RÉPUBLIQUE
FRANÇAISE

Liberté
Égalité
Fraternité

VIGINUM

Leveraging *OpenCTI* to combat Foreign Information Manipulation and Interference: A Practical Guide

Part 1: Create and enrich knowledge

A VIGINUM's guidebook
Version 2.0 | April 2025

Contents

1. FIMI THREAT KNOWLEDGE CAPITALIZATION

1.1 UNIFYING THE BEST PRACTICES, A PREREQUISITE FOR INTEROPERABILITY

Over the last years, the global landscape in the fight against information manipulation has witnessed significant changes, characterized by a growing number of initiatives from the civil society, the private sector, and governments. As these efforts expand, the key challenge is now to adopt a common framework that ensures both a unified description of the information threat and effective information sharing.

As a technical and operational service, VIGINUM has been dedicated for several years to promoting and developing common standards to reach this objective. In 2024, VIGINUM released a French translation of the DISARM framework[1], which facilitates the structured description of information manipulation campaigns, and is also contributing to the Defending Against Deception Common Data Model (DAD-CDM) project[2]. This initiative, led by the U.S.-based non-profit organization OASIS, seeks to adapt the STIX language to the fight against information manipulation.

Following the release by VIGINUM of the first guidebook on the use of *OpenCTI* in January 2024, this new guidebook aims at proposing a framework on the information threat knowledge capitalization ("TKC") within the *OpenCTI*[3] platform. This framework is oriented toward unifying the TKC best practices in the tool, in order to ensure that valuable knowledge is retained in the future, remains accessible, and can be easily shared, in order to foster interoperability between different stakeholders.

N.B.: VIGINUM distinguishes the process of imputation, which consists in technically linking observables to an information manipulation set (IMS), and the process of attribution, which is a matter of political choice. In this document, the term "attribution" is used generically to refer to the STIX relationship "attributed to". Moreover, the following examples are fictional or have been adjusted to facilitate the description of the different steps. They should therefore not be considered as VIGINUM's official knowledge.

1.2 THREAT KNOWLEDGE CAPITALIZATION: KNOWLEDGE AS A LONG-TERM RESOURCE

As part of the defensive fight against Foreign Information Manipulation and Interference (FIMI), the Threat Knowledge Capitalization (TKC) refers to the process of identifying, modeling, storing and enriching digital threats knowledge in a structured, sustainable and shareable format.

TKC can take various forms, such as textual notes, visual graphs, or threat-specialized platforms, such as *OpenCTI*.

TKC with *OpenCTI* offers several advantages, including:

- minimizing strategic knowledge loss over time;
- standardizing informational threat modeling across analysts;

[1] https://fr.linkedin.com/posts/viginum_disarm-traduction-fran%C3%A7aise-activity-7159231318042550273-J5eO

[2] https://www.oasis-open.org/2023/11/16/oasis-defending-against-disinformation-dad-cdm/

[3] *OpenCTI* is an open source platform initially co-developed by ANSSI, CERT-EU and the *Luatix* non-profit organization. Today, the French company *Filigran* contributes to its development.

- simplifying searching and analysis in large volumes of contextual and technical threat data;
- enhancing the efficiency of data sharing among stakeholders involved in the defensive ecosystem against FIMI (governments, researchers, media, NGOs, industries, etc.).

The knowledge base typically includes both internal resources, such as internal analysis reports, and external resources, such as technical reports from companies, press articles, official statements, as well as other documents and materials related to FIMI.

2. GENERAL PRINCIPLES

2.1 STIX DATA FORMAT

STIX (*Structured Threat Information eXpression*) is a standardized format developed by OASIS. It is designed to represent information through **objects** (entities or observables) that are connected by **links** (relationships). Originally developed for Cyber Threat Intelligence (CTI), this model is now being progressively adapted to address FIMI.

On *OpenCTI*, the STIX format is used to build knowledge graphs. The objects represent entities such as Intrusion sets, countries, organizations, or individuals. The links, on the other hand, provide meaning between objects: "An Intrusion set **is attributed to** a threat actor", "an organization **is located** in a country", "a campaign **uses** a narrative," etc. To make these relationships possible, every link is:

- **designated**: with a limited set of choices: "uses", "attributed to", "located in", "targets", "part of", “related to”, etc.;
- **directional**: an intrusion set is attributed to a threat actor, not the other way around. Therefore, the correct direction must be chosen.

In collaboration with *OASIS*, the French company *Filigran* has proposed adaptations of the STIX format to meet the specific needs of the fight against FIMI. The list of native and added objects by *Filigran* is available in appendix 4.8.

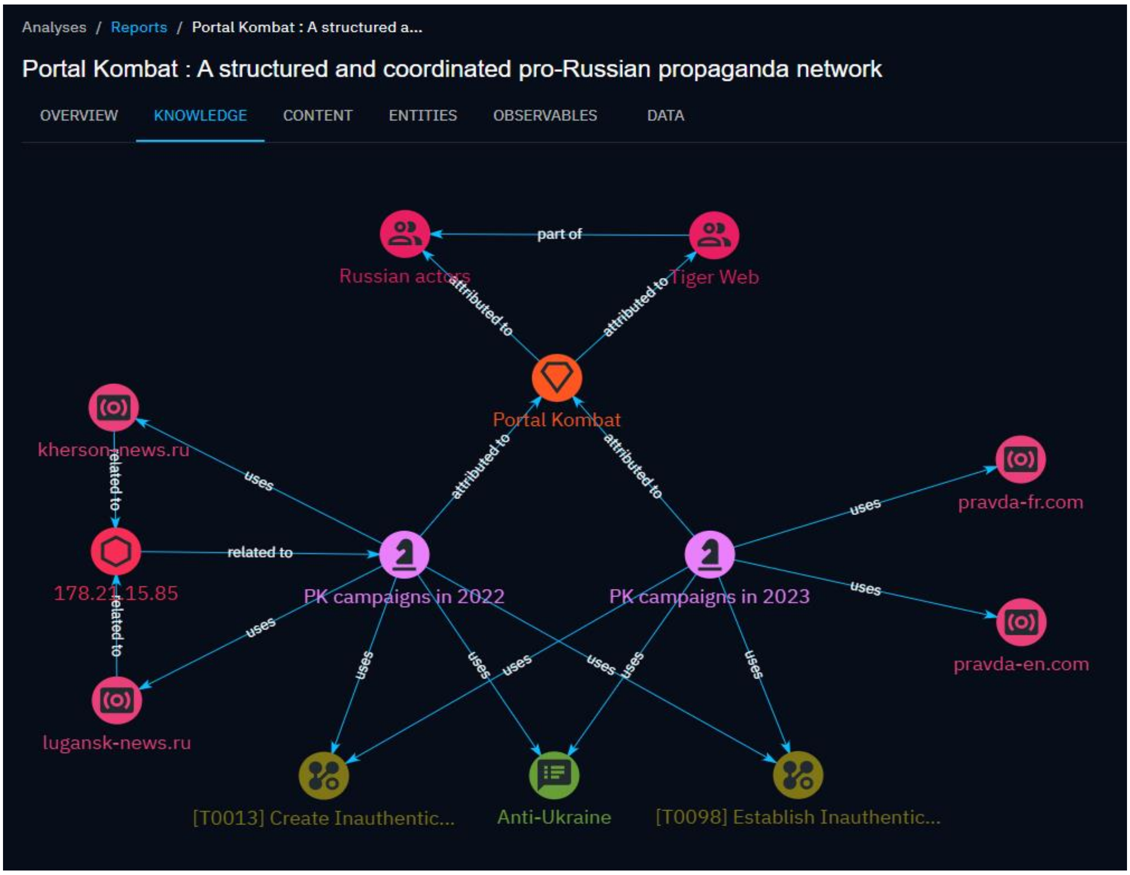

Example of knowledge graph

2.2 INFERENCE RULES

OpenCTI offers a set of predefined rules (named "Rules engine") that govern how new relationships are inferred based on existing data.

For instance, a user can activate the following inference rule: if an Entity A is attributed to an Entity B and this Entity B is itself attributed to an Entity C, then the Entity A is also attributed to Entity C.

In the same way, if the Intrusion set "RRN" uses the narrative "France in Africa", and that the Intrusion set "RRN" is attributed to the Threat Actor "Russian Actors", then "Russian Actors" uses the narrative "France in Africa".

Inferences are crucial for easily identifying:

- all the Intrusion sets that have targeted the same organization, sector, or individual;
- all the tools, narratives, tactics, techniques, and procedures (TTP) used by an Intrusion set;
- all the Intrusion sets related to a specific Threat Actor;
- all the technical elements (IP addresses, emails, phone numbers, etc.) associated with a Campaign, an Intrusion set or a Threat Actor.

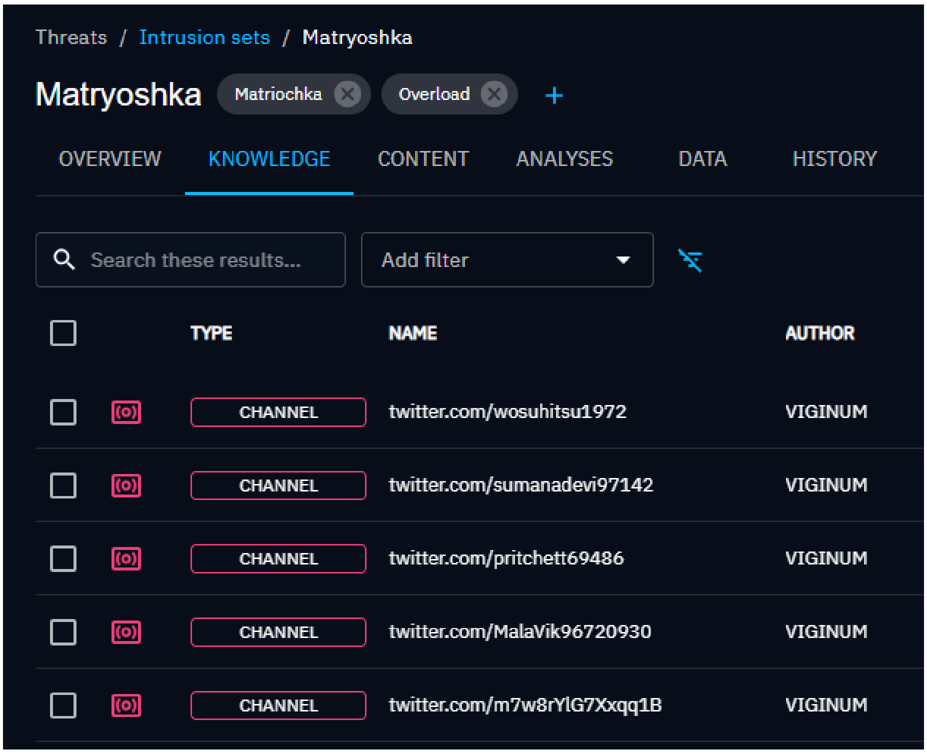

Threats / Intrusion sets / Matryoshka

Matryoshka Matriochka ⊗ Overload ⊗ +

OVERVIEW KNOWLEDGE CONTENT ANALYSES DATA HISTORY

Search these results... Add filter

□		TYPE	NAME	AUTHOR
□	((o))	CHANNEL	twitter.com/wosuhitsu1972	VIGINUM
□	((o))	CHANNEL	twitter.com/sumanadevi97142	VIGINUM
□	((o))	CHANNEL	twitter.com/pritchett69486	VIGINUM
□	((o))	CHANNEL	twitter.com/MalaVik96720930	VIGINUM
□	((o))	CHANNEL	twitter.com/m7w8rYlG7Xxqq1B	VIGINUM

Example of Channels used by an Intrusion set

3. MODELING THE INFORMATION THREAT KNOWLEDGE

Knowledge is imported on *OpenCTI*, both manually or automatically, by creating a new entry in the “Reports” section.

3.1 OVERVIEW OF THE REPORT DASHBOARD

The following entity details and basic information must be completed in the dashboard:

1. **Name**: the exact title of the report that will be imported on *OpenCTI*;
2. **Description**: a brief summary of the document's content;
3. **Report types**: internal report or external documents;
4. **Publication date**: the publication date of the document;
5. **External references**: a link to the original document (either from an open source or an internal database);
6. **Marking**: the classification or marking of the document (for instance, TLP:CLEAR or TLP:RED[4]);
7. **Status**: set to "New" by default, but can be adapted to internal workflow organization;
8. **Author**: the author(s) of the document.

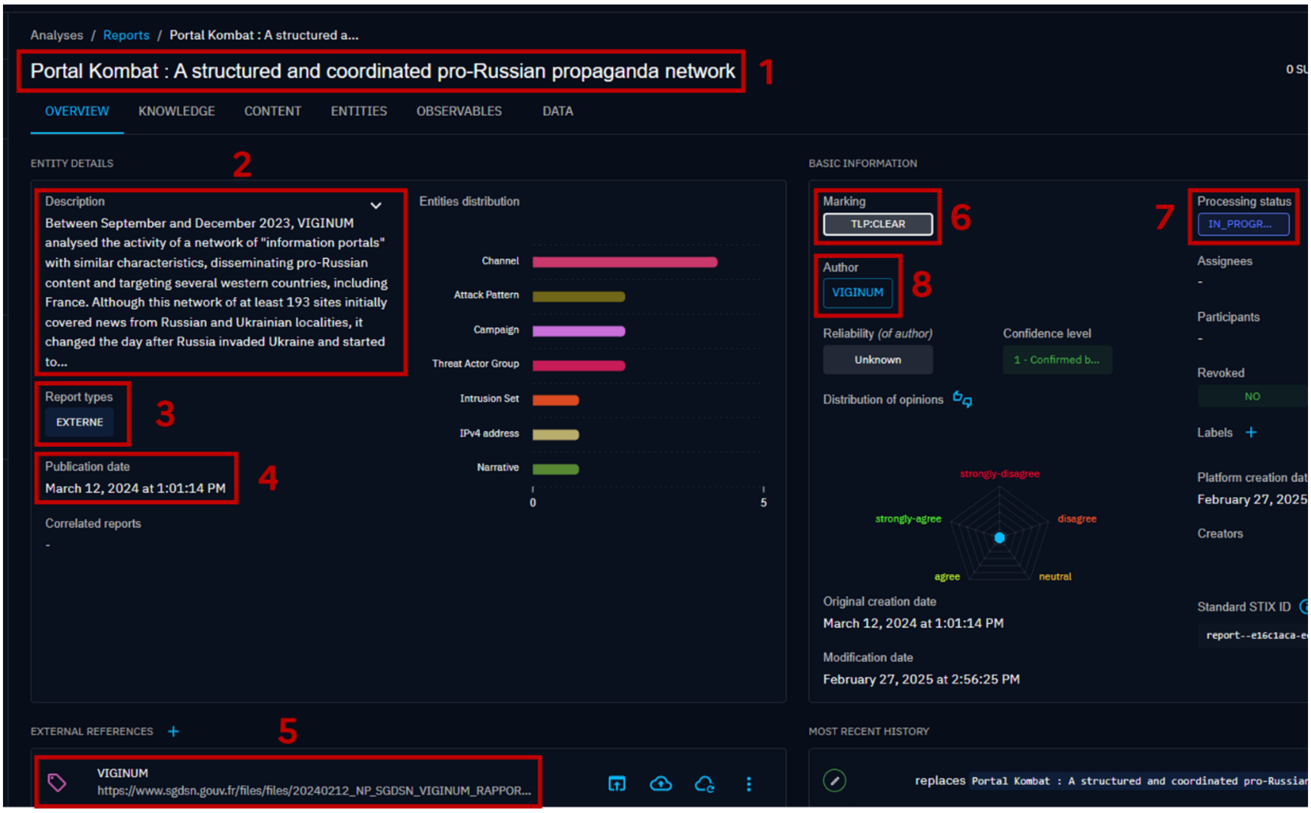

Portal Kombat : A structured and coordinated pro-Russian propaganda network

Description
Between September and December 2023, VIGINUM analysed the activity of a network of "information portals" with similar characteristics, disseminating pro-Russian content and targeting several western countries, including France. Although this network of at least 193 sites initially covered news from Russian and Ukrainian localities, it changed the day after Russia invaded Ukraine and started to...

Report types: EXTERNE

Publication date: March 12, 2024 at 1:01:14 PM

Correlated reports: -

Marking: TLP:CLEAR

Author: VIGINUM

Reliability *(of author)*: Unknown

Confidence level: 1 - Confirmed b...

Original creation date: March 12, 2024 at 1:01:14 PM

Modification date: February 27, 2025 at 2:56:25 PM

Processing status: IN_PROGR...

Assignees: -

Participants: -

Revoked: NO

Platform creation dat
February 27, 2025

Creators

Standard STIX ID
`report--e16c1aca-e`

VIGINUM
https://www.sgdsn.gouv.fr/files/files/20240212_NP_SGDSN_VIGINUM_RAPPOR...

replaces `Portal Kombat : A structured and coordinated pro-Russian`

Example of a report “Overview” window

[4] More information on the policy for sharing and using operational information on the following page (in French): https://www.cert.ssi.gouv.fr/csirt/politique-partage.

If relevant, additional elements can be completed:

- a reliability and confidence score;
- one or more labels[5];
- assignee(s);

3.2 OVERVIEW OF THE KNOWLEDGE DASHBOARD

The "knowledge" tab is where the graph that represents the threat knowledge is built. Entities and links created in this graph will then generate inferences, as mentioned above.

The information which can be encoded in the "knowledge" tab include attribution, victimology, operating methods and TTPs, along with the technical insights linked to a Campaign or an Intrusion set.

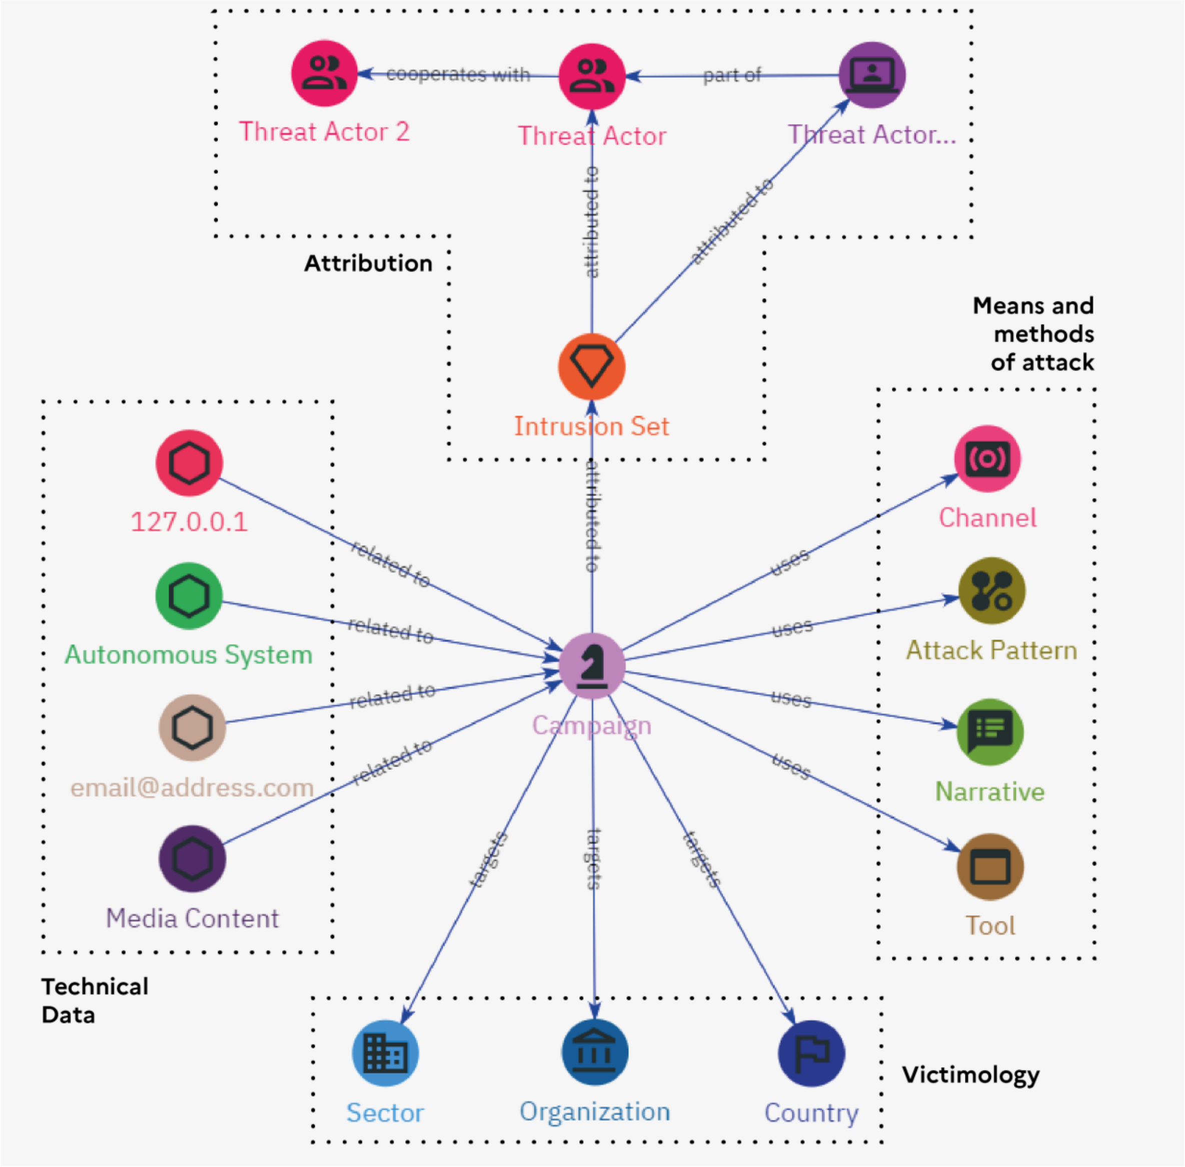

[5] Cf. Appendix 4.2.

The entities (IP addresses, narratives, targeted countries, etc.) **must all be linked to an annual Campaign**[6], which helps maintaining a temporal indicator of the activity of the Campaign / Intrusion sets and facilitates their analysis over time.

Below are the cases illustrating the different possible links (and their direction) between the relevant objects to model properly threat knowledge on *OpenCTI*. Other links between objects exist, but it is important to note that **it is always better to create a "qualified" link** (e.g., *attributed to, cooperates with, located at, targets*, etc.) rather than using the default link (*related to*), which does not generate inference rules.

3.2.1 Modeling attribution

Threat modeling is structured around a central "Campaign" object, which can be linked to an Intrusion set or directly to a Threat Actor, based on investigations and findings.

For instance, an Intrusion set may be attributed to one or more Threat Actors (organization or individual) which can belong to other Threat Actors or cooperate with each other. This model is particularly useful for representing organizational charts, such as attributing an Intrusion set to a known individual or a specific unit belonging to a foreign intelligence service.

Principles for attribution modeling:

1) **Always represent the original threat**

A report presenting a presumed state activity should always include a Threat Actor entity related to the information operation, as described in the screenshot above.

For instance, when importing an external report on the Intrusion set "RRN", publicly attributed to the Russian government, the Threat Actor "Russian Actors" entity must be included in the knowledge graph. Mentioning this entity in the graph allows users to find knowledge from this report on the related "Russian Actors" Threat Actor page.

2) **Linking an Intrusion set to a Threat Actor**

In the graph, intrusion sets are not necessarily attributed to a Threat Actor. In fact, this link is only relevant if the document **provides sufficient evidences** to establish a relationship between the Intrusion set and the Threat Actor.

This approach ensures the traceability of important technical attributions and official attribution statements. By doing so, users can closely track which governments, companies or researchers have attributed the Intrusion set "RRN" to the Russian government, when this attribution was made, for what reason, and its level of marking.

Examples:

- if the document provides technical elements linking RRN to Russian actors, a relationship "*attributed to*" between the Intrusion set "RRN" and the "Russian Actors" entity must be created;

[6] Cf. appendix 4.1.

- if the document is an official statement regarding political attribution, whether substantiated or not, it is important to keep a record of it: a relationship between the "RRN" Intrusion set and the "Russian Actors" entity must be created;
- if the document states that "RRN is likely related to the Russian government" but does not support this attribution, there is no need to create a relationship between the "RRN" Intrusion set and to "Russian Actors" entity.

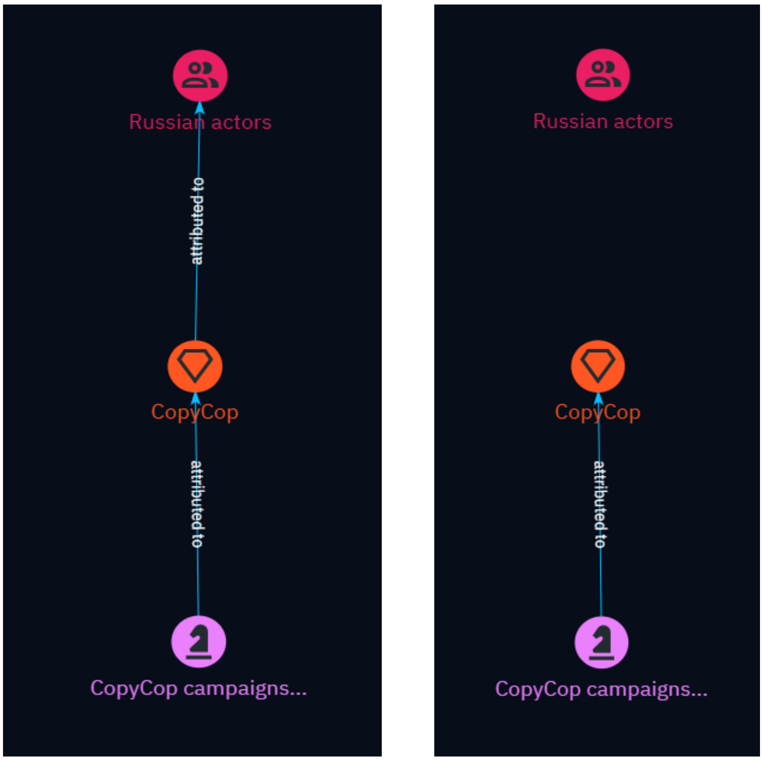

Modeling attribution

3.2.2. Modeling means of action & TTPs

The Campaign ***uses***:

- TTPs;
- Channels;
- Tools;
- Narratives.

Links can also be created between the different means or tools used, e.g. Channel ***uses*** a Narrative.

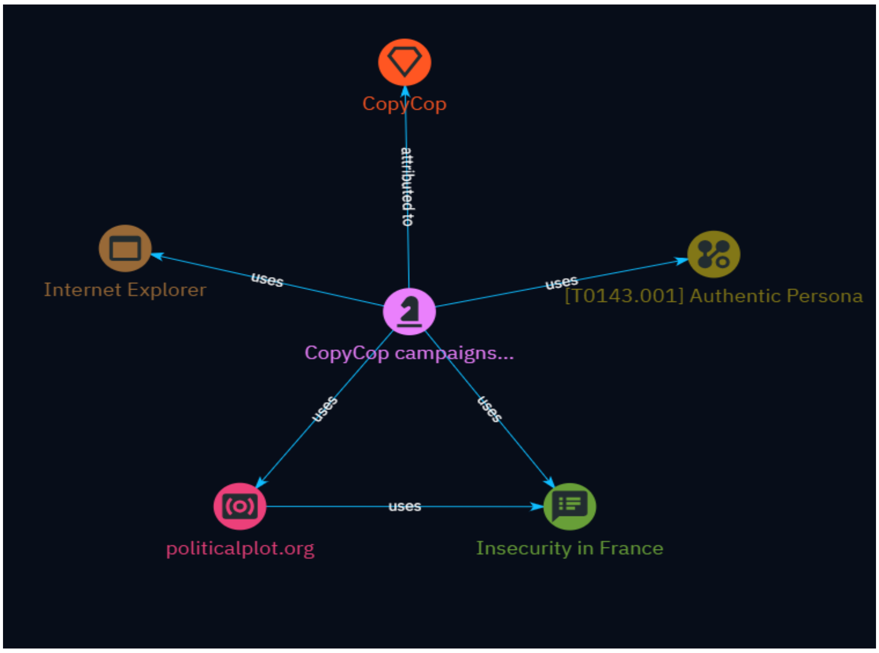

Modeling means of action & TTPs

3.2.3. Modeling technical elements (observables)

Observables are ***related to*** a Campaign:

- Email address;
- IPv4 address;
- Domain name;
- Telephone number;
- URL.

Technical information can also be linked together, for instance to specify that a Domain name ***resolves to*** an IPv4 address.

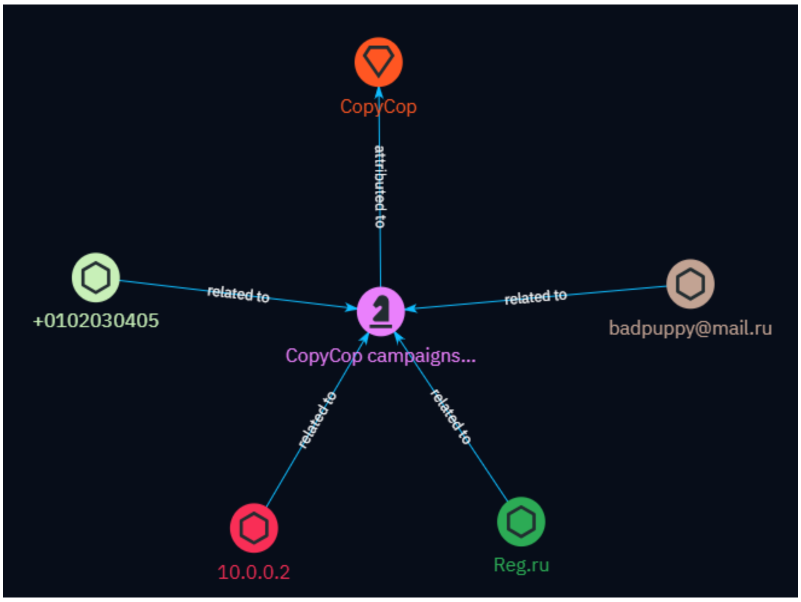

Modeling technical elements

3.2.4 Modeling activity on social media

A Campaign:

- ***uses*** one or more Channels;
- these Channels ***amplify*** other Channels.

Links can also be created to specify that a Channel ***publishes*** a media content, an URL or a Domain name.

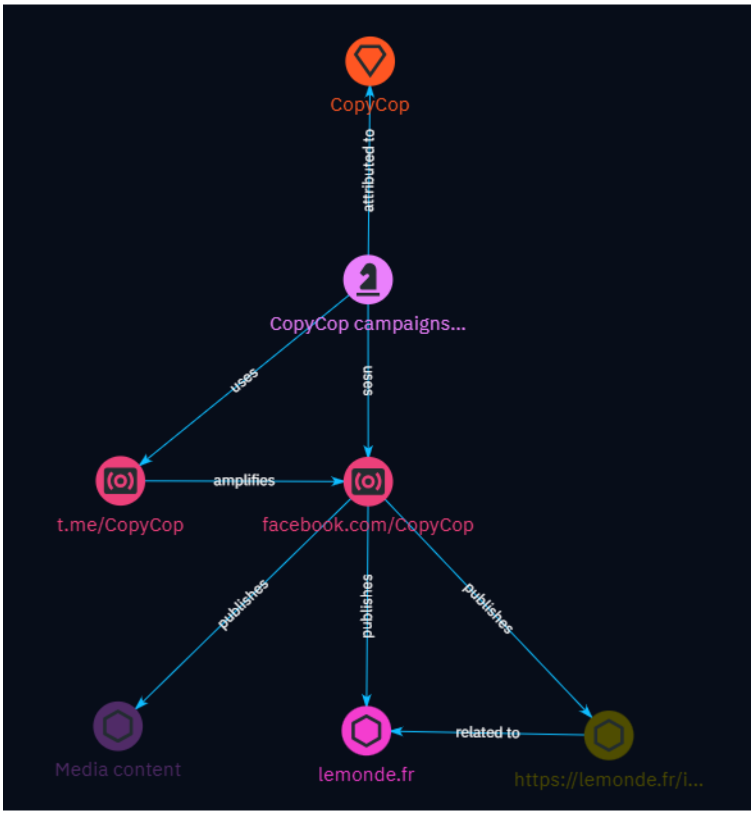

Modeling activity on social media

3.2.5 Modeling victimology

A Campaign ***Targets***:

- Sectors;
- Countries;
- Events;
- Individuals;
- Organizations.

Information on victimology may also be linked to each other, for instance to specify that an Event ***is located*** in a Country.

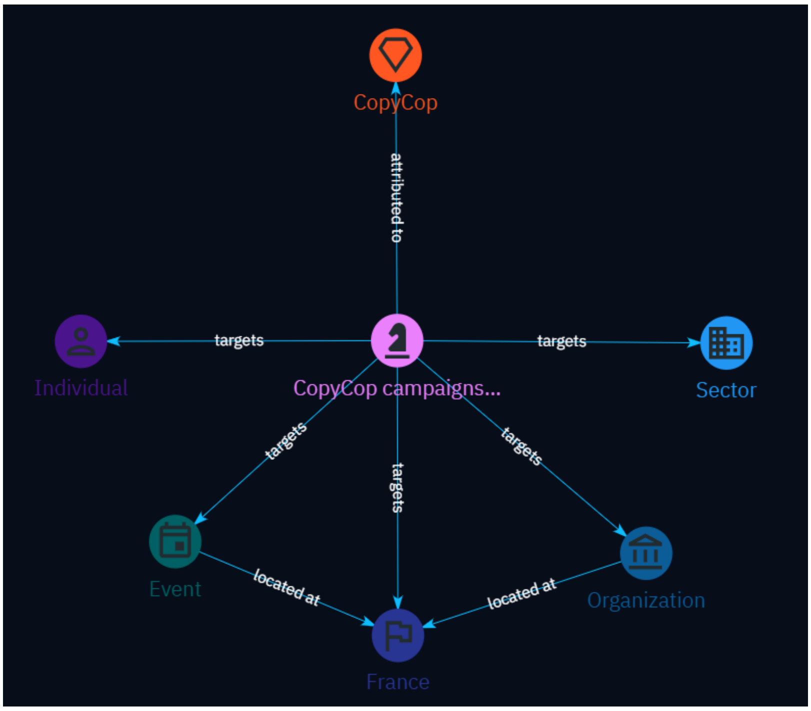

Modeling victimology

4. APPENDIX

4.1 HOW TO CREATE A CAMPAIGN

In order to standardize the Campaign titles, which are central to the knowledge management process, a naming convention has been established: activities of the Intrusion set must be linked to annual Campaigns that allow for chronological tracking of the evolution of their victimology, TTPs, and attributions.

Examples:

- "RRN Campaigns in 2023";
- "Storm-1516 Campaigns in 2019";
- "Portal Kombat Campaigns in 2036".

If an information operation has not been attributed yet, it is possible to create a campaign specifying that it is unattributed, the nature of observed malicious activities and a timeframe, for example:

- "Unattributed campaign using deepfakes of celebrities on *Bluesky* in December 2023";
- "Unattributed campaign targeting Emmanuel MACRON's *Wikipedia* page in July 2019".

If the campaign is subsequently attributed, it can be merged with the corresponding annual Campaign of a Threat Actor or Intrusion set. For example, if the first aforementioned unattributed campaign is later attributed to RRN, it will be merged with "RRN Campaigns in 2023".

4.2 HOW TO CHOOSE A LABEL

Labels can be added to almost any type of entity: Intrusion set, report, organization, etc. **They are primarily used to facilitate cross requests** on *OpenCTI*.

For now, we recommend using labels primarily for **cross-cutting themes** (e.g. "unattributed", "elections", "artificial intelligence", "Olympic Games", and the language leveraged by attackers, etc.).

4.3 HOW TO CHOOSE THE PROCESSING STATUS OF A REPORT

The processing status – also called the status –allows for quickly knowing whether a report has been finalized or not.

Currently, there are four processing statuses:

- NEW;
- IN_PROGRESS;
- ANALYZED;
- CLOSED.

NEW is the default status indicated for a report. It should be kept as it is until the report is capitalized.

IN_PROGRESS indicates that the report is being capitalized. This status allows the user to indicate that a report is not yet fully completed.

ANALYZED indicates that the report is fully completed. The status must be chosen when a report is fully capitalized.

CLOSED is a status reserved for errors and duplicates. The analyst can indicate this status to signal that the report should not be considered and should be deleted from *OpenCTI*.

4.4 HOW TO CREATE RELATIONSHIPS

For inferences to be made correctly, the links between the objects in a knowledge graph must:

1. be drawn in the "correct" direction;
2. be named appropriately.

Except for observables, **the links are drawn in a star pattern from the central entity, the Campaign**.

The main directions and naming types are listed below:

- a Campaign *is attributed* to an Intrusion set or a Threat Actor;
- a Campaign *targets* an Organization, a Sector, an Individual, or a Country;
- a Campaign *uses* an Attack pattern, a Tool, a Narrative, or a Channel;
- a Threat Actor *is located in* a Country;
- a Threat Actor *cooperates* with another Threat Actor;
- a Threat Actor *is part of* another Threat Actor;
- an Intrusion set *is attributed to* a Threat Actor;
- an IPv4 address, an Email address, a Phone number, a Domain name, or a URL *is related* to a Campaign;
- an organization *is part of* a Sector;
- an organization *is located in* a Country.

4.5 HOW TO FILL IN THE RELATIONSHIPS FIELDS

Links must be correctly created to produce relevant inferences. Most relationships need to be filled in, meaning justified by elements from the initial document.

The analyst must correctly fill in the following fields based on the elements contained in the document:

- The correct start and/or end dates of observation;
- The author of the document;
- The document marking.

4.6 How to date links

Every relationship can be dated. The **dates need to be systematically changed** when creating a relationship, as they will display otherwise, by default, the date of publication of the document in the fields "Start date of observation" and "End date of observation”.

The dating is very important as it helps maintaining a temporal indicator of the activity of the Intrusion sets and Campaigns.

4.6.1 Dating of victimology, means of action and TTPs

The dating of relationships regarding victimology, mean of actions and TTPs is the same. It focuses on the information provided in the document. There are four cases:

Case 1: The author observed the start of an activity (targeting an Organization, a Country, a Sector, use of a tool or attack pattern, etc.) and the end of this activity.

→ *The* "Start date of observation" *and* "End date of observation” *must be filled in*.

Case 2: The author observed the start of an activity, but it is ongoing.

→ *Only the* "Start date of observation" *should be filled in*.

Case 3: The author observed the end of an ongoing activity, but the start date is unknown.

→ *Only the* "End date of observation” *should be filled in*.

Case 4: The author observed an activity but did not mention any specific dates.

→ *No dates should be filled in*.

4.6.2 Dating of attribution

The dating of relationships regarding attribution differs from the dating of relationships concerning victimology and tools and methods: **it focuses on the document publication date.**

Regardless of the entity to which an Intrusion set or Campaign is attributed (actor, organization, individual), this implies:

- The "Start date of observation" always corresponds to the publication date of the capitalized document, which is normally indicated by default when creating the relationship;
- The attribution never has an "End date of observation", which should therefore be removed.

4.7 Description of the main STIX objects for FIMI

4.7.1 Entities (STIX Domain object)

Campaign

(Native STIX object)

The Campaign entity is the central element of the knowledge graphs on *OpenCTI*. All technical elements, attribution information, and other indicators are structured in a star pattern around an annual Campaign attributed to an Intrusion set, a Threat actor, or unattributed.

Threat Actor (group)

(Native STIX object)

On *OpenCTI*, a Threat Actor (group) refers to an organization (administration, company, informal group, etc.) that is responsible for, supervises, sponsors, or supports malicious activities.

The entity is used to:

- represent malicious organizations themselves: foreign intelligence service, company involved in information operations, hacktivist groups, etc.

- represent meta- Threat Actor or categories of actors. For instance, all organizations participating in pro-Russian information operations can be *part of* a meta-Threat Actor "Russian Actors", so they have a dedicated page on *OpenCTI*.

Threat Actor (individual)

(Native STIX object)

On *OpenCTI*, a Threat Actor (individual) refers to an individual who is responsible for, supervises, sponsors, or supports malicious activities. Examples include: a known member of a foreign intelligence service, an employee of a company, an influencer, etc.

Intrusion set

(Native STIX object)

When a specific actor behind a malicious activity is not known (see section 3.2.1), or when the tracking of information operations is managed by Information Manipulation Set (IMS), the Campaign is attributed to an Intrusion set.

VIGINUM defines an IMS as a collection of behaviors, tools, tactics, techniques, procedures and adversary resources used by a malicious actor or group of actors as part of one or more information operations. It should not be confused with the Threat Actors (group), consisting of organizations or individuals (cf. above).

Channel

(Added by Filigran*)*

A channel is a social media account or a domain name used by a Threat Actor or an Intrusion set in the context of an information operation.

To easily request Channels on *OpenCTI* and avoid creating duplicates (for example, when an actor uses both an *X* account and a *Telegram* account with the same string in the name), it is important to specify the type of Channel. Example of naming convention:

- *Facebook*: facebook.com/username (UserID, groupID, or pageID as aliases);

- *X*: twitter.com/username or x.com/username (pseudonym as alias);

- *Telegram*: t.me/channelorgroupname (group name as alias);

- *Domain nam*e: mechantemanoeuvre.ru.

If the analyst does not know any specific account involved in an information operation but still wants to represent, for instance, that it used *Bluesky* or *Twitter*, a generic Channel entity named “bluesky.com” or “x.com” can be used.

Event

(Added by Filigran)

The Event describes a real-life event exploited to conduct an information operation. These events should be broad enough to be targeted by multiple operations.

For example, it is preferable to use "Paris 2024 Olympic Games" rather than "Men's individual kayaking event on August 13, 2024, during the Paris Olympic Games."

Besides a sporting event, it could for example be an international meeting, an election, a diplomatic visit, a natural disaster, etc.

Individual

(Native STIX object)

A person who is not considered as a Threat Actor: target of a Campaign, individual whose identity has been impersonated, etc.

Organization

(Native STIX object)

A legal entity that is not considered as a Threat Actor: companies, NGOs, or associations targeted by information operations, media whose identity has been impersonated, etc.

Sector

(Native STIX object)

"Sector" entity allows to represent the targeting of some specific interests. Sectors must be quite broad and general. Here are examples of sectors that can be targeted by an information operation:

- *State Institutions;*
- *Economic, industrial, and scientific interests;*
- *Cultural heritage;*
- *State personalities;*
- *Foreign policy;*
- *National security and defense;*

This entity can also represent sectors of activity, such as the media ecosystem or NGOs.

Infrastructure

(Native STIX object)

This STIX object models the physical or virtual resources exploited by attackers, including hosting providers (*GoDaddy, NameSilo*, etc.) used during an information operation.

Narrative

(Added by Filigran*)*

This entity describes the narratives exploited by an Intrusion set or a Threat Actor during an information operation. These narratives must be broad enough to be exploited by multiple actors or over a relatively long period of time.

For example, favoring "Neocolonialism" rather than "Accusing France of sending migrants to the front in Ukraine in June 2024."

Attack pattern

(Native STIX object)

Entity used to describe attackers' actions based on the DISARM framework[7]: amplification by trolls [T0049.001], creation of personas [T0097], facilitate state propaganda [T0002], etc.

Tool

(Native STIX object)

This entity describes legitimate software or online services exploited by an Intrusion set or a Threat Actor during an information operation, for example a commercial solution for setting up redirect links or managing bots on social media.

4.7.2 Observables

Observables are designed to model the technical elements observed during an information operation. They are particularly useful for representing a digital infrastructure exploited to conduct an operation. Here is a non-exhaustive list of the most relevant observables to document information operations:

- Email address;
- IPv4 address;
- Media content;
- Domain name (note that **attacking sites are not listed as Domain names, but as Channel**, see 4.8.1);
- Phone number;
- Autonomous system;
- URL.

	Malicious	**Legitimate**
Entity type	Threat Actor (Organization)	Organization
	Threat Actor (Individual)	Individual
	Channel	Domain name

Summary table of the usage of entities

[7] Cf. https://www.disarm.foundation/.

Media Content

(*Added by* Filigran)

This observable models types of media content, such as social networks publications, videos, audios, pictures, sponsored content, etc. Naming convention:

- [Video] Title of the video;
- [Article] Title of the article;
- [Podcast] Title of the podcast;
- [Tweet] Title of the tweet;
- [Facebook] Title of the post;
- [Sponsored content] Title of the post;

To create a Media Content observable, an URL must be provided.

ABOUT VIGINUM

Created on July 13, 2021, and attached to the General Secretariat for Defence and National Security (SGDSN), France's service for vigilance and protection against foreign digital interference (VIGINUM) is intented to protect online public debate which affects France's fundamental interests.

This technical and operational state agency is responsible for monitoring and defining information manipulation campaigns on digital platforms, involving foreign actors with the aim of damaging France and its interests

Service de vigilance et protection contre les ingérences numériques étrangères | SGDSN

Cover photo credits: Cai Fang on Unsplash