…	…to carry out this action
Edit	Makes the entity detail pane fields writable, so that you can edit and modify the entity.

Select this menu option…	**…to carry out this action**
Delete	Permanently removes the entity from the platform. If you confirm the entity deletion in the pop-up confirmation dialog, the operation is completed, and it cannot be undone.
Add to dataset	Adds the entity to the specified dataset(s).
Add to graph	Adds the selected entity to the graph to further examine it in a visual environment.
Create a Task	Makes the entity data actionable by creating a task, and by assigning it to a user.
Export to JSON	Exports the entity as a JSON object that can be saved, for example locally.
Download Original	Downloads the selected entity in its original data format.