ThreatQuotient

MITRE ATLAS CDF

Version 1.0.0

February 10, 2025

ThreatQuotient
20130 Lakeview Center Plaza Suite 400
Ashburn, VA 20147

ThreatQ Supported

Support
Email: support@threatq.com
Web: support.threatq.com
Phone: 703.574.9893

Contents

Warning and Disclaimer

ThreatQuotient, Inc. provides this document “as is”, without representation or warranty of any kind, express or implied, including without limitation any warranty concerning the accuracy, adequacy, or completeness of such information contained herein. ThreatQuotient, Inc. does not assume responsibility for the use or inability to use the software product as a result of providing this information.

Support

This integration is designated as **ThreatQ Supported**.

Support Email: support@threatq.com
Support Web: https://support.threatq.com
Support Phone: 703.574.9893

Integrations/apps/add-ons designated as **ThreatQ Supported** are fully supported by ThreatQuotient's Customer Support team.

ThreatQuotient strives to ensure all ThreatQ Supported integrations will work with the current version of ThreatQuotient software at the time of initial publishing. This applies for both Hosted instance and Non-Hosted instance customers.

> ThreatQuotient does not provide support or maintenance for integrations, apps, or add-ons published by any party other than ThreatQuotient, including third-party developers.

Integration Details

ThreatQuotient provides the following details for this integration:

Current Integration Version	1.0.0
Compatible with ThreatQ Versions	>= 5.12.1
Support Tier	ThreatQ Supported

Introduction

The MITRE Adversarial Threat Landscape for AI Systems (ATLAS) CDF integration allows teams to ingest the MITRE ATLAS knowledgebase of artificial intelligence techniques used by threat actors. These techniques are mapped alongside the MITRE ATT&CK framework, providing a view of the tactics, techniques, and procedures (TTPs) used by adversaries to target AI systems.

MITRE ATLAS is a comprehensive knowledge base designed to help protect artificial intelligence systems by cataloging real-world attacks on AI, providing insights into adversary tactics and techniques.

The integration provides the following feeds:

- **MITRE ATLAS** - ingests techniques used by threat actors targeting AI systems.
- **MITRE ATLAS Reports** - ingests the case studies/reports generated by OpenCTI that are linked to the MITRE ATLAS framework.

The integration ingests the following object types:

- Attack Patterns
- Courses of Action
- Incidents
- Identities
- Reports
- Vulnerabilities

Prerequisites

The following is not required but recommended for the integration:

- **GitHub API Token -** In order to pull data from the GitHub API with the MITRE ATLAS Reports feed, you may need to authenticate with a GitHub API Token. See the Known Issues / Limitations section for more details.

Installation

Perform the following steps to install the integration:

> The same steps can be used to upgrade the integration to a new version.

1. Log into https://marketplace.threatq.com/.
2. Locate and download the integration yaml file.
3. Navigate to the integrations management page on your ThreatQ instance.
4. Click on the **Add New Integration** button.
5. Upload the integration yaml file using one of the following methods:
   - Drag and drop the file into the dialog box
   - Select **Click to Browse** to locate the file on your local machine
6. Select the individual feeds to install, when prompted, and click **Install**.

   > ThreatQ will inform you if the feed already exists on the platform and will require user confirmation before proceeding. ThreatQ will also inform you if the new version of the feed contains changes to the user configuration. The new user configurations will overwrite the existing ones for the feed and will require user confirmation before proceeding.

The feed(s) will be added to the integrations page. You will still need to configure and then enable the feed.

Configuration

> ThreatQuotient does not issue API keys for third-party vendors. Contact the specific vendor to obtain API keys and other integration-related credentials.

To configure the integration:

1. Navigate to your integrations management page in ThreatQ.
2. Select the **OSINT** option from the *Category* dropdown (optional).

   > If you are installing the integration for the first time, it will be located under the **Disabled** tab.

3. Click on the integration entry to open its details page.
4. Enter the following parameters under the **Configuration** tab:

MITRE ATLAS Parameters

PARAMETER	DESCRIPTION
Save Intrusion Set As	Select how to ingest Intrusion Sets into the ThreatQ platform. Options include: ◦ Adversaries ◦ Intrusion Sets ThreatQuotient recommends users select the Adversaries option to create a consolidated profile.
Disable Proxies	Enable this option if the feed should not honor proxies set in the ThreatQ UI.
Enable SSL Certificate Verification	Enable this for the feed to validate the host-provided SSL certificate.

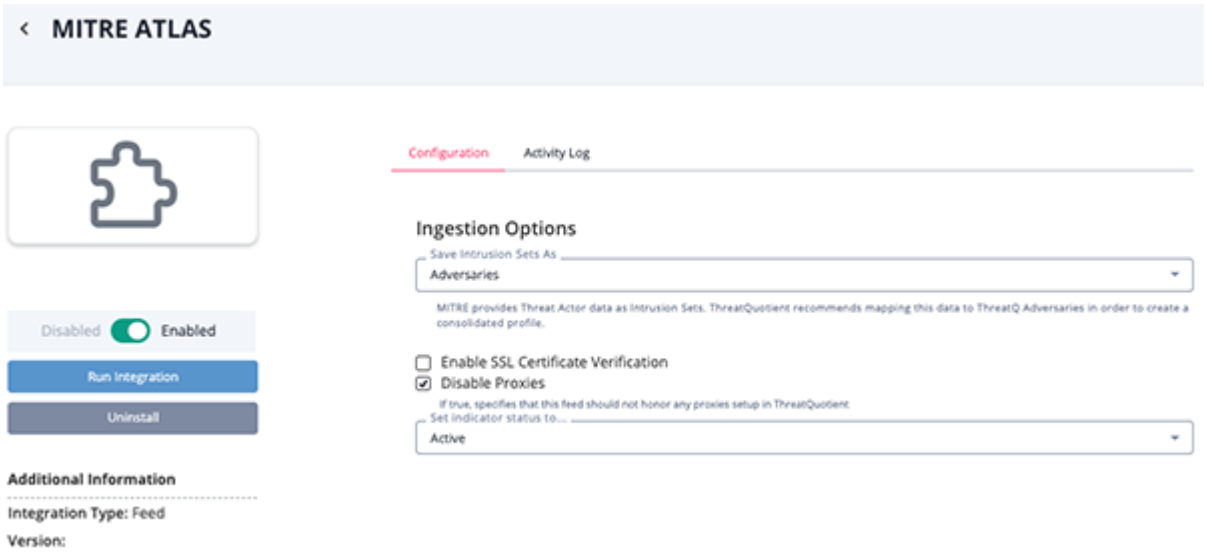
MITRE ATLAS
Disabled
Enabled
Run Integration
Uninstall
Additional Information
Integration Type: Feed
Version:
Configuration
Activity Log
Ingestion Options
Save Intrusion Sets As
Adversaries
MITRE provides Threat Actor data as Intrusion Sets. ThreatQuotient recommends mapping this data to ThreatQ Adversaries in order to create a consolidated profile.
Enable SSL Certificate Verification
Disable Proxies
If true, specifies that this feed should not honor any proxies setup in ThreatQuotient
Set Indicator status to...
Active

MITRE ATLAS Reports Parameters

PARAMETER	DESCRIPTION
GitHub API Token	Optional and Recommended - Enter your GitHub API Token (Personal Access Token). The MITRE ATLAS Reports feed will fetch the MITRE ATLAS Report data from their public GitHub repository. These reports are generated by OpenCTI and contain information linked to the MITRE ATLAS framework. The ATLAS Reports reside within a GitHub repository and are updated on a semi-regular basis. In order to pull data from the GitHub API, you may need to authenticate with a GitHub API Token. This will increase the rate limit from 60 requests per hour to 5000 requests per hour. If you are experiencing rate limiting issues, please consider adding your GitHub API Token to the feed configuration. Authentication is not required to access the MITRE ATLAS data, but is highly recommended. If you are experiencing rate limiting issues, you can add a GitHub API Token to the feed configuration. Unauthenticated users have a rate limit of 60 requests per hour, while authenticated users have a rate limit of 5000 requests per hour.
Save Intrusion Set As	Select how to ingest Intrusion Sets into the ThreatQ platform. Options include: ◦ Adversaries ◦ Intrusion Sets ThreatQuotient recommends users select the Adversaries option to create a consolidated profile.
Disable Proxies	Enable this option if the feed should not honor proxies set in the ThreatQ UI.

PARAMETER	DESCRIPTION
Enable SSL Certificate Verification	Enable this for the feed to validate the host-provided SSL certificate.

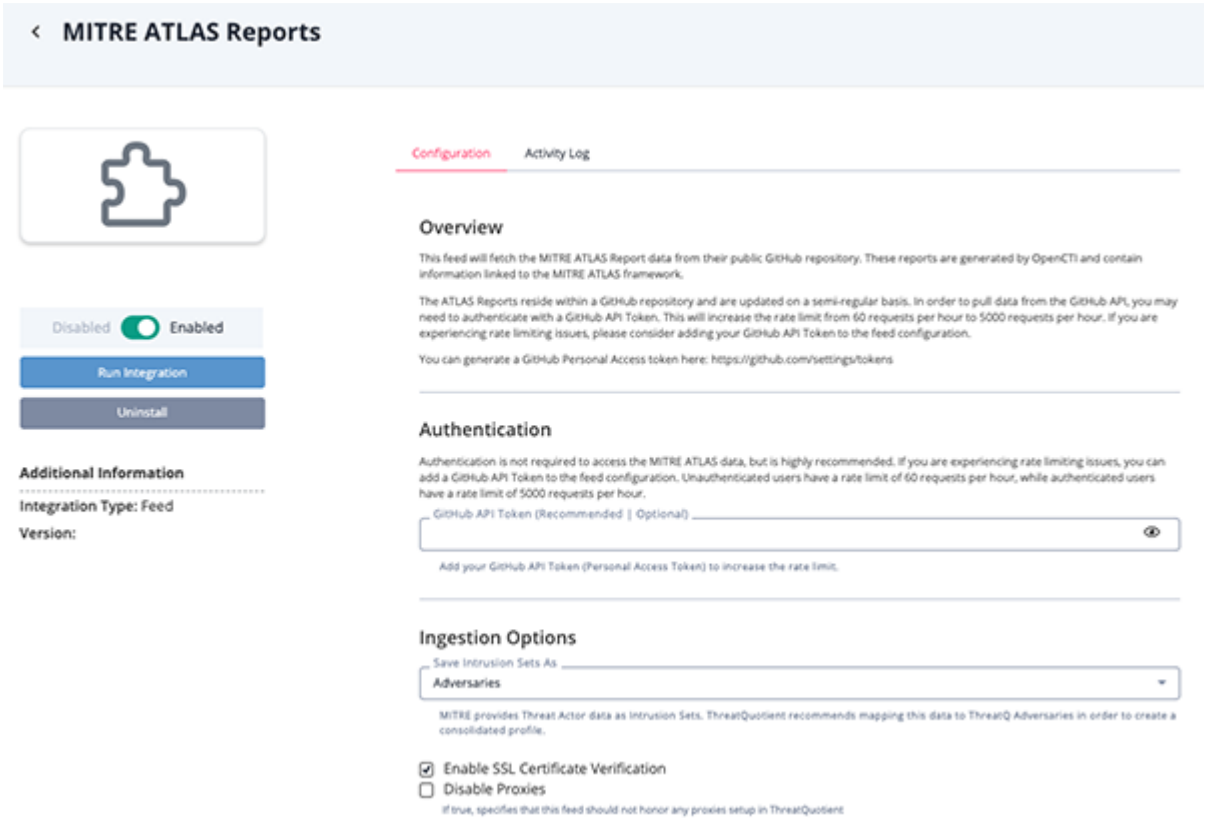

5. Review any additional settings, make any changes if needed, and click on **Save**.
6. Click on the toggle switch, located above the *Additional Information* section, to enable it.

ThreatQ Mapping

MITRE ATLAS

The MTIRE ATLAS feed ingests techniques used by threat actors targeting AI systems. These techniques will be ingested into ThreatQ as Attack Pattern objects, and will include context such as tactics, references, and descriptions.

```
GET https://raw.githubusercontent.com/mitre-atlas/atlas-navigator-data/main/dist/stix-atlas.json
```

Sample Response:


```
{
    "type": "bundle",
    "id": "bundle--a9061128-f24c-414d-ac6b-4dc6a3329e3b",
    "objects": [
        {
            "type": "x-mitre-tactic",
            "spec_version": "2.1",
            "id": "x-mitre-tactic--4bf0ef0b-3de4-46fe-9a0d-3b1dea352a30",
            "created": "2024-06-24T20:23:47.456097Z",
            "modified": "2024-06-24T20:23:47.456097Z",
            "name": "Reconnaissance",
            "description": "The adversary is trying to gather information about the machine learning system they can use to plan future operations.\n\nReconnaissance consists of techniques that involve adversaries actively or passively gathering information that can be used to support targeting.\nSuch information may include details of the victim organizations' machine learning capabilities and research efforts.\nThis information can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using gathered information to obtain relevant ML artifacts, targeting ML capabilities used by the victim, tailoring attacks to the particular models used by the victim, or to drive and lead further Reconnaissance efforts.\n",
            "external_references": [
                {
                    "source_name": "mitre-atlas",
                    "url": "https://atlas.mitre.org/tactics/AML.TA0002",
                    "external_id": "AML.TA0002"
                }
            ],
            "x_mitre_shortname": "reconnaissance"
        },
        {
            "type": "x-mitre-tactic",
            "spec_version": "2.1",
            "id": "x-mitre-tactic--6b232c1e-ada7-4cd4-b538-7a1ef6193e2f",
            "created": "2024-06-24T20:23:47.45631Z",
```

```
            "modified": "2024-06-24T20:23:47.45631Z",
            "name": "Resource Development",
            "description": "The adversary is trying to establish resources they
can use to support operations.\n\nResource Development consists of techniques
that involve adversaries creating,\npurchasing, or compromising/stealing
resources that can be used to support targeting.\nSuch resources include
machine learning artifacts, infrastructure, accounts, or capabilities.\nThese
resources can be leveraged by the adversary to aid in other phases of the
adversary lifecycle, such as [ML Attack Staging](/tactics/AML.TA0001).\n",
            "external_references": [
                {
                    "source_name": "mitre-atlas",
                    "url": "https://atlas.mitre.org/tactics/AML.TA0003",
                    "external_id": "AML.TA0003"
                }
            ],
            "x_mitre_shortname": "resource-development"
        }
    ]
}
```


The data is mapped via the default STIX 2.1 mapping, which can be found on the ThreatQ Help Center: https://helpcenter.threatq.com/ThreatQ_Platform/System_Objects/STIX/STIX_2.1.htm.

MITRE ATLAS Reports

The MITRE ATLAS Reports feed ingests the case studies/reports generated by OpenCTI that are linked to the MITRE ATLAS framework. These reports will be ingested into ThreatQ as Report objects, and will include context such as the report title, description, and references.

The feed makes requests to the following GitHub API Endpoints:

`GET https://api.github.com/rate_limit` - fetches the rate limits for the current user (or no user).
`GET https://api.github.com/repos/mitre-atlas/atlas-navigator-data/commits` - fetches the commits for the repository's "dist/opencti-bundles" content directory.
`GET https://api.github.com/repos/mitre-atlas/atlas-navigator-data/commits/{{ sha_hash }}` - fetches the changed files for each commit.
`GET https://raw.githubusercontent.com/mitre-atlas/atlas-navigator-data/main/dist/opencti-bundles/{{ file_name }}` - fetches the raw JSON file for each case study.

Sample Response:

```
{
    "type": "bundle",
    "id": "bundle--ce0e48fe-3174-4e6d-9887-d87f2b286f63",
    "objects": [
        {
            "id": "report--2a5fab52-5c12-56b2-8ca2-5e1e3c8ae6bb",
            "spec_version": "2.1",
            "revoked": false,
            "x_opencti_reliability": "A - Completely reliable",
            "confidence": 75,
            "created": "2020-01-12T12:00:00.000Z",
            "modified": "2024-05-22T17:24:46.467Z",
            "name": "MITRE ATLAS Case Study: Microsoft Azure Service
Disruption",
            "description": "The Microsoft AI Red Team performed a red team
exercise on an internal Azure service with the intention of disrupting its
service. This operation had a combination of traditional ATT&CK enterprise
techniques such as finding valid account, and exfiltrating data -- all
interleaved with adversarial ML specific steps such as offline and online
evasion examples.",
            "report_types": [
                "internal-report"
            ],
            "published": "2020-01-12T12:00:00.000Z",
            "x_opencti_workflow_id": "4c154301-4863-42d1-ba1a-bcf5cff32385",
            "labels": [
                "mitre atlas source",
                "aml.cs0010"
            ],
            "external_references": [
                {
```

```
                    "source_name": "AML.CS0010",
                    "url": "https://atlas.mitre.org/studies/AML.CS0010/"
                }
            ],
            "x_opencti_id": "4f8938d4-2da4-445b-8770-5f7f18cd9248",
            "x_opencti_type": "Report",
            "type": "report",
            "created_by_ref": "identity--48b4f20a-ddd3-5524-a646-55dd68b62ede",
            "object_marking_refs": [
                "marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9"
            ],
            "object_refs": [
                "attack-pattern--18c33065-849f-5705-b4d4-28a08470fba6",
                "attack-pattern--2dc305cb-dd66-55b6-a0c8-87dbf69bae66"
                "incident--8458a7d3-dc00-5460-9d32-8db303ab7415"
            ]
        },
        {
            "id": "attack-pattern--18c33065-849f-5705-b4d4-28a08470fba6",
            "spec_version": "2.1",
            "revoked": false,
            "confidence": 0,
            "created": "2023-10-31T13:48:07.631Z",
            "modified": "2023-11-20T18:19:47.016Z",
            "name": "Exfiltration via Cyber Means",
            "description": "Adversaries may exfiltrate ML artifacts or other
information relevant to their goals via traditional cyber means.\n\nSee the
ATT&CK [Exfiltration](https://attack.mitre.org/tactics/TA0010/) tactic for more
information.",
            "x_mitre_id": "AML.T0025",
            "labels": [
                "atlas"
            ],
            "kill_chain_phases": [
                {
                    "kill_chain_name": "mitre-atlas",
                    "phase_name": "exfiltration",
                    "x_opencti_order": 13
                }
            ],
            "external_references": [
                {
                    "source_name": "mitre-atlas",
                    "url": "https://atlas.mitre.org/techniques/AML.T0025",
                    "external_id": "AML.T0025"
                }
            ],
            "x_opencti_id": "ec43ad24-eebc-4442-b1cd-7d82123b9313",
            "x_opencti_type": "Attack-Pattern",
            "type": "attack-pattern",
```

```
            "object_marking_refs": [
                "marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9",
                "marking-definition--c4099854-e00c-5dbc-b0fe-4c9909920101"
            ]
        },
        {
            "id": "attack-pattern--2dc305cb-dd66-55b6-a0c8-87dbf69bae66",
            "spec_version": "2.1",
            "revoked": false,
            "confidence": 0,
            "created": "2023-10-31T13:48:07.628Z",
            "modified": "2023-11-20T18:19:43.967Z",
            "name": "Valid Accounts",
            "description": "Adversaries may obtain and abuse credentials of
existing accounts as a means of gaining Initial Access.\nCredentials may take
the form of usernames and passwords of individual user accounts or API keys
that provide access to various ML resources and services.\n\nCompromised
credentials may provide access to additional ML artifacts and allow the
adversary to perform [Discover ML Artifacts](/techniques/AML.T0007).
\nCompromised credentials may also grant and adversary increased privileges
such as write access to ML artifacts used during development or production.",
            "x_mitre_id": "AML.T0012",
            "labels": [
                "atlas"
            ],
            "kill_chain_phases": [
                {
                    "kill_chain_name": "mitre-atlas",
                    "phase_name": "initial-access",
                    "x_opencti_order": 3
                }
            ],
            "external_references": [
                {
                    "source_name": "mitre-atlas",
                    "url": "https://atlas.mitre.org/techniques/AML.T0012",
                    "external_id": "AML.T0012"
                }
            ],
            "x_opencti_id": "eb0a6b17-ae24-4d15-be8f-c70ade49345d",
            "x_opencti_type": "Attack-Pattern",
            "type": "attack-pattern",
            "object_marking_refs": [
                "marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9",
                "marking-definition--c4099854-e00c-5dbc-b0fe-4c9909920101"
            ]
        },
        {
            "id": "incident--8458a7d3-dc00-5460-9d32-8db303ab7415",
            "spec_version": "2.1",
```

```
            "revoked": false,
            "confidence": 75,
            "created": "2024-02-06T15:29:37.168Z",
            "modified": "2024-02-06T16:44:38.985Z",
            "name": "2020 Microsoft Azure Service Disruption",
            "description": "The Microsoft AI Red Team performed a red team
exercise on an internal Azure service with the intention of disrupting its
service. This operation had a combination of traditional ATT&CK enterprise
techniques such as finding valid account, and exfiltrating data -- all
interleaved with adversarial ML specific steps such as offline and online
evasion examples.",
            "first_seen": "2020-01-12T12:00:00.000Z",
            "incident_type": "research-finding",
            "labels": [
                "aml.cs0010"
            ],
            "external_references": [
                {
                    "source_name": "AML.CS0010",
                    "url": "https://atlas.mitre.org/studies/AML.CS0010/"
                }
            ],
            "x_opencti_id": "7f59d9c4-f27e-4553-91be-788af57929c7",
            "x_opencti_type": "Incident",
            "type": "incident",
            "created_by_ref": "identity--48b4f20a-ddd3-5524-a646-55dd68b62ede",
            "object_marking_refs": [
                "marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9"
            ]
        }
    ]
}
```


The data is mapped via the default STIX 2.1 mapping, which can be found on the ThreatQ Help Center: https://helpcenter.threatq.com/ThreatQ_Platform/System_Objects/STIX/STIX_2.1.htm.

Average Feed Run

Object counts and Feed runtime are supplied as generalities only - objects returned by a provider can differ based on credential configurations and Feed runtime may vary based on system resources and load.

MITRE ATLAS

METRIC	RESULT
Run Time	1 minute
Attack Patterns	82
Attack Pattern Attributes	428
Courses of Action	20
Course of Action Attributes	60

MITRE ATLAS Reports

METRIC	RESULT
Run Time	1 minute
Attack Patterns	32
Attack Pattern Attributes	234
Identities	2
Incidents	7
Incident Attributes	30
Reports	8
Report Attributes	43
Vulnerabilities	2
Vulnerability Attributes	8

Known Issues / Limitations

- **MITRE ATLAS Reports** - The ATLAS Reports reside within a GitHub repository and are updated on a semi-regular basis. In order to pull data from the GitHub API, you may need to authenticate with a GitHub API Token. This will increase the rate limit from 60 requests per hour to 5000 requests per hour. If you are experiencing rate limiting issues, please consider adding your GitHub API Token to the feed configuration.
  While authentication is not required, it is highly recommended when accessing the MITRE ATLAS data. If you are experiencing rate limiting issues, you can add a **GitHub API Token** to the feed configuration. Unauthenticated users have a rate limit of 60 requests per hour, while authenticated users have a rate limit of 5000 requests per hour.

Change Log

- **Version 1.0.0**
  - Initial release