CERT-EU

Table of contents

Introduction

The CERT-EU Cyber Threat Intelligence framework defines the analytical and operational standards CERT-EU uses to classify, assess, and prioritise malicious cyber activities relevant to our constituents, the European Union institutions, bodies, offices, and agencies (Union entities) and their ecosystem. The framework provides a shared reference model for us and our constituents to support consistent reporting, alerting and awareness raising on cyber threat intelligence.

The framework is also a key enabler for our Full-Spectrum Adversary Approach, our own flavour of threat-informed defence, as it supports consistent and holistic modelling of threats to Union entities across both strategic and technical dimensions. By facilitating the translation of threat observations and analyses into structured data, it strengthens situational awareness and operational coherence. This, in turn, enables faster reaction, clearer communication, and more effective response.

The framework introduces core concepts such as malicious activities of interest, ecosystem, threat categories, domains, and threat (actor) levels. It also outlines scoring mechanisms for adversaries and mitigation. These elements are designed to facilitate handling of cyber threats at various levels in Union entities, including by our primary operational contacts (POCs) and local cybersecurity officers (LCOs).

All components of this framework are aligned with recognised intelligence and cybersecurity standards and internal good practices of CERT-EU. Where applicable, terminology and methods follow practices from EU cybersecurity regulations, FIRST, NATO and threat intelligence industry good practices. The framework may evolve in response to regulatory changes, stakeholder feedback and other factors.

Malicious activities of interest

We define a malicious activity of interest (MAI) as any adversarial cyber activity with a potential impact for Union entities or their ecosystem. This includes confirmed compromise, suspicious

attempts, adversarial resource development, or reconnaissance activities. We are tracking MAIs to support alerting and awareness raising, and to support, where relevant, the response and mitigation of threats by Union entities.

Ecosystem

Limiting our monitoring to malicious cyber activity within our constituents' networks would risk missing relevant threats. At the other extreme, attempting to analyse all malicious activity across cyberspace would be prohibitively resource-intensive and beyond our capabilities. Consequently, we devised the concept of ecosystem to identify malicious activity that may affect our constituents directly or indirectly. Based on past observations, we defined the ecosystem as a set of components that reflect the exposure of our constituents to supply-chain risks, geopolitical developments, regional threats, and risks related to their business activity.

In our data-centric approach, we translated this concept into the following components: countries of operation, sectors of activity, geopolitical events of interest, partners, providers, systems and software, as defined in the table below.

Ecosystem component	Definition and examples
Countries	Countries in which Union entities operate. This includes all EU Member States as well as non-EU countries where Union entities have a physical presence. Each Union entity is located in one or more countries. Targeting these countries can affect constituents through local infrastructure or service breaches, and any campaigns with a geographical focus.
Sectors	Sectors in which Union entities operate. They are listed in the Sectors of interest chapter. A Union entity may belong to one or more sectors. Targeting a sector can expose constituents through shared dependencies and attack surfaces.
Events	Events of a geopolitical nature in which Union entities are involved and which may trigger or be targeted by malicious cyber activity. Examples include conferences, summits, disputes, international negotiations, conflicts or elections. The nature and level of a Union entity's involvement can vary. For instance, an entity may organise or participate in a conference or summit, or it may support or sanction a party to a conflict. As a result, event-related malicious cyber activity may target constituents directly or indirectly.
Partners	Organisations with which Union entities cooperate or exchange information. Each Union entity can have several partners, in EU countries or third countries. These partners can be permanent stakeholders of Union entities or may cooperate on ad hoc initiatives or projects. Examples include other Union entities, ministries or agencies in EU Member States, international organisations (for example NATO or the ICC), or non-profit organisations. Targeting partners can affect constituents through trusted channels, shared projects or information exchange.
Providers	Information technology (IT) companies providing services to Union entities. These include but are not limited to cloud service providers (CSPs), managed service providers (MSPs) and internet service providers (ISPs). The breach of a provider can affect constituents through service disruption, breach of data confidentiality, or malicious access to systems.
Software	Software products used by Union entities. These include but are not limited to operating systems, browsers, edge devices, security software, business software and AI software. Software products may be internet-facing or not. Targeting software used by constituents can affect them in various ways such as initial access via vulnerability exploitation, infection via trojanised software, and exfiltration or phishing via legitimate software.

Ecosystem component	Definition and examples
Systems	Information systems composed of technologies and software assembled by an organisation, or by a group of organisations, to support collaborative or shared purposes and for their exclusive use. Examples include Union entities' public websites and special-purpose services such as EU Login and EU Survey. Targeting shared or critical systems can directly affect service continuity, data integrity and user trust.

The classification of an event as a MAI is based on a combination of these factors. A single criterion may be sufficient where the impact is direct and significant; in other cases, several weaker indicators may collectively justify attention.

Threat and counter-threat categories

This section defines the core threat and counter-threat categories used to classify MAIs based on the intent of the threat actor or the nature of the action. Most categories describe adversarial intent; one category (Policy & law enforcement) captures non-adversarial context. Note that certain activities as well as threat actors may overlap across multiple categories, in some cases to hinder attribution.

Category	Definition
Policy & law enforcement	**(Non-adversarial context.)** Undertakings that aim to address malicious cyber activity. These include policies, regulations, cooperation, arrests, seizures, takedowns, bans, etc.
Cyberespionage & prepositioning	Threat actors steal sensitive information for intelligence purposes or covertly compromise an information system for future exploitation.
Cybercrime	Threat actors compromise systems for financial benefits. This includes ransomware breaches, compromising an IT system to sell access or deploying malware to steal credentials and resell them.
Hacktivism	Threat actors target systems to promote an ideological or political agenda. This includes certain website attacks such as DDoS, defacement, or hack-and-leak operations when they are carried out to draw attention to a political or ideological cause.
Opportunistic	Non-targeted malicious activity aiming at identifying and exploiting vulnerable systems in the wild. This includes spreading a worm through unpatched routers worldwide, or scanning and attempting automated exploitation of vulnerabilities in publicly exposed assets.
Digital foreign interference	The goal of the threat actor is to influence public opinion or sow discord via unauthorised cyber means. This includes fake accounts spreading disinformation during an election, leaking selectively altered documents to mislead the public, or bots amplifying polarising content on social media.
Disruption & destruction	The goal of the threat actor is to disrupt the operations of a victim's information system, destroy the system or destroy data. This includes wiper malware attacks, or DDoS on critical infrastructure.
Data exposure and leaks	The activity leads to information exposure or leaks, thereby causing damage to reputation, or facilitating further cyberattacks. This includes hack-and-leak operations by threat actors, or purposeful exposure or leaks from insider threats. Data exposure and leaks can also happen accidentally.

Category	Definition
Unknown	The purpose of the activity is unknown.

Threat domains

This section defines a hierarchical model for classifying the geographical or institutional scope affected by malicious cyber activity. Domains are ranked from the innermost institutional core to the broadest global context, as listed in the table below (from highest to lowest priority). When multiple domains apply, the highest-ranking domain takes precedence.

Domain	Definition
Union entities	The activity targeted one or more organisations as identified in Regulation 2023/2841.
EU	The activity targeted entities in one or more EU Member States, including national governments, infrastructure, or private entities.
Europe	The activity targeted entities in one or more European countries outside the EU. This includes some NATO countries, EFTA members, EU candidate and potential candidate countries.
EU Civilian Mission Area	The activity targeted one or more countries outside of Europe hosting an EU civilian mission.
World	The activity targeted any country not falling under the above domains.

Threat levels

This section defines the threat level scale used to assess the criticality and proximity of malicious cyber activity in relation to Union entities. These levels reflect analytical judgement based on threat actor intent, technical impact, and known targeting of Union entities. Threat levels are used particularly in the Threat Alerts we provide to Union entities. The scale below guides the urgency and prioritisation of mitigation and response.

Threat level	Definition
High	An immediate threat to Union entities. Verification and action are required without delay. Examples: • Significant Incidents affecting Union entities. • Exploitation in the wild of a zero-day in an internet-facing system deployed by multiple Union entities. • State-sponsored spearphishing campaign detected in at least one Union entity or in close partners.

Threat level	Definition
Medium	A close threat to Union entities. Close monitoring and checking are strongly recommended. Examples: • Focused cyberespionage campaign against sectors of interest (see Chapter - Sectors of interest) in the EU. • Opportunistic exploitation of a known vulnerability in software used by Union entities. • Threat actor activity targeting critical infrastructure within the EU.
Low	A distant or indirect threat with no immediately identified link to Union entities. Monitoring is advised, and action is recommended depending on available resources and priorities. Examples: • Opportunistic scanning or enumeration activity. • Global cyberespionage campaign targeting multiple continents with no apparent EU focus. • Indicators related to a non-EU incident reused in opportunistic malware campaigns.

Threat actor levels

This section defines the threat actor levels used to assess and prioritise adversaries based on their recent impact on Union entities and their ecosystem. The classification considers both the **period of interest** (e.g. last three months, last 12 months, or a defined timeframe such as 2025-Q1) and the **scope** (e.g. a specific constituent or the broader EU constituency).

Threat actor level	Definition
Critical	The threat actor caused at least one **Significant Incident** affecting one or more Union entities during the period of interest.
High	The threat actor is responsible for at least one MAI that has not been qualified as a Significant Incident, affecting **one or more Union entities** during the period of interest.
Medium	The threat actor is responsible for at least one MAI affecting **two or more elements of the ecosystem** during the period of interest.
Low	The threat actor is responsible for at least one MAI affecting **exactly one element of the ecosystem** during the period of interest.

Tactics, Techniques and Procedures (TTPs)

We use the MITRE ATT&CK framework to map techniques to the malicious activities of interest. This framework provides a shared, behaviour-based taxonomy that links observable actions to known adversary methods — making detection, threat-hunting and prioritised mitigation systematic and repeatable for CERT-EU and Union entities.

Sectors of interest

This section defines the sectors relevant to Union entities. Sectors are sorted in alphabetical order and not by importance. The list includes the sectors defined in the NIS2 Directive, plus additional sectors relevant to Union entities that are not covered by that directive.

The sector list supports structured analysis and classification of malicious activity. New sectors may be added as EU operational, regulatory, or policy priorities evolve.

Sector
Agriculture
Air transport
Chemicals
Cybersecurity
Defence
Diplomacy
Education
Energy
Environment
Finance
Fisheries
Food
Fundamental rights
Health
Intellectual property
Justice
Labour
Law enforcement
Maritime transport
Parliamentary administration
Pharmaceuticals
Public administration
Rail transport
Research
Space

Sector
Technology
Telecommunications
Transport

Confidence and uncertainties

Adhering to common norms for expressing confidence and uncertainties in CTI reporting ensures consistent interpretation, reduces miscommunication, and enhances the credibility and usability of our CTI products for Union entities. This section explains how we assess and express confidence in the information we use in our reporting and how we express uncertainties.

Confidence in information

We use the Admiralty Code, a NATO-standard system that rates the reliability of the source and the credibility of the information independently.

The Admiralty Code is based on two dimensions:

- **Source reliability**: An assessment of the trustworthiness of the source providing the information, based on their track record, access, and consistency. It is rated from A (completely reliable) to F (unreliable or untested).
- **Information credibility**: An assessment of the plausibility and confirmability of the information itself, regardless of the source. It is rated from 1 (confirmed by multiple sources) to 6 (cannot be judged).

The final confidence level is expressed as a combination of both dimensions (e.g. A1, B2).

We will use information in our threat intelligence products only if they match one of the **authorised combinations** marked as "Yes" in the table below. This threshold ensures that our CTI products are based on information from sources with a demonstrated track record (A or B) and with sufficient corroboration or plausibility (credibility 1 or 2). Combinations below this threshold are excluded to maintain the reliability and actionability of our reporting.

Credibility of information	A (Completely reliable)	B (Usually reliable)	C (Fairly reliable)	D (Not usually reliable)	E (Unreliable)	F (Reliability cannot be judged)
1 (Confirmed by other sources)	Yes	Yes	No	No	No	No
2 (Probably true)	Yes	Yes	No	No	No	No
3 (Possibly true)	No	No	No	No	No	No
4 (Doubtful)	No	No	No	No	No	No
5 (Improbable)	No	No	No	No	No	No

Credibility of information	A (Completely reliable)	B (Usually reliable)	C (Fairly reliable)	D (Not usually reliable)	E (Unreliable)	F (Reliability cannot be judged)
6 (Cannot be judged)	No	No	No	No	No	No

Communicating on uncertainties

We implement FIRST guidelines in our CTI reporting to address imperfect information and uncertainty by using standardised language — Levels of Confidence in Assessment (LCA) and Words of Estimative Probability (WEP). This ensures clarity, consistency, and usability for Union entities using our CTI products.

- **Levels of Confidence in Assessment (LCA)** express how confident we are in an analytical judgement (e.g. *low confidence*, *moderate confidence*, *high confidence*). They reflect the quality and quantity of supporting evidence and the strength of the analytical reasoning.
- **Words of Estimative Probability (WEP)** convey the likelihood of a future event or the accuracy of a current assessment using calibrated language (e.g. *unlikely*, *likely*, *very likely*, *almost certainly*). They help recipients interpret our assessments without overstating or understating probability.

Attribution

This section outlines the principles guiding our approach to attributing MAIs to threat actors. Attribution is the analytical process of linking observed activity to a threat actor, an intrusion set, a state, or an organisation. It is essential to clarify that we engage **only in technical attribution**, on an ad hoc basis only, and under strict conditions. We do **not engage in political attribution**.

- **Political attribution** refers to assigning accountability to a state or an organisation for malicious cyber operations — this falls outside our remit and is the responsibility of national or institutional decision-makers.
- **Technical attribution** involves linking malicious activity to known threat actors based on behavioural patterns, infrastructure reuse, malware indicators, and targeting profiles.

Technical attribution principles

- **Strictly technical:** We do not attribute activity to states or organisations. Our focus is on identifying threat actors based on technical indicators and behavioural consistency.
- **Where required:** We pursue technical attribution only where required to strengthen our Full-Spectrum Adversary Approach.
- **Evidence-based:** Attribution is grounded in observable characteristics, such as TTPs (tactics, techniques, procedures), infrastructure overlaps, malware artefacts, and targeting.
- **Confidence-driven:** We only attribute activity when supported by sufficient evidence and express a level of confidence. We reference open-source or partner analysis when deemed credible.
- **Contextual:** Attribution is valid for a defined period and scope, and may be updated as new information emerges.

Unattributed threat actors

When it is impossible to attribute a MAI to a known threat actor, particularly if it is qualified as Significant Incident, we link the MAI to an Unattributed Threat Actor (UTA) to which we append a numeric suffix (example: UTA-53). Depending on further analysis and information received, we might later merge a UTA with a known threat actor or with another UTA.

Scoring

This chapter explains how we calculate and apply scores to prioritise adversaries and defensive measures in our CTI products. These scores help determine which threats and mitigation are most relevant to the operational environment of Union entities. Information shared in this chapter covers the high-level principles used for threat scoring.

Threat scoring

The threat score is a numeric value used to measure the criticality of a threat, support prioritisation, and identify threat trends over time. It is based on MAIs linked to a given threat during a defined period. In simple terms, scores increase when MAIs are more frequent, more severe, closer to our constituency, and more recent.

The threat score model is based on five components:

1. Occurrences
2. Targeting
3. Severity
4. Time period
5. Decay over time.

Occurrences in scope

Occurrences in scope are the MAI-linked observations connecting a threat to one or more targets within the selected time window. They are the base input of the score: the higher the number of occurrences, the higher the score.

Targeting

Targets represent who or what is affected by the MAI. Direct targeting of our constituents weighs more than targeting of ecosystem elements. Broader targeting across multiple ecosystem components also increases the score.

Severity

Severity reflects the impact level of each occurrence. Higher-severity occurrences contribute more strongly than lower-severity ones.

Time period

The time period defines the observation window used for scoring. Shorter windows provide a more tactical view, while longer windows provide a more strategic view.

Decay

Decay reflects recency. Recent occurrences have stronger influence, while older occurrences progressively lose weight and eventually stop contributing.

Mitigation scoring

Mitigations are also scored to support prioritised defence planning. The score measures how well a mitigation addresses adversary techniques, protects initial access vectors, and aligns with recognised baseline practices.

We use the following formula:

```
Mitigation Score = K₁ × MMW + K₂ × MIA + K₃ × ME8
```

Where:

- `MMW` (Mitigation Weight): Total impact across observed adversary techniques and incidents.
- `MIA` (Mitigation Initial Access): Number of initial access techniques addressed.
- `ME8` (Mitigation Essential Eight): Number of linked Essential Eight controls.
- `K₁`, `K₂`, and `K₃` are weights periodically reviewed and adjusted based on new information and the evolving threat landscape.

These scores help determine which mitigations offer the greatest security value given observed threat activity. The `mitigations` file ranks defensive measures accordingly.

TLP definition

TLP	Disclosure	Message
RED	Not for disclosure, restricted to participants only.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed.
AMBER+STRICT	Limited disclosure, restricted to participants' organisations.	Recipients may share TLP:AMBER+STRICT information only with members of their own organisation.
AMBER	Limited disclosure, restricted to participants' organisations and their clients.	Recipients may share TLP:AMBER information only with members of their own organisation and its clients.

TLP	Disclosure	Message
GREEN	Limited disclosure, restricted to the community.	Subject to standard copyright rules, TLP:GREEN information may be distributed with peers and partner organisations within their sector or community, but not via publicly accessible channels.
CLEAR	Disclosure is not limited.	TLP:CLEAR information may be distributed freely.