



TAMM:

Threat Actor Mapping Methodology

DISRUPT Toolkit

Author: Saman Nazari

**Alliance
4Europe**



Table of Content

Executive Summary		/01
Introduction		/02
Definitions		/03
Threat Actor		/11
Influence Operations		/11
Threat Actor Attribution		/11
Mapping		/11
Information Manipulation Set (IMS)		/11
Taxonomies		/05
Taxonomy of Relationships		/05
Our Precise Labels		/06
Owned Entity		/06
Controlled Entity		/07
Funded Entity		/07
Supported Entity		/08
Shared Infrastructure		/08
Service Provider		/09
Systematic Amplification		/09
Methodology for Mapping Threat Actor Ecosystems		/10
Process and Evidence Indicators		/11
1. Establish Known Threat Actors and their Assets		
2. Investigate Known Assets for Indicators		
3. Draft Profiles and Mapping		



Taxonomy of Influence Operations		/16
The Threat Actor Ecosystem		/17
Threat Actor		/17
Information Manipulation Set (IMS)		/17
Assets		/18
Behaviours		/18
Influence Campaign		/19
Influence Operation		/19
Taxonomy of New Information Manipulation Set (IMS)		/21
Threat Actor Ecosystem		/22
IMS Type		/22
Sources		/25



Executive Summary

Across the counter-FIMI and influence operations community, defenders attribute the same channels, map the same networks and name the same operations, and much of this work is never published.

Threat actors are growing more sophisticated while defender funding shrinks, so the community can no longer afford this duplication.

The Threat Actor Mapping Methodology (TAMM) sets out a common framework for conducting, classifying and sharing threat actor attributions. It is a shared taxonomy that will let attribution data be exchanged through Alliance4Europe's TRANSCRIPT platform.

It builds on the ADAC.IO attribution frameworks (Cardiff and Lund Universities), VIGINUM's Information Manipulation Set (IMS) concept, EEAS FIMI reporting and the DISARM Red Framework.

TAMM proposes **three taxonomies**:

- **Relationships:** neutral, granular labels (owned, controlled, funded, supported, shared infrastructure, collaborates, service provider, systematic amplification), each marked as overt or covert. Limiting claims to relationships that can be proven reduces both room for interpretation and legal exposure.
- **Influence operations:** a structure linking threat actor ecosystems, threat actors, IMS, campaigns, operations, assets and behaviours. It adapts VIGINUM's model to cases where a direct sponsor link cannot be established.
- **IMS naming:** a descriptive convention (Ecosystem + IMS Type) to replace today's ad hoc names.

TAMM also sets out **a mapping method built on content, behavioural, technical and financial signals and the ABCDE framework**. No attribution rests on narrative alignment alone, or on a single non-technical indicator.

TAMM is published as the basis for a public consultation. Its taxonomies are proposals to be refined with practitioners.



Introduction

Throughout the counter-disinformation, FIMI, and influence operations community, democracy defenders are generating information about threat actors and their activities. As the threat actors are becoming more sophisticated and funding more scarce, **the need for efficiency has never been more apparent**. As the defender community works to counter these threats, duplication is happening across different activities. Multiple researchers end up working on the same investigations, conducting the same attributions and threat actor mappings, much of which is never made public.


There is a clear need to build upon each other's knowledge, instead of spending the few resources we have on attributing the same channels and actors over and over again.

To simplify this work, **Alliance4Europe has launched the Threat Intelligence Database and Coordination Platform (TRANSCRIPT)**, a threat intelligence, knowledge base, and disruption coordination system that aims to, amongst other things, enable the sharing of attribution data.

To effectively use this system to share threat actor data between different organisations and countries, **a common taxonomy is needed**.

This methodology paper outlines the **Threat Actor Mapping Methodology (TAMM)**, a methodology that contains a framework and taxonomy to help organisations conduct and classify attributions. It also contains a taxonomy for classifying Information Manipulation Sets (IMS). This paper will act as the foundation of a public consultation, aiming to develop a general approach that fits the needs of all actors.

The methodology builds on the **ADAC.IO Horizon project** publications by the University of Cardiff and Lund University (Innes & Ahonen, **2025**; Palmertz, Isaksson, & Pamment, **2025**), which cover an evidence framework for attribution and best practices on attribution.



Definitions

Before diving into outlining the methodology and taxonomy, we will define the concepts that underlie our work.

Threat Actor

We define a threat actor as an entity (state, organisation, group, or person) utilising influence operations to achieve their objectives, whether political, financial, reputational, or any other reason.

Influence Operations

An influence operation is an organised, deliberate effort to manipulate the beliefs, behaviours, or decisions of a target audience using illegal content, illegal activity, inauthentic behaviour, covert activity, or foreign interference ([Nazari & Sarvela 2026](#), p.11).

We do not utilise the term information operation, as the term indicates that the operation is trying to spread information or inform. The concept that the term refers to does no such thing, instead referring to an activity that tries to manipulate its target audience.

Threat Actor Attribution

We define Threat actor attribution as the act of showing the relationships between one actor and another or tying a social media asset (e.g., a social media page or account) to an actor.

We do not use actor attribution to refer to the act of attributing a specific actor to an action.

Mapping

We define mapping as the act of analysing an information space to identify the different social media assets and web assets affiliated with threat actors. Mapping can also identify coordinated inauthentic behaviour networks that display a specific behaviour but only utilise disposable accounts.

Information Manipulation Set (IMS)


IMS is a new concept by VIGINUM that tries to expand upon the Coordinated Inauthentic Behaviour Network concept.

VIGINUM defines an Information Manipulation Set as a “collection of adversarial behaviours, tools, and Tactics, Techniques, and Procedures (TTPs) presumed to be linked to the same Threat Actor or group of Threat Actors, which may be unknown” (Viginum, 2026, **p.5**)

VIGINUM outlined in their definition that two criteria need to be fulfilled for the observed collection of activities, actors and infrastructures to be an IMS.

1. Clandestinity - malicious intent and the covert nature of the structure and activity observed
- trying to avoid attribution.
2. Coordinated - Social media or web assets presumed to be operated in a coordinated manner by the same actor or group of actors working for the same sponsor.

An IMS cannot operate under their own names or be unattributed and utilising different infrastructures, while sharing the same branding (Viginum, 2026, p. ~~6~~-7)





Taxonomies

This methodology will propose three taxonomies.

The **first taxonomy** will neutrally describe the relationships between different actors. This taxonomy differentiates itself from other taxonomies, such as the EEAS actor taxonomy, which is divided into Official Communication Channels, State-Controlled Channels, and State-Linked Channels (EEAS, 2023, p. [20](#)), by being more granular and neutral.

The taxonomy is designed so you can easily and descriptively show the relationships, leaving little room for interpretation. This provides the researcher a degree of legal protection, as it's hard to dispute factual and proven claims about an actor, e.g. funding or partnerships.

The **second taxonomy** aims to outline the relationship between threat actors, IMS, influence operations, and assets. We will build on the efforts of the EEAS and VIGINUM to develop an easy-to-use and versatile taxonomy.

The **third taxonomy** proposes a naming convention for IMS, making it easier to name IMS when they appear. The design philosophy of this new IMS naming taxonomy is to ensure clarity and simplicity, making it accessible to a range of practitioners.

Taxonomy of relationships

In 2025, the EEAS proposed a 4 level relationship taxonomy (EEAS, 2025, p. [15](#)):

Official State Channels: "Official communication channels Directly operated by a government and its representatives, openly reflecting the state's official voice" (Ibid).

State-Controlled Outlets: "Media outlets funded, managed, and editorially controlled by state-appointed bodies or the ruling party. These outlets deliver the editorial lines set by the state" (Ibid).

State-linked Channels: "Channels that operate under state oversight without publicly disclosing their affiliation. Uncovering their connection to the state requires a combination of

evidence to reveal hidden patterns of influence or indirect control” (Ibid).

State-aligned Channel: “Channels that cannot be directly attributed to state control or funding and therefore remain unattributed but show systematic signs of alignment with a state entity. They fall into this category until more evidence emerges to confirm attribution to state infrastructure” (Ibid).

The EEAS approach is not granular about the relationships that exist between the actors, leaving room for interpretation. Legal liability also emerges when it comes to state-aligned channels, as it can be seen as a subjective assessment.

Our Precise Labels

The following labels describe the relationship between different actors. One actor may have relationships with more than one actor, including different threat actors.

The labels can also be combined. The nature of the threat actor should be added:

E.g. Chinese State-funded and Editorially Controlled Media.

Each relationship can also be covert or overt.

A **covert relationship** means that they are not publicly divulging the relationship. You can’t find it on their website or in their communication. They are taking efforts to exclude or hide the relationship. A relationship that is known through third parties can still be covert if the actor is still not confirming or declaring it.

An **overt relationship** means that it has been announced and is publicly available. If you look it up, you can find the information.

In case you are not sure about the attribution, add “unknown”.

Owned Entity

An actor can be owned by another actor. Ownership usually entails the legal entity being owned fully or partially by another actor. This can be through shares or other legal structures. A government legally owns its agencies, unless they are set up as independent bodies.

If a minority of the legal entity is owned by the actor, make sure to declare it clearly. You can call it a minority ownership.



Examples include one actor owning another actor, including a company being owned by an oligarch.

Note: watch out for satire and impersonation channels.

While rare in the context of IO, an owned entity might still retain its independence and not be controlled if the board and management of the entity are left alone.

Controlled entity

An entity can be controlled without being owned. Control can stem from a series of different structures, laws, or relationships. You can tag a report with both owned and controlled, or choose only one of them if the other does not apply.

Examples:

- A media outlet operating in a country without editorial freedom, where the state provides editorial instructions, e.g. China.
- An entity's board is controlled by members of a political party.
- Journalists in repressive environments whose online presence needs to be approved by the state and is monitored.

Examples of Indicators:

- The actor publicly admits to being state-controlled (easily accessible or legacy pages stating that they are the national media of a country or the official media of an entity).
- Has been attributed by trusted sources using proprietary or classified sources.
- The legal framework of the country forces control.
- The board of directors or the president is set by the threat actor.

Funded entity

An entity funded wholly by an actor. Funding may not necessarily mean that they are controlled.

Examples:

- A media entity gets its funding completely covered by a state.
- A person employed by a threat actor-controlled entity.
- A researcher or journalist at a state-controlled university or media.

Examples of Indicators:

- Tax or other administrative documents.
- Annual report.
- Public declaration of funding on their website or social media.
- A dedicated page on a threat actor's website.

- Description in content byline.

A freelancer is not a funded entity, but rather a service provider (see below).

Supported entity

Partially funded or supported by an actor.

Funding can come in the shape of grants, subsidies, and donations.

Examples: an actor provides another actor for free:

- IT infrastructure,
- Offices.
- Venues (repeated).
- Recording studio.
- Regular trips.

Examples of Indicators:

- Address of the venue for an event,
- Office addresses,
- IP address.
- Tax documents and other official documents concerning subsidies and donations.
- Public statements about funding.

Be mindful that there is a service provider relationship that you can use for purely financial relationships.

Shared Infrastructure

Shares infrastructure with another actor, not as a donation. This can be used when you are not sure if the actor is providing the infrastructure for free.

Examples: actors sharing:

- IT infrastructure.
- Offices.
- Venues (repeated).
- Recording studio.

Examples of Indicators:

- Address of the venue for an event.
- Office addresses.
- IP address.



Collaborates

An entity that has a partnership with another actor without the exchange of money.

Examples:

- Shared offices.
- Organise events together.
- Content sharing agreement.

Examples of Indicators:

- Address of the venue for an event.
- Office addresses.
- Public statements about partnerships.

Service Provider

An entity that has partnerships or commercial relations with a threat actor.

Examples:

- Supplies another actor with staff.
- One actor provides another with content production.

Examples of Indicators:

- Public tender results
- Tax documents and other official documents concerning business transactions.
- Public statements about partnerships or commercial deals.

Systematic Amplification

Support can also come from one actor systematically promoting a channel or actor, or systematically sharing their content, giving them exposure to the actor's existing audiences.

Examples:

- An expert is regularly being asked for comments by state-controlled media.
- Diplomats regularly share content from specific social media pages.
- State-controlled influencers are doing regular collaborations with unattributed influencers.

Examples of Indicators:

- Threat actor-owned channels share content at regular intervals from the amplified entity.
- Threat actor-owned channels share a regular amount of content from the amplified entity, e.g. every month.
- A threat actor-owned channel's promotion of content from the amplified entity stands out in comparison to other channels being amplified.



Methodology for Mapping Threat Actor Ecosystems

The following section will outline the Onion Threat Actor Mapping Methodology. It relies on starting from known information about the threat actors and digging deeper from there, finding more layers of threat actors working for them, and the evidence that connects them.

To start mapping a threat actor ecosystem, you need to identify the central point of the ecosystem, like a government agency, and work outwards. A good best practice is to start with the Ministries, state media, embassies, and already attributed threat actors, and work to identify who these threat actors are engaging with in different ways.

From this starting point, you can use a myriad of evidence to map the relationship between different actors.

The [Lund](#) paper outlined 4 types of evidence:

1. Content Signals

Analysing images or video data contained within messages and/or websites to identify who created them.

2. Behavioural Signals

Behavioural analysis is concerned with how inauthentic activities are organised and conducted.

3. Technical Signals

Focusing on analysing technological infrastructure used in influence operations, e.g. domains, servers, IP addresses, SSL certificates, and other technical components that make a website or web service run. When previous investigations have linked these components to an actor, they can be used to attribute future operations.

4. Financial Signals

Following monetary flows allows us to connect one actor to another, and establish the direction of



he relationship, e.g. who funds / controls who.

The EEAS in their 2025 report (EEAS, 2025, p. [17](#)) lump these into two categories.

- **Technical Evidence:** Including Public affiliations, financial records, shared infrastructure (e.g. IPs, social credit numbers).
- **Behavioural Evidence:** Systemic interaction and amplification patterns, reiterated use of inauthentic media assets, Coordinated messaging, Historical consistency and interaction with other operations.

Both of these are solid models for connecting actors. The EEAS model is the most operational of the two.

Process and Evidence Indicators

The following section outlines the process and indicators used to map threat actor ecosystems.

Mapping a threat actor ecosystem entails identifying relationships between threat actors and other threat actors, but also identifying their assets, such as social media pages and websites.

As with everything else, this is an initial draft to start a conversation.

1. Establish Known Threat Actors and their Assets

The first step is to establish an understanding of the Threat Actors that have been identified by previous researchers. From this point, you can establish an understanding of the threat actors known assets, e.g. social media accounts, websites, servers, known addresses, and other similar assets.

2. Investigate known assets for indicators.

After you have identified the known actors and their assets, you can set up monitoring of those assets and analyse their behaviours and components.

Most commonly, you will, through known assets, identify other assets, through which you can identify threat actors to investigate and place within the threat actor ecosystem.

Using the sets of evidence outlined by Lund and the EEAS, we propose the following indicators that can be used to establish a relationship between one threat actor and another.

Public-Affiliation

The actor announced their relationship with another actor on social media, to the press, on their website and through other means. The best way to identify this type of affiliation is by using a Search Engine and social listening tools to check for any keywords, such as “Threat Actor A” AND “Signed” AND “Threat Actor B”.

Such messages might also be archived by other researchers, so do check websites such as the Wayback Machine.

Website footers or dedicated sections on websites might also outline the partnerships the threat actor has.

Financial Records

Financial records can also help you identify any financial links between one actor and another.

Invoices, tax records, ownership records, public stock listings, end of year reporting and other similar documents might contain important information of ownership, business dealings, and partnership agreements. They can also contain information on members of the board, addresses, and other indicators that are useful for tying one actor to another.

In some countries, some of these documents are public. In other countries, you have to request them from the tax authorities or other instances.

Court rulings might also publicise financial records, which can be relevant in relation to sanctioned actors.

Financial Assets

Influence operations also commonly use financial assets of different sorts that can tie them to a threat actor, or tie threat actors to each other.

A strong indicator is the re-use of payment provider assets, such as paypal accounts, crypto wallets, Patreon, Credit Card payment provider accounts, and other means of receiving payments. These can be found on websites or linked to via social media.

The creation and verification of some of these are cumbersome, which leads to threat actors re-using them across different threat actors that they own.

Ad IDs is another common indicator that an asset or actor is affiliated with another.

Ad IDs are identifiers used by advertisement providers to ensure that ad revenue from a website goes to the right user. Threat Actors commonly use the same Ad Id across many different assets, and if the same threat actor owns several other threat actor organisations, such as subsidiaries, they might use the same Ad Id across all of them.



Physical Evidence

A threat actor might own, rent, or use physical property that tie them to another actor. Shared offices, the use of the same PO box, and repeated operations within a specific area might provide indicators of relationships with other actors.

Personnel

At the end of the day, threat actors and influence operations are owned and run by people.

While it is common for contractors to work for several threat actors, identifying these and other types of personnel might be of great help to tie one actor to another.

A threat actor might employ a journalist who works under a separate brand, but it is paid by the threat actor, trying to obfuscate the relationship.

Knowing the shareholders, CEOs, editors, and other senior staff of a threat actor organisation might also help identify links between them and other threat actors, as these actors might interact with or even operate other threat actor organisations.

Website

Threat actors might own and operate a website. A website has several components that make it work, which we can analyse to connect one actor to another. These components include IP addresses, server provider, domain provider, SSL certificate, content delivery networks, usernames of administrators, folder names, and software used to run the website.

These components are good starting points for identifying other actors, as sharing infrastructure within a threat actor ecosystem is not uncommon.

Together with other indicators, these can show potential ownership of one threat actor by another threat actor, or other types of relationships.

Social media

Social media accounts have many different indicators related to it that can give us evidence of a relationship between threat actors.

Following the Russian full-scale invasion of Ukraine and China's move to more clandestine influence operations, a myriad of assets changed their names and removed indicators of their affiliations. Checking old names, the first posts made by the channels, the dates of creation, and the origin of their initial logos can be helpful clues for identifying how a threat actor is placed within the threat actor ecosystem.

Furthermore, threat actors often use social media to share information about events, meetings, interviews, and other pieces of information that can be used.

Systematic interactions and Amplification

If a threat actor shares content from another actor once or twice, it can be seen as a coincidence. If they continuously share their content, continuously translating it to other languages, or republish content while removing its origin (information laundering), then one can establish that the threat actor is a systemic amplifier of the other threat actor.

Other types of systemic interactions are also relevant, e.g. continuously collaborations, joint content production, or event organisation.

Systemic amplification or interactions do not necessarily mean that there is any strong relationship between the actors beyond these actions themselves. The behaviour is still noteworthy and warrants further monitoring and analysis.

Coordinated Inauthentic Behaviour

Coordinated Inauthentic Behaviour (CIB) networks, or IMS, usually have distinct patterns of behaviours that can be used to track them and attribute them to a threat actor.

These can range from posting timing, social media profile composition, website re-routing methods, as well as sharing, liking, viewing, and posting patterns.

The content itself can also be a strong indicator.

Some IMS continuously target the same small group of people, or rely on impersonation.

When one tries to connect an IMS to a threat actor, CIB patterns of behaviour and infrastructure can be used to identify common service providers, tools used by both actors, and other indicators that the threat actors are part of the same legal structure.

If an IMS is using the same patterns of behaviours, technologies, and tools as another, it is likely that this IMS is being run by the same threat actor.

3. Draft Profiles and Mapping

Make sure to preserve all the evidence you have identified in the previous step through archiving. This can be done through offline or online archiving solutions, such as Hunchly or the Wayback Machine.

When you have mapped the threat actor ecosystem, having established the threat actors in it, their relationships, and assets, you need to aggregate this information.

The Threat Intelligence and Coordination Platform (TRANSCRIPT) aims to provide a common database for democracy defenders to aggregate this information into one place.




TRANSCRIPT allows you to aggregate your threat actor mapping in one place, visualizing it and enabling others to build upon it, similar to how a Wiki works.

In the platform, you can create profiles for each threat actor, provide a summary of who they are, which can be used in future reports, connect them to specific assets, and show the relationships they have with other threat actors, using the taxonomy covered earlier in this report.

Transcript also enables researchers to quickly pull in information about threat actors to their investigations. Whenever a researcher links to a piece of content from an already identified threat actor asset, the system automatically tags the threat actor and their relevant asset to the case.

When covering threat actor mapping in a report it is important to be very precise in the language you use. Never publicly call an actor a threat actor without having clear and verifiable proof of them being one. The relationships also need to be clearly outlined, leaving no room for interpretation. Threat actors might take legal action against you, trying to suppress your work.

By identifying the relationship between threat actors, you might be able to also open up new legal actions that you can take against them. E.g. if a threat actor is supplying a sanctioned threat actor with services, that is most likely a violation of the sanctions, and the supplier can be fined or jailed. By knowing who owns a threat actor, one can also enact new sanctions against these entities, or take other legal actions against them.





Taxonomy of influence operations

From the threat actor ecosystem to the influence operation

01 The Threat Actor Ecosystem

02 Threat Actor

03 Information Manipulation Set IMS

04 Assets

05 Behaviours

06 Influence Campaign

07 Influence Operation

In the VIGINUM taxonomy of an IMS, they outline 5 levels, divided into Policy (Sponsor), Strategic (threat actor and Information Manipulation Set - IMS), Operational (information campaign), Tactical (information Operation), and Technical (Digital Assets).

For the purposes of VIGINUM, the concept is well defined and delimited, but for the objectives of our work, we propose another taxonomy that includes IMS.



The Threat Actor Ecosystem

The threat actor ecosystem encompasses a cluster of threat actors who have a common denominator, such as affiliation with a specific state. More than anything, this is a way to cluster actors and make it easier to navigate. It cannot be purely ideological, but has to be behavioural.

The sponsor approach proposed by VIGINUM is more objective and granular, but limits the scope of research to cases where we can directly establish financial connections.

Examples of Threat Actor Ecosystems include the Russian State, the Chinese Communist Party, and Ilan Shor.

Each ecosystem has some legal entity or person that is at the centre of it, such as a government, a political party, a company, or a person (e.g., an oligarch).

Threat Actor Ecosystems might eventually merge, which is what has happened with the Yevgeny Prigozhin and Ilan Shor ecosystem that has now merged with the Russian State ecosystem.

Threat Actor

A threat actor is an actor who conducts influence operations. By identifying who these are, we can take legal action against them and also make a stronger case for why we need to stop them, e.g. if it is a foreign or terrorist actor.

The threat actor ecosystem contains several threat actors. An actor can have a myriad of relationships with the threat actor ecosystem, but also with others. A threat actor might be a part of several threat actor ecosystems, e.g. a company that provides services to several different states.

Examples include Russia Today, Sputnik, and CGTN, which all hide their relationship with the Russian state and the CCP while also operating different influence operations.

While a threat actor can be operating or being part of an IMS, they can also be running influence operations independently, using only their own assets.

Information Manipulation Set (IMS)

An IMS consists of manipulative behaviours conducted by one or several threat actors.

Mapping an IMS enables us to track their activities, understand how they circumvent social media platform policies or other safeguards, and coordinate disruption measures against them.

They can operate persistent or temporary assets to spread or create their content.

As stated before, these operations need to be malicious in intent and covert in nature - trying to avoid attribution. An IMS also needs to have one or a group of threat actors working in the same threat actor ecosystem, preferably for the same funder. The threat actor does not need to have been identified before you can label the activity as an IMS. It also needs to conduct coordinated activity using social media or web assets.

Both an IMS and a threat actor can conduct an influence campaign.

Assets

Assets are resources used by Threat Actors to create or execute influence operations. By understanding what assets they use, we can disrupt their operations.

Assets are divided into Web Assets, Social Media Assets, Communication Assets, Physical Assets, Financial Assets, Software Assets, and Legal Assets. This section builds upon the DISARM Red Framework's Establish Asset Section.

Web assets examples: servers, domains, SSL certificates, Content Delivery Networks, and other similar types of technical components that make a website or web service function.

Social Media Assets examples: social media accounts, channels, pages, or groups. Any type of persistent entity that can post content or interact with content.

Communication Assets examples: emails, phone numbers, addresses, and voice over internet protocol (VOIP) accounts.

Financial Assets examples: payment provider accounts, bank accounts, credit cards, crypto wallets, and advertiser accounts.

Software Assets examples: computer programs, Mobile Apps, Games, Game Mods, and malware.

Physical Asset examples: offices, studios, billboards (digital and analogue), printing press, speakers, cameras, and microphones.

Behaviours

Behaviours are the tactics and techniques used by the threat actors to conduct influence operations.

To cover these, we rely on the DISARM Red Framework. The DISARM Foundation is developing a 2.0 framework, but until its completion, we'll rely on V1.7 and its future updates.



By understanding and aggregating the behaviours used during influence operations, we can:

- Identify patterns of behaviours commonly used by a threat actor or IMS,
- Identify vulnerabilities they are exploiting in, e.g. social media platforms.
- Expose the behaviours to citizens to demonstrate their manipulative nature - educating them but also providing legitimacy to the takedown of the operation.

We have also made a slimmer version of the DISARM Framework that removes the techniques that are already covered in other parts of our Taxonomy of influence operations or do not provide added value for IO Disruption. You can find it [here](#).

As stated before, these operations need to be malicious in intent and covert in nature - trying to avoid attribution. An IMS also needs to have one or a group of threat actors working in the same threat actor ecosystem, preferably for the same funder. The threat actor does not need to have been identified before you can label the activity as an IMS. It also needs to conduct coordinated activity using social media or web assets.

Both an IMS and a threat actor can conduct an influence campaign.

Influence Campaign

An influence campaign is a collection of influence operations conducted by one or several threat actors part of the same Threat Actor Ecosystem and IMS.

The campaign either has to target (attack) a specific actor or country (e.g. working to discredit a political party), or be working persistently towards a long-term objective (polarising a country).

While assessing the objective of an influence operation is not always straightforward, it can be deduced based on the narratives they utilise and their target audience.

Several influence campaigns can be run by the same IMS, threat actor, or group of threat actors targeting, e.g. the same elections.

Influence Operation

An influence operation does not necessarily need to be part of a campaign.

Threat actors can run an influence operation as a singular effort. This is common and related to specific events or crises.


An influence operation not originating from an IMS is characterised by a threat actor using digital assets that are not anonymous disposable accounts that aims to manipulate the beliefs, behaviours, or decisions of a target audience using illegal content, illegal activity, inauthentic behaviour, covert activity, or foreign interference.

An example of a non-IMS influence operation includes state-owned media running a story spreading deepfake videos of a politician, or an influencer hiding that they are funded by a foreign ministry to travel to an illegally occupied region.

An operation can last several days or weeks, but usually doesn't last longer than that.

Delimited based on it having to be:

- Same threat actor or group of threat actors
- Same target or narrative.
- Same time period.





Taxonomy of New Information Manipulation Set (IMS)

No unified system for naming Information Manipulation Sets or CIB networks exists today.

Doppelganger, Overload, Storm-1516, CopyCopy, and Spamouflage are all different types of IMS that have previously been mapped by different parts of the Defender Community.

Storm-1516 is named through Microsoft's system (Microsoft, [N.A.](#)), which relies on a combination of words related to weather (e.g. Storm and Heatwave). These weather classifications relate to a specific nation-state, such as China (Typhoon) and Russia (Blizzard), or specific motives (e.g., financially motivated - Tempest). They use Storm and a four-digit number for groups that are still emerging where the operation is newly discovered, unknown, emerging, or developing - and yet not properly attributed.

Other naming conventions have been more ad hoc.

While the Microsoft system provides us with one component of attribution, we aim to shape this proposal to be easy to use and understand, and easily implemented.

A major challenge to this exercise is that standardised names are rarely very attractive and catchy, in the same way that, e.g. Doppelganger and Overload are.

The proposed structure is this:

<Threat actor Ecosystem> <IMS Type> <IMS Type 2>

IMS Types can be combined to showcase the different types of IMS operations.

This initial taxonomy is very much meant to start a conversation - it is far from perfect or a finished proposal.



Threat Actor Ecosystem

Threat Actor Ecosystem taxonomy aims to help name what ecosystem the IMS is originating from.

Microsoft's approach is per se not a bad one, but it is not very recognisable.

CrowdStrike proposes a Taxonomy that works around naming countries after animals, which creates engaging but not very clear names (CrowdStrike, N.A).

We don't want to complicate this beyond simply naming the threat actor ecosystems, e.g. Russia, China, and Iran.

IMS Type

IMS come in different shapes. The taxonomy below aims to help us name these.

Each IMS type has a short name that is combined with the threat actor name, e.g. Russian combined with e.g. Mimic.

These are later combined with a unique identifier that fits the characteristics of the operation.

Examples of existing operations:

Russian Web Mimic Bot Content Network - Doppelganger.

Russian Flood Social Network, Targeted Messaging - Operation Overload.

Bot Content Networks - Flood.

- No re-routing.
- Publish large amounts of content

Bot Re-routing Network - Re-router.

- Amplify links that lead to other sites.

Bot Interaction Network - Amplifier

- Amplify posts by sharing and/or liking them

Bot Comment Network - Commenter

- Amplify narratives by flooding replies and comment sections under targeted posts, often those of journalists, politicians, or media.

Bot Viewing Network - Viewer

- Amplify posts by viewing them.

Bot Following Network - Follower

- Amplify or legitimise accounts by following them.

Impersonation Websites - Web Mimic

- Impersonate legitimate entities to use their legitimacy to spread the operations narrative or discredit the impersonated actor.

Laundering Websites - Laundromat

- Republishing content from threat actor sources while hiding its origin.

Covert Websites - Shadow Web

- Publishing “news” and other content while hiding their funding or ownership by a threat actor.

Astroturfing Network - Astroturf Web

- Manually operated inauthentic accounts with developed identities (fake journalists, experts, local residents), used to give content credibility. Distinct from bots because of their depth and human operation.

Hijacked Account Network - Hijack Socials

- Compromised or purchased accounts, often verified or with an established audience, repurposed to spread operation content.

Targeted Outreach Network - Targeted Messaging

- Emails and direct messages sent to journalists, fact-checkers, and researchers to waste resources or bait coverage (the Operation Overload model).

Channel Network - Messaging channels

- Coordinated messaging-app or video channels (Telegram, YouTube, TikTok) that cross-post and amplify each other.

Influencer Network - Influencer

- Real individuals paid or directed to spread operation narratives, wittingly or not.

Front Organisations - Front Orgs

- NGOs, institutes, think tanks, or election observation missions that hide their links to a threat actor.


Advertising Network - Ad network

- Paid ads or sponsored content used to push operation content past organic reach limits.



AI Content Farms - AI Web

- Websites that mass-produce AI-generated articles with little human input, often mixing operation narratives with filler content. You could also treat these as a subtype of Covert or Laundering Websites.





Sources

CrowdStrike (N.A.), “Global Threat Landscape”. Retrieved 15/10/2026 from <https://www.crowdstrike.com/en-us/adversaries/>

EEAS. (2023). “1st EEAS Report on Foreign Information Manipulation and Interference Threats”. Retrieved 09/09/2026 from https://www.eeas.europa.eu/eeas/1st-eeas-report-foreign-information-manipulation-and-interference-threats_en

EEAS (2025). “3rd EEAS Report of Foreign Information Manipulation and Interference Threats”. Retrieved 15/10/2026 from <https://www.eeas.europa.eu/sites/default/files/documents/2025/EEAS-3nd-ThreatReport-March-2025-05-Digital-HD.pdf>

Innes, M. & Ahonen, A. (2025). “Attribution and Information Influence Operations: A ‘Field Guide’ for Open-Source Investigators and Researchers”. Retrieved 08/09/2026 from <https://adacio.eu/attribution-and-information-influence-operations-a-field-guide-for-open-source-investigators-and-researchers>.

Microsoft (N.A.). “How Microsoft names threat actors”. Retrived 15/10/2026 from <https://learn.microsoft.com/en-us/defender-xdr/microsoft-threat-actor-naming>

Nazari, S. & Sarvela, L. (2026). “DISRUPT White Paper: A Whole-of-Society Approach to Defending Democracy Against Influence Operations”. Retrieved 08/09/2026 from <https://alliance4europe.eu/report/disruption-framework-white-paper>

Palmertz, B., Isaksson, E., & Pamment, J. (2025). “A Framework for Attribution of Information Influence Operations”. Retrieved 09/09/2026 from <https://adacio.eu/a-framework-for-attribution-of-information-influence-operations>

Viginum (2026). “Information Manipulation Set (IMS)”. Retrieved 10/10/2026 from https://www.sgdsn.gouv.fr/files/files/Publications/20260122_NP_TLP-CLEAR_SGDSN_VIGINUM_IMS_0.pdf