A good practice guide of using taxonomies in incident prevention and detection

DECEMBER 2016

www.enisa.europa.eu

European Union Agency For Network And Information Security

About ENISA

The European Union Agency for Network and Information Security (ENISA) is a centre of network and information security expertise for the EU, its member states, the private sector and Europe's citizens. ENISA works with these groups to develop advice and recommendations on good practice in information security. It assists EU member states in implementing relevant EU legislation and works to improve the resilience of Europe's critical information infrastructure and networks. ENISA seeks to enhance existing expertise in EU member states by supporting the development of cross-border communities committed to improving network and information security throughout the EU. More information about ENISA and its work can be found at www.enisa.europa.eu.

Contact

For contacting the authors please use opsec@enisa.europa.eu
For media enquiries about this paper, please use press@enisa.europa.eu

Acknowledgements

The project team would like to express the gratitude to all the experts who contributed to the guide, who provided their time and insight into the guide.

Emmanuel Bouillon (NCIRC - NATO)
Europol
Frédéric Garnier (CERT-EU)
Javier Candau (CCN-CERT)
Lino Santos (CERT.PT)
Jorge Ruão (CERT.PT)
Koen Van Impe (Cudeso.be)
Paavo Kuiv (CERT-EE)
Pavel Kácha (CESNET)
Remi Seguy (Consilium)
Rita Forsi (Ministero dello Sviluppo, MISE)
Sandro Mari (Ministero dello Sviluppo, MISE)
Anna Passeggia (Ministero dello Sviluppo, MISE)

Legal notice

Notice must be taken that this publication represents the views and interpretations of the authors and editors, unless stated otherwise. This publication should not be construed to be a legal action of ENISA or the ENISA bodies unless adopted pursuant to the Regulation (EU) No 526/2013. This publication does not necessarily represent state-of the-art and ENISA may update it from time to time.

Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication.

This publication is intended for information purposes only. It must be accessible free of charge. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication.

Copyright Notice

ISBN: 978-92-9204-194-6, DOI: 10.2824/780536

Table of Contents

Executive summary

The aim of this document is to provide good practices on using taxonomies for incident detection and prevention by taking into account the input received from the CSIRT community and relevant information from previous ENISA studies. In addition, it provides conclusions and recommendations on improvements that can be made on current taxonomies.

The methodology followed to collect and assess the information for this study included a stock taking and desk research, discussions with CSIRTs during the 11th CSIRT ENISA Workshop, interviews with the CSIRT community, a qualitative assessment of taxonomies (and other formats and schemes relevant to mention) and a validation call with CSIRTs on 22 September 2016.

A qualitative assessment was performed on an indicative taxonomy landscape, which illustrates various ways of comparing and assessing taxonomies to determine, for instance, what fields that could be added or extended in current taxonomies.

In addition, use cases that would benefit from the use of taxonomies have been identified. These use cases include; recording events from different sources, automatic deduplication, ability to export in other taxonomies, ability to aggregate and search events in the data, ability to exchange data with other CSIRTs, feeding threat intelligence, incident report management. For each use case a requirement was identified that a taxonomy should fulfil. An indicative linking of the right taxonomy for the right use case is also provided both for threat intelligence and incident management purposes.

Since a comparison of every single taxonomy is not feasible, a couple of indicative examples of comparing taxonomies are provided. This includes an indicative comparison of fields of taxonomies similar to or based on the eCSIRT taxonomy [1]to determine which fields are common. In addition, a comparison of the complexity of taxonomies in terms of malware incidents is provided to illustrate different ways of describing the same context available in the current taxonomy landscape.

Furthermore, a number of good practices that have been identified by various CSIRTs during this study are described while taking into account shortcomings of taxonomies as identified by CSIRTs. The most relevant good practices include (a) that the top level categorisation of a taxonomy should be simple, (b) the categories within a taxonomy should be mutually exclusive, (c) taxonomies should support performance measurement and (d) that taxonomies should have an appropriate of level of ease of use. In addition, the report describes which good practices of high and medium relevance apply to which use cases according to the CSIRTs that have been involved in this study.

Three case studies illustrate the use of taxonomies in CSIRT operational activities while taking into account the use cases established in this study. These include a case study on using taxonomies for a website for major NIS incidents occurring across the EU for the public, using taxonomies to minimise the re-categorisation of cyber incidents and using taxonomies for incident handling metrics.

A number of conclusions were reached as a result of this study that should be taken into account in any future work to be performed by CSIRTs. The following are the most relevant to mention.

- There is currently no consensus on concepts and definitions related to taxonomies. Clear definitions reflecting the operational interpretation of the CSIRTs should be considered as a key success factor towards increasing cooperation between EU Member States. To this end, sections 8.1, 8.2 and 8.3 reflect some of the discussions held on concepts and definitions related to taxonomies with CSIRTs throughout this study.

[1] https://www.enisa.europa.eu/topics/csirt-cert-services/community-projects/existing-taxonomies

These sections have been added as food for thought and as possible point of departure for future work to be performed by CSIRTs.

- Taxonomies currently lack terms to properly handle the impact of an incident, incidents with no malice intended, explicit fields for ransomware, whether the incident is confirmed, and the differentiation between intrusion attempts and intrusions.
- The identified areas for potential improvement of existing taxonomies are based on the complexity, contextual information, mutual exclusivity or ambiguity, performance measurement, impact, sensitivity, confidentiality, and purpose of taxonomies.

The following recommendations based on CSIRT input and the good practices are provided:

- A centralised repository for hosting all relevant taxonomies along with their versions should be set up by ENISA. This would be a great benefit to the CSIRTs community as it would not only allow the selection of appropriate taxonomies for specific use cases, but it may also provide a general overview of what taxonomies or variations thereof are used by CSIRTs, which may be particularly useful in keeping statistics.
- A small set of common taxonomies for specific use cases should be agreed upon by CSIRTs at the EU level. This would provide examples of taxonomies based on the requirements of the CSIRTs network, which can be either implemented or used to implement a modified version of the taxonomy, saving time and effort that would be spent into researching taxonomies.
- An “Other” or “Unknown”, “Tag” field should be used by the owners of taxonomies as an indicator to revise taxonomies, if there is an increase in that category with incidents or events of the same type. For example, in a case involving ransomware, it is relevant that it should be categorised ransomware, but also the type of ransomware (such as crypto locker, etc.), if the same tag repeatedly used then it might also indicate the need for a new field.
- A roadmap towards standardised exchange formats in the CSIRTs community should be established at the EU level by the CSIRTs network. Such a roadmap should at least consider having CSIRTs agree use cases, definitions and concepts from an operational point of view for each use case; perform quantitative assessment (in addition to the qualitative assessment in this study) on the taxonomies used, a centralised repository for taxonomies, and a list of tags/values that can apply across taxonomies.

Finally, this document provides statistical data about taxonomy usage in a real Malware Information Sharing Platform (MISP) instance maintained by CIRCL. It allows the collection of statistical conclusions such as the relevance of the terms used by CSIRTs during day-to-day incident handling. While these statistics are composed of data from only one community, meaning they may not be representative of all CSIRTs, they illustrate the applicability of some of the conclusions made in this study.

1. Introduction

1.1 Purpose

The main objective of this report is to provide relevant good practices in terms of taxonomies for incident detection and prevention for the CSIRT community. Additionally, it aims to provide conclusions and recommendations based on the qualitative assessment of taxonomies within the current taxonomy landscape on improvements that can be made on current taxonomies, such as what fields can be extended or added to existing taxonomies.

1.2 Key Concepts and Definitions

It is important to highlight that the following terms have been used in different ways in various studies: ”information exchange standard”, “ontology”, “taxonomy”, “data type”, “data/field format standard”, “data/Field representation format”, “classification”, “semantic vocabulary”, “field” and “knowledge map”. In addition, based on feedback received by various CSIRTs during this study, it seems that there is no clear consensus on the exact interpretation of each one of these terms in the operational environment. Therefore, although the title and main focus of this report is on “taxonomies” used in incident detection and prevention, the reader is invited to consider the term “taxonomy” and other directly and indirectly related terms as something that is to be further discussed and agreed by the CSIRT community. However, chapter 8.1 does provide the reader with some preliminary insights on the above mentioned concepts gathered during this study that could serve as a possible basis for further work to be conducted in line with recommendation 1 “A centralised repository for hosting all relevant taxonomies along with their versions should be set up”. It is worth mentioning that the definition of “taxonomy” in previous ENISA studies is as follows:

A taxonomy[2] is defined as a classification of terms. Three characteristics define a taxonomy:

- a form of **classification scheme** to group related things together and to define the relationship these things have to each other;
- a **semantic vocabulary** to describe knowledge and information assets; and
- a **knowledge map** to give users an immediately grasp of the overall structure of the knowledge domain covered by the taxonomy, which should be comprehensive, predictable and easy to navigate.

In addition to this definition, according to some CSIRTs a taxonomy exposes inheritance and differentiation:

- **Inheritance**: Different kinds of objects often have a certain amount in common with each other;
- **Differentiation**: Characteristics of an object, which allow to differentiate an object from another object. For example, in case of malware you can have the category “malware” with the related characteristics “downloaders”, “rootkits”.

In addition, according to some CSIRTs, a taxonomy can be considered to be 2- or 3-dimensional. For example, TLP is an example of a 2-dimensional taxonomy while the classification of malware is usually 3-dimensional as it tends to have a “namespace” a “predicate” and a “value” (for example, SANS Malware Classification: malware type= value).

[2] https://www.enisa.europa.eu/activities/Resilience-and-CIIP/Incidents-reporting/metrics/ontology/ontology_taxonomies

2. Methodology

This chapter outlines the methodology applied to collect and assess the information for this study, including:

- Stocktaking and desk research;
- Discussions with CSIRTs during the 11th CSIRT ENISA Workshop;
- Interviews with the CSIRT community;
- A qualitative assessment of taxonomies;
- A validation call with CSIRTs on 22 September 2016.

2.1 Stocktaking and desktop research

The objective of the stocktaking and desktop research was to identify the relevant taxonomies and relevant studies for this study. In addition, the findings of the following ENISA reports were taken into account: "Actionable information for Security Incident Response"[3]; "Standards and tools used in exchange & processing of actionable information"[4]; "Detect, Share and Protect"[5]; "Information sharing and common taxonomies between CSIRTs and Law Enforcement"[6].

2.2 Discussions with CSIRTs during the 11th CSIRT ENISA Workshop

The 11th CSIRT ENISA Workshop took place on 10 & 11 May 2016 in Riga, Latvia. On 11 May 2016, the study team organised a half-day workshop to discuss the following topics with the CSIRTs for the benefit of this study:

- Current taxonomies used for incident detection and prevention within CSIRTs;
- Use cases describing the context in which the taxonomies are used;
- Current gaps vs. good practices related to the current taxonomies.

2.3 Interviews with the CSIRT community

Interviews were conducted with 15 CSIRTs to gain operational information related to the use of taxonomies. The key notions that were investigated are the following:

- Type of operational activities performed by CSIRTs that could benefit from the use of taxonomies;
- For what purposes CSIRTs use taxonomies; which of them are easy to use, what are the associated benefits, etc;
- Good practices as regards taxonomy use, creation, or revision.

2.4 Qualitative assessment and comparison of the collected taxonomies

A qualitative assessment was performed to evaluate the type and number of fields/terms, structure of taxonomies and to initiate the good practice guide on taxonomies used in incident detection and prevention.

2.5 A validation call with CSIRTs on 28 September 2016

A call was organised with various CSIRTs to validate the main outcomes of this study including its conclusions and recommendations.

[3] https://www.enisa.europa.eu/publications/actionable-information-for-security

[4] https://www.enisa.europa.eu/publications/standards-and-tools-for-exchange-and-processing-of-actionable-information

[5] https://www.enisa.europa.eu/publications/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs

[6] https://www.enisa.europa.eu/publications/information-sharing-and-common-taxonomies-between-csirts-and-law-enforcement

3. Qualitative assessment of taxonomies used by CSIRTs

This chapter presents the qualitative assessment of taxonomies (and other relevant formats and schemes) performed on an indicative taxonomy landscape. It illustrates various ways of comparing and assessing taxonomies to determine things such as common fields shared by taxonomies and what type of fields/terms that could be added or extended in current taxonomies.

3.1 Taxonomy inventory

This chapter contains a non-exhaustive list of the taxonomies that were analysed during the desk research. The following table is a brief explanation of each column of the table that will follow below:

TAXONOMY	PURPOSE
Admiralty Scale[7]	Describes the reliability of information
EUCI - EU classified information marking[8]	Describes any information or material designated by an EU security classification
FR classification[9]	French government information classification system
NATO Classification Marking[10]	Classification marking for classifying documents and information
OSINT Open Source Intelligence - Classification[11]	Open Source Intelligence - Classification (MISP taxonomies) which categorizes reliability of information
TLP - Traffic Light Protocol[12]	Used as categorisation scheme for sharing sensitive information while keeping control of its distribution at the same time
AVOIDIT Taxonomy[13]	Cyber-attack taxonomy based on a symposium paper by the University of Memphis, USA
Adversary[14]	Describes adversary infrastructure status, type, as well as action
The common language[15]	Developed at the Sandia National Laboratories, uses 3 main terms: event, attack, and incident
US-CERT[16]	US Federal incident notification guidelines. Contains terms to describe Impact (functional, information impact, and recoverability) and threat vector
Veris[17]	Vocabulary for Event Recording and Incident Sharing
dhs-ciip-sectors[18]	DHS critical sectors as in https://www.dhs.gov/critical-infrastructure-sectors
EU Marketop and Public Admin[19]	Market operators and public administrations that must comply to some notifications requirements under NIS directive
FIRST CSIRT Case Classification[20]	Guidelines needed for CSIRT Incident Managers (IM) to classify the case category, criticality level, and sensitivity level for each CSIRT case. This

[7] https://www.circl.lu/doc/misp-taxonomies/#_admiralty_scale
[8] http://www.consilium.europa.eu/en/general-secretariat/corporate-policies/classified-information/
[9] https://www.circl.lu/doc/misp-taxonomies/#_fr_classif
[10] http://www.dtic.mil/whs/directives/corres/pdf/520001_vol2.pdf
[11] https://fas.org/irp/doddir/army/atp2-22-9.pdf
[12] https://www.circl.lu/pub/traffic-light-protocol/
[13] http://www.albany.edu/iasymposium/proceedings/2014/6-SimmonsEtAl.pdf
[14] https://www.circl.lu/doc/misp-taxonomies/#_adversary
[15] https://www.enisa.europa.eu/topics/csirt-cert-services/community-projects/existing-taxonomies
[16] https://www.us-cert.gov/incident-notification-guidelines
[17] http://veriscommunity.net/
[18] https://www.circl.lu/doc/misp-taxonomies/#_dhs_ciip_sectors
[19] https://www.circl.lu/doc/misp-taxonomies/#_eu_marketop_and_publicadmin
[20] https://www.first.org/_assets/resources/guides/csirt_case_classification.html

	information will be entered into the Incident Tracking System (ITS) when a case is created
Hungarian Taxonomy[21]	Taxonomy describing who reported an incident. I.e. national CIIP, CIIP of partners with SLA, incidents reported by international partners, threats and incidents reported by cooperating organisations
Malware Categorisation based on a SANS document[22]	Categorisation based on different categories to describe malware
SURFcert taxonomy - KISS[23]	Uses the following terms only: (Administrative), Content, Vulnerable, Spam, Abusive, Probe, Denial

3.2 Linking taxonomies to use cases based on defined requirements and metrics

This section considers relevant use cases that would benefit from the use of taxonomies. For each use case a requirement was identified that a taxonomy should fulfil. This linking takes into account the various attributes of a taxonomy (such as complexity, size, etc. defined below), and insures that the taxonomies selected can adequately meet the requirements of the use cases.

3.2.1 Use cases and use case requirements

A use case is a list of actions or event steps, typically defining the interactions between a role and a system, to achieve a goal. In this case, "the system" would consists of the taxonomy and/or an exchange platform. The "goal" could for example relate to either incident detection or incident prevention.

In use case UC.TI.101 - Recording events from different sources, the goal is to adequately record events with the applicable system using a taxonomy, where recording the events is the action.

From the use cases, requirements are created in terms of what a taxonomy should provide to fulfil the task or activity at hand. This is performed because setting requirements in line with the overall goal of the use cases aids in the identification of which taxonomy that is the most appropriate. The requirements depend on the nature of the data in question. The requirements in day-to-day operational activities may differ entirely from those set out in the table below, they are not meant to represent a concrete example, but merely a suggestion of the type of requirement and qualities of taxonomies that would be suitable for a given situation. The use cases in the table below are based on interviews with the CSIRT community and research[24] conducted on use cases related to incident detection and prevention specifically for CSIRTs. The use cases employ the following naming convention:

- The prefix UC refers to Use Case;
- TI and IM after the prefix refer to Threat Intelligence and Incident Management respectively;
- 0XX and 1XX after represent whether the use case is based on interviews with the CSIRT community (marked with a 1), or on research (marked with a 0). XX represents the unique identifier of the use case.

The title of the use cases is also representative of the defined goal of the use case. This is followed by a brief description of the use cases in terms of CSIRT activity and the corresponding requirement that needs to be fulfilled by the taxonomy:

[21] https://www.enisa.europa.eu/topics/csirt-cert-services/community-projects/existing-taxonomies

[22] https://www.sans.org/reading-room/whitepapers/incident/malware-101-viruses-32848
[23] https://www.terena.org/activities/tf-csirt/meeting39/20130523-DV1.pdf
[24] http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=14253

3.2.1.1 Linking use cases and use case requirement for threat intelligence

USE CASE	DESCRIPTION	REQUIREMENT
UC.TI.101 Recording events from different sources	A CSIRT member should be able to describe the incident, vulnerability, or threat information from different sources, while taking into account that different sources (such as honeypots, sinkholes, etc.) require different terminology to adequately describe the event or source.	**More terms and complexity**. To adequately account for the varying terminology that different sources of information may present, a more advanced taxonomy in terms of semantic vocabulary and number of terms may be preferable .
UC.TI.002 Automatic deduplication	Whenever new event, vulnerability, or threat information is received by the system after a sharing process by another CSIRT, the system should be able to recognize whether the received information has already been ingested or not.	**Maximise mutual exclusivity**. The deduplication process (either by general or specific automatic “server-side” deduplication of even information) may be affected by the number of terms and the mutual exclusivity of them. If there are many terms in a taxonomy, but they are not overly mutually exclusive, then an identical incident may be recorded under a different category, causing a duplicate to remain. This of course depends on the duplication process itself as well.
UC.TI.003 Ability to export in other taxonomies	A CSIRT team member should be able to require event, vulnerability, or threat information to be exported in structured and open formats provided by the platform (e.g. OpenIoC, STIX or TAXII, and CVE or CVRF, if these specifications are supported).	**High complexity, high number of terms, but simple top level categories preferable**. Exporting in other taxonomies is affected by the complexity and number of terms that are being exported from and exported to. Exporting in a more complex taxonomy requires more effort than exporting in a simple taxonomy, but causes less loss of information and context overall.
UC.TI.104 Ability to aggregate and search events in the data	A CSIRT team member should be able to aggregate events using appropriate terms to classify the event elements, and should also be able to search those events effectively.	**Multiple levels of categories with simple top-level categories.** For aggregation and search of events in data. It should be simpler to use a taxonomy with multiple levels of categorisation, as it would speed up the aggregation and searching of the event. Furthermore, top-level categories should be simpler and not numerous, to facilitate the aggregation and searching process.
UC.TI.105 Ability to exchange data with other CSIRTs	A CSIRT team member should, when required, be able to effectively exchange information with another CSIRT conveying information accurately and with no loss of information. Potentially being able to map taxonomies to enable the information transfer.	**High number of terms and high complexity.** It would be preferable to maximize the information exchange; as such using a complex taxonomy with larger amount of terms allows more specific and contextual information to be exchanged.

3.2.1.2 Linking use cases and use case requirement for incident management

ID AND TITLE	DESCRIPTION	REQUIREMENT
UC.IM.006 Feeding threat intelligence	A CSIRT team member (i.e. CSIRT incident handler) indicates an incident as solved or is in the ongoing process of being solved. The system should extract information from the incident, structure that information in a threat format supported by the system (e.g. STIX) and feed the threat intelligence platform.	**High number of terms and high complexity.** To ensure that the full incident information is available upon requirement, a maximum amount of information should be provided. As such using a taxonomy with high complexity and number of terms may be appropriate.

ID AND TITLE	DESCRIPTION	REQUIREMENT
UC.IM.007 Incident report management	A CSIRT team member (i.e. CSIRT incident handler) should be able to access automatically generated incident reports and define new reports (to enable automatic generation) by specifying properties, such as, but not limited to, scope, type of incident, priority, time (daily, weekly, etc.); and visualization (graphs, trends, or maps).	**Low number of terms and low complexity.** The target audience of the report dictates the requirement of the taxonomy. If the report is to management or none technical personnel, then the taxonomy should reflect that by using simple and general terms, whereas if it is a technical report, the complexity of the taxonomy should reflect that.

Table 1 - Table of use case as defined by previous research and input from the CSIRT community.

3.2.2 Indicative linking of the right taxonomy for the right use case

To facilitate the linking of the taxonomies to the appropriate use case, the taxonomies are chosen from a reduced taxonomy pool in the table below. The criteria for the reduced pool is to maximise the variation in taxonomy type and format. The relevant taxonomies are listed and then categorised as either technical or general. In the table below, "technical" in the second column refers to whether the taxonomy utilises technical terms that may not be understood by all audiences. While terms such as "virus" or "worm" are generally understood, some others such as "XSS" or "CRFS" may not be. As some taxonomies contain both general and technical terms, if a taxonomy contains many technical terms it will be categorised as technical. The "size of the semantic vocabulary", in the third column, is defined in relation to the overall size using eCSIRT.net MkII as a baseline due to its popularity in the CSIRT community. "Overall complexity" is defined based on the "technical or general" and the number of terms of the taxonomy.

TAXONOMY	TECHNICAL OR GENERAL	SIZE OF SEMANTIC VOCABULARY	OVERALL COMPLEXITY
CERT.pt	General	Medium	Low
CIRCL	General	Small	Low
eCSIRT.net MkII	General	Medium	Medium
ENISA Taxonomy	General	Large	Medium
First CSIRT Case Classification	General	Small	Low
CCN-STIC 817	Technical	Medium	Medium
Veris	Technical	Large	High

Table 2 - List of reduced taxonomy pool to perform an indicative linking between use cases and taxonomies.

The table below takes the use cases previously defined in this section and takes into account the taxonomy attributes defined in the previous table. For each use case and goal, an appropriate taxonomy is selected.

3.2.2.1 Indicative linking of the right taxonomy for the right use case for Threat Intelligence

USE CASE	TAXONOMY	REASON
UC.TI.101 Recording events from different sources	Veris	Veris provides a comprehensive list of both technical and non-technical terms that are general enough to record a number of events. This enables effective recording from different sources.
	FIRST CSIRT Case Classification	FIRST CSIRT Case Classification Also is technical enough to gain an understanding of the transpired event, but it is not restrictive enough to cause loss of meaning. This allows the recording of events from different sources on a higher level effectively.

USE CASE	TAXONOMY	REASON
	ENISA Taxonomy	The ENISA taxonomy provides a wide range of different incident types. As such, it would be well suited to handle an incident from various sources. Aspects covered by the taxonomy include, but are not limited to eavesdropping, interception, hijacking, nefarious activity or abuse, etc. This allows the handling of event information regardless of whether it originated from a honeypot, sinkhole, etc.
UC.TI.002 Automatic deduplication	CERT.pt	CERT.pt is low in complexity and general in its terms. This does however mean that each concept represented in the terms is often mutually exclusive. As this taxonomy has two levels of categorisation, it further enables mutual exclusivity, which helps in the deduplication process by avoiding scenarios where the same incident is present more than once due to a different categorisation.
	CCN-STIC 817	CCN-STIC 817 uses a comprehensive list of incident types. The taxonomy also covers a relatively wide aspect including more technical terms and general terms (“XSS” versus “Human Error”).
UC.TI.003 Ability to export in other taxonomies	Veris	Veris contains a large number of terms that can be used to maximise the information transferred with minimal loss of context. Exporting into other taxonomies is more effective when a larger set of terms, as it avoids miss-categorisation.
UC.TI.104 Ability to aggregate and search events in the data	CCN-STIC 817	Due to CCN-STIC 817 having multiple levels of categorisation, the aggregation and search can be performed on the different levels. The top-level categories provided by this taxonomy also allow to search by action or target, making it particularly useful in searching of event information. For example, making a search on the result of an incident to map what incidents have similar outcomes and consequences.
	CIRCL	CIRCL provides a small number of general terms, as such it makes the process of aggregating and searching data simpler.
	eCSIRT.net MkII	eCSIRT.net MkII has multiple levels of categorisation, allowing for effective aggregation and searching of terms. I.e. events can be aggregated and search using the top-level categorisation, and the next level down of categorisation can be used to refine the aggregation and search.
UC.TI.105 Ability to exchange data with other CSIRTs	eCSIRT.net MkII	Depending on the main type of information exchanged, the eCSIRT.net MkII taxonomy provides a general semantic vocabulary that could facilitate the sharing of threat information between CSIRTs.
	Veris	Veris is the taxonomy used to collect the data volunteered by various CSIRTs for the yearly Verizon reports. As such, it is tried and tested to work with exchanging data effectively with other CSIRTs. It contains a great number of technical terms, effectively providing a comprehensive picture of an incident.

3.2.2.2 Indicative linking of the right taxonomy for the right use case for Incident Management

USE CASE	TAXONOMY	REASON
UC.IM.006 Feeding threat intelligence	Veris	Veris provides an extensive vocabulary, which provides a very comprehensive list of terms.
UC.IM.007 Incident report management	eCSIRT.net MkII	This taxonomy provides simple and comprehensible terms that can be used effectively for a non-technical environment. Additionally, it provides a multi-level categorisation, where it is possible to abstract to the top-level should the bottom level be too technical. It allows for use in a more technical environment, as the bottom level categorisation contains more detail than the top level. For example, if there was a phishing incident that is registered by a CSIRT, it could be described “Abusive Content -> SPAM”, “Fraud -> Phishing”.

USE CASE	TAXONOMY	REASON
	CCN-STIC 817	CCN-STIC 817 provides a comprehensive map of terms and a simple top-level category that can be used to provide different levels of technicality when reporting requires it.

Table 3 - Table linking the taxonomies to specific use cases based on established requirements and metrics

3.3 Example comparison of related taxonomies to evaluate common terms

This section provides two examples of comparing related taxonomies:

- Comparison of fields of taxonomies similar to or based on the eCSIRT taxonomy to determine which fields are common;
- Comparision of complexity of taxonomies in terms of malware incidents to illustrate different ways of describing the same context available in the current taxonomy landscape.

Note that comparing every single taxonomy is not feasible as some taxonomies are too different to make comparison possible (for instance, comparing TLP to Veris). In this respect, an example of a comparison between different versions of the eCSIRT.net taxonomy is also provided.

3.3.1 Example of comparison of fields of taxonomies similar to or based on the eCSIRT taxonomy

Common fields shared among taxonomies can be determined by comparing fields of popular taxonomies. This comparison can also illustrate any new fields that can be added to existing taxonomies, as well as extending current fields according to the needs of the CSIRT community in order to enhance taxonomies.

The following four taxonomies are similar in the sense that they are either used by eCSIRT or part of it as the basis for the taxonomy:

CERT.PT	CERT.BE	CESNET CERT	ECSIRT.NET MKII
Malware	Spam	Spam	Spam
Botnet Drone	Abusive Content	Bounce	Harassment
Ransomware	Malware	Virus	Child/Sexual/Violence/…
Malware Configuration	Scan	Malware	Virus
C&C	System/Account Compromised	Trojan	Trojan
DDoS	(D)DoS	Malware	Spyware
Scanner	Phishing	Probe	Dialler
Exploit	Vulnerability Report	Crack	Rootkit
Brute-force	Other	Botnet	Scanning
IDS alert		Dos	Sniffing
Defacement		Copyright	Social Engineering
Compromised		Scam	Exploiting of known Vulnerabilities
Backdoor		Phishing	Login attempts
Drop zone		Pharming	New attack signature
Phishing		Other	Privileged Account Compromise
SPAM		Unknown	Unprivileged Account Compromise
Vulnerability			Application Compromise
Service			Bot
Other			DoS

			DDoS
			Sabotage

It is worth noting that a term from the bottom level category of one taxonomy may be a top category in another taxonomy. This is the case for “Abusive Content” in CERT.be, which is present as a top-level category in the eCSIRT.net.

The Venn diagram below displays the terms present in all of the above taxonomies. Illustrating the fields shared amongst these taxonomies, each of the ovals contain the terms of the taxonomy, each intersection in the ovals illustrates that a term, an equivalent term, or similar term are present in multiple taxonomies (such as the term “Spam” being contained in all the taxonomies below). Where the wording of the term is different, but the context is similar or the same in different taxonomies, they are grouped in squares.

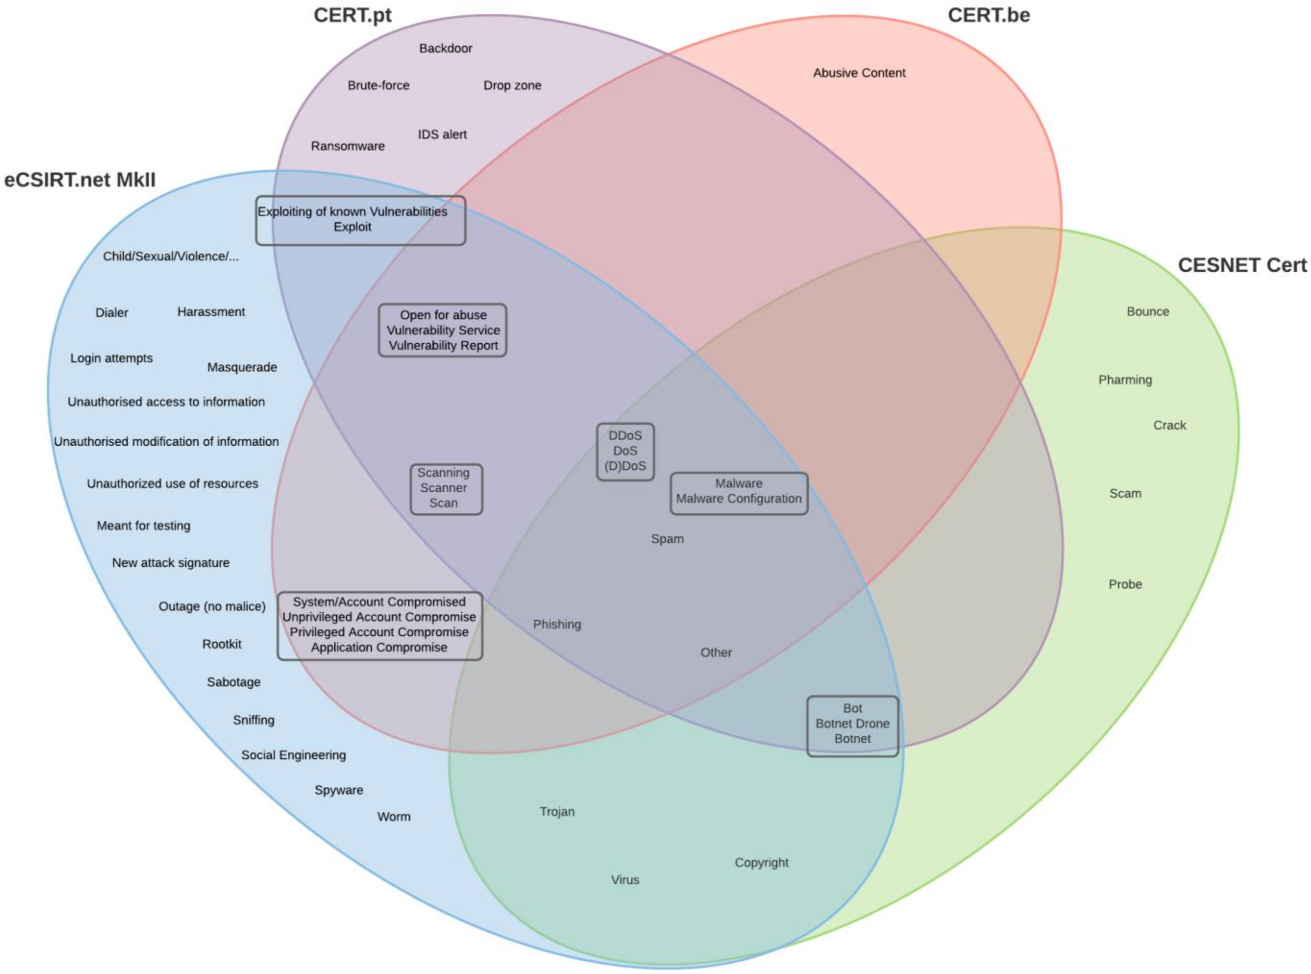

Figure 1 - Venn diagram comparing the terms of the eCSIRT.net MkII, CERT.pt, CERT.be, and CESNET Cert taxonomies. Similar concepts with different wording are grouped in squares.

Based on the Venn diagram above the following observations are made:

- The most common fields (present in all four taxonomies) shared by the taxonomies are "DDoS", "Spam", "Phishing", "Malware", terms for vulnerabilities, "Scan", and "Other";
- Some of the terms are worded differently or describe different aspects (such as action, process, or entity). For example, "Scan", "Scanner", and "Scanning" represent a similar concept but are worded differently and do not refer to the concept in the same manner;
- CERT.pt is the only taxonomy to mention "Ransomware", which has been a rapidly growing threat[25]. As the likelihood of ransomware is increasing, and the high impact it causes it could be considered as a term to add.

3.3.2 Example of comparing complexity of taxonomies in terms of malware

While many taxonomies describe malware, the level of detail of malware varies among them. To have a clear view of the difference in the level of detail, the table represents the fields of the taxonomies (for the Veris taxonomy, only the appropriate category is displayed due to its relatively high number of terms). This comparison allows conclusions to be drawn, as all of these taxonomies cover the concept of malware somehow, the way in which they are represented differ. The compared taxonomies are CIRCLE, CERT.pt, SANS, and Veris:

CIRCL	CERT.PT[26]		SANS		VERIS	
Spam	Malicious Code	Malware	Malware Categories	Virus	ACTION MALWARE VARIETY	Adware
System Compromise		Botnet Drone		Worm		Backdoor
Scan		Ransomware		Trojan		Brute Force
Denial of service		Malware Configuration		Ransomware		Capture App Data
Copyright Issues		C&C		Rootkit		Capture Stored Data
Phishing	Availability	DDoS		Downloader		Client-Side Attack
Malware	Information Gathering	Scanner		Adware		Click Fraud
XSS	Intrusion Attempts	Exploit		Spyware		C2
Vulnerability		Brute-force	Obfuscation categorisation	No-obfuscation		Destroy Data
Fastflux		IDS Alert		Encryption		Disable Controls
SQL Injection	Intrusion	Defacement		Oligomorphism		DoS
Information Leak		Compromised		Metamorphism		Downloader
Scam		Backdoor		Stealth		Exploit Vulnerability
	Information Content Security	Dropzone		Armouring		Export Data
	Fraud	Phishing		Tunnelling		Packet Sniffer
	Abusive Content	SPAM		Retro		Password Dumper

[25] http://www.symantec.com/connect/blogs/report-organizations-must-respond-increasing-threat-ransomware

[26] http://www.cncs.gov.pt/cert-pt-2/documents-2/index.html

CIRCL	CERT.PT[26]		SANS		VERIS	
	Vulnerable	Vulnerability Service	Payload Categorisation	No Payload		Ram Scraper
	Other	Other		Non-Destructive Payload		Ransomware
				Droppers		Rootkit
			Memory Categorisation	Resident		Scan Network
				Temporary Resident		Spam
				Swapping Mode		Spyware / Keylogger
				Non-Resident		SQL Injection
				Swapping Mode		Adminware
				User Process		Worm
				Kernel process		Unknown

Table 4 - Comparison of possible ways to represent malware with a taxonomy.

From the example above, a number of observations that can be made regarding the difference between taxonomy fields, especially those concerning the same or a similar concept. The following observations include:

CIRCL

- CIRCL has the most general semantic vocabulary. Malware incidents or events involving malware that lead to an incident are categorised as malware. This has some benefits in situations where the use of high level terms is appropriate (reporting, for example).

SANS

- Provides the most comprehensive information regarding malware, not only categorising into malware types (e.g. worm), as well as obfuscation techniques, payload, and memory categorisation in the taxonomy. In other words, the SANS malware categorisation taxonomy not only classifies malware, but also its evasion and operation mechanics. In terms of use cases, SANS is technically orientated, meaning it is well suited for technical reports or personnel, but less well suited for general use.
- It is possible SANS can give more insight into new detailed trends of attack methodologies, which might not be clear from the other taxonomies. This is because SANS describes the malware itself, whereas Veris describes the action of the malware.

Veris

- Veris has the most comprehensive list of terms (of the compared taxonomies). There is a distinction between malware and hacking, but also with each respective distinction being split into variety and vector. This taxonomy classifies the malware according to what the malware is or does, but also includes terms to classify propagation (I.e. the threat vector, not shown in the table).

3.3.2.1 Contextual differences between similar terms

There is a distinction in the context of the terms in the above example. While semantically the terms “Malware”, “Malicious Code”, “Malware Category”, and “Malware Variety” are similar, their respective sub-categories can be divided into two distinct concepts for this specific example:

- The action that caused the incident or event - such as “Destroy data” or “SQL Injection”;
- The element that made the action that cause the incident or event - “Virus” or “Worm”.

It is also important to consider that in both SANS and Veris the field named “Worm” exists, although the terms are identical, they are contextually different. This is because SANS describes the malware itself, whereas Veris describes the action of the malware. While both aspects are important to know, this may also lead to misunderstandings or different understandings of terms.

3.3.3 Example of a comparison between different versions of the eCSIRT.net taxonomy

This section compares the eCSIRT.net dating from 2003 with the eCSIRT.net MkII[27]. The following tables illustrate the additions of the eCSIRT taxonomy upon upgrading it. It contains the incident classes present (the top-level category), and the incident types from the 2003 version and the incident example from the 2013 version. The following are incident classes and their sub- categories, which have not had their terms updated or changed:

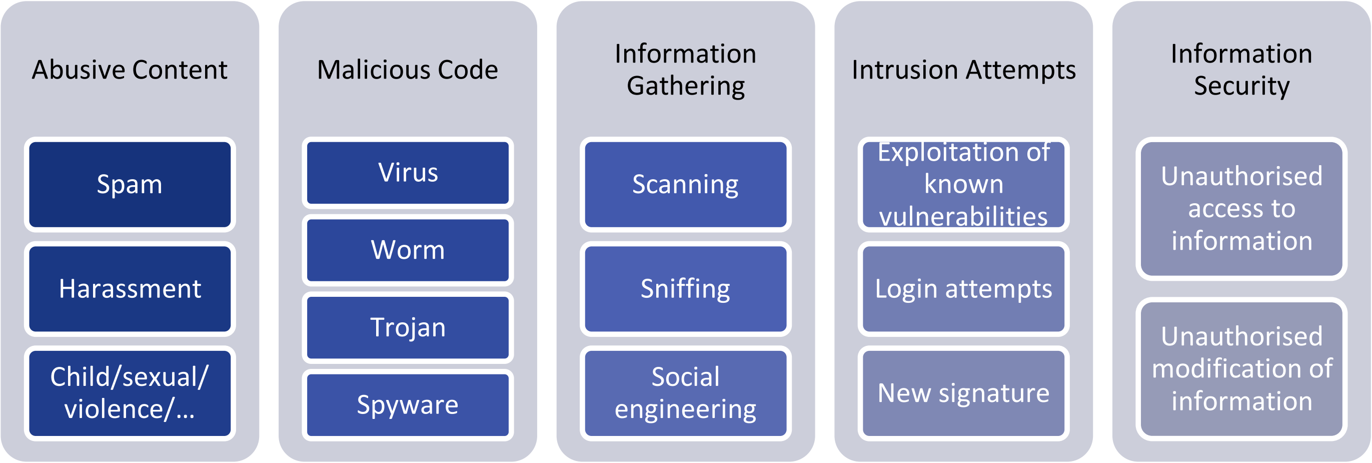

Abusive Content	Malicious Code	Information Gathering	Intrusion Attempts	Information Security
Spam	Virus	Scanning	Exploitation of known vulnerabilities	Unauthorised access to information
Harassment	Worm	Sniffing	Login attempts	Unauthorised modification of information
Child/sexual/ violence/…	Trojan	Social engineering	New signature	
	Spyware			

New terms or modifications are additionally highlighted in bold.

INCIDENT CLASS	INCIDENT TYPE (2003)	INCIDENT EXAMPLE (2013)
Intrusions	Privileged Account Compromise	Privileged Account Compromise
	Unprivileged Account Compromise	Unprivileged Account Compromise
	Application Compromise	Application Compromise
		Bot
Availability	DoS	DoS
	DDoS	DDoS
	Sabotage	Sabotage
		Outage (no malice)
Fraud	Unauthorized use of resources	Unauthorized use of resources
	Copyright	Copyright
		Masquerade

[27] https://www.terena.org/activities/tf-csirt/meeting39/20130523-DV1.pdf

INCIDENT CLASS	INCIDENT TYPE (2003)	INCIDENT EXAMPLE (2013)
		Phishing
Vulnerable		Open for abuse
Other		All incidents which don't fit in one of the given categories should be put into this class
Test		Meant for testing

Table 5 - Comparison of modified categories eCSIRT.net and eCSIRT.net Mk II with additions in bold.

The following observations were made from the above comparison:

- This is good example of a taxonomy which explicitly separates the concepts between an intrusion attempt and a successful one, which can provide useful information in incident analysis;
- Terms have been added to account for new threats (i.e. such as "Phishing");
- A category has been added to also keep track of vulnerabilities, this is useful as it is a concept CSIRTs often need to contend with;
- The "Other" category allows the categorisation of any incident or event that was not accounted for. The major advantage of this category is that if there is an influx of incidents of a similar (new) nature, it can be used as a point of consideration to revise the taxonomy and account for it. It is also often used for things that are not incidents or events, such as conference requests, job requests, press, etc.;
- The addition of the term of "Outage (no malice)", allows the taxonomy to cover a broader scope of incidents;
- The addition of the term "Bot" is also relevant as it has become a greater and more frequent threat. It may also be worth considering other more modern terms for all taxonomies to be up to date with the most current threat landscape.

4. Good practices for taxonomies used for incident detection and prevention

This chapter describes a number of good practices identified by various CSIRTs during this study. A good practice can be described as a practice that has been proven to work well and which produces good results, and therefore can be recommended as a model. In addition, it is a successful experience, which has been tested, validated and repeated and deserves to be shared with the CSIRT community to consider it for adoption. Each good practice has been defined while taking into account shortcomings of taxonomies as identified by CSIRTs. This chapter is organised as depicted below.

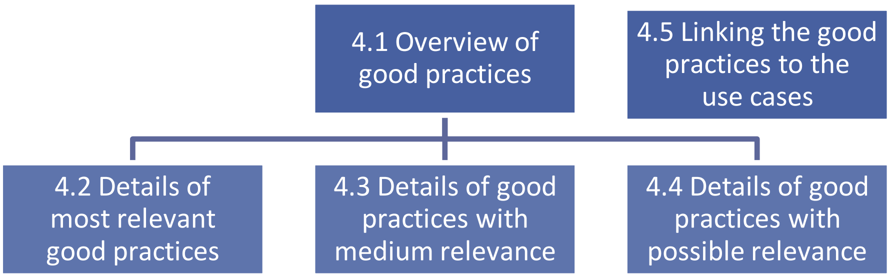

4.1 Overview of good practices

4.5 Linking the good practices to the use cases

4.2 Details of most relevant good practices

4.3 Details of good practices with medium relevance

4.4 Details of good practices with possible relevance

- Chapter “4.1 Overview of good practices” provides a high-level overview of the identified good practices grouped by relevance;
- Chapters 4.2, 4.3 and 4.4 provide details of the good practices with high relevance, medium relevance and possible relevance respectively. For each good practice details are provided on what problem or gap has been taken into consideration, a summary of the good practice, how the good practice addressed the issue at hand and a practice example;

Chapter “0

- Linking the good practices to the use cases" described which good practice is considered to support which use case as identified in chapter "3.2 Linking taxonomies to use cases based on defined requirements and metrics".

4.1 Overview of good practices

The good practices have been grouped by relevance to this guide based on the qualitative assessment. However, actual relevance depends on the reader and his/her context. Good practices established by the interviewed CSIRT experts that focus on the taxonomies themselves (and less on the CSIRT operational activity practices) are also assigned a higher priority. This chapter is organised as depicted below.

4.1.1 Overview of the most relevant good practices

4.1.2 Overview of the medium relevance good practices

4.1.3 Overview of other good practices with possible relevance

4.1.1 Overview of the most relevant good practices

The following table contains a quick overview of the good practices identified as "highly relevant" to the CSIRT community. Note that each good practice is described in more detail in chapter "04.2 Details of most relevant good practices".

GOOD PRACTICE	REASONING
GP1: The top-level categorisation of a taxonomy should be simple	With simple top-level categorisation we take into account the complexity of the taxonomy. As simple top level also implies more than one level of categorisation. With a multi-level categorisation system, the preferred level of complexity can easily be selected. If a non-technical report is required, a higher and more general categorisation level can be used. If a technical report is required, then using the bottom level category enables that. In addition, having a low number of top-level categories can reduce the mutual exclusivity of the categories. Furthermore, a taxonomy with at least 2-3 levels of categorisation provides the most versatility and scalability, as it gives the choice of adding a branch to a tree, or adding a leaf to the branch.
GP2: The categories within a taxonomy should be mutually exclusive	This good practice is related to an issue brought forth by a few CSIRTs. The multiple categorisation or re-categorisation of an incident during the incident lifetime. While sometimes this may be unavoidable, as laid out in the use case B of a spear phishing campaign, the first time categorisation of an incident should not cause any doubt. The re-categorisation may influence reporting, as it may be unclear whether to only take the final category in the report or take account of the previous categories. (E.g. reporting could include statistics such as "X cases of type A have eventually resulted in Y cases of type B after final categorisation). If the taxonomy terms are defined too loosely and if they do not have enough constraints, the same incident could potentially be categorised differently by two different analysts. It would also make the machine reading harder due to the added ambiguity.
GP3: Taxonomies should support performance measurement	There are many benefits to allowing a taxonomy to support the measurement of performance in terms of time taken to solve an incident (both in terms of complexity and person hours to solve it). In many cases it is for example relevant to know how quick an incident was closed. Some reports need to include the time needed to close the incident. This good practice was mentioned by several CSIRTs. A taxonomy with such a feature would allow management to make operational decisions, such as resource allocation for a type of incident whose estimate complexity and solving time is known. It is also useful for keeping statistics.
GP4: Taxonomies should have an	Ease of use needs to be clearly defined and depends on the operational requirement of the CSIRT, and the use case for which the taxonomy is required. Although ease of use can be interpreted in

GOOD PRACTICE	REASONING
appropriate level of ease of use	many different ways by different entities, the following attributes may provide relevant and general aspects of ease of use: • Completeness of the taxonomy; • Complexity of the taxonomy; • Simple first level of categorisation (for example, when triage happens at the emergency desk in a hospital, the initial intake should be simple, swift and self-explanatory); • If the taxonomy is to be used to report to decision makers, the terms have to be understood by the decision makers. These aspects are very important in the day-to-day usage of the taxonomy independently of the use cases.

Table 6 – Identified good practices with a high relevance to the CSIRT community.

4.1.2 Overview of the medium relevance good practices

The following table contains a quick overview of the good practices identified by the project team as “fairly relevant” to the CSIRT community. Note that each good practice is described in more detail in chapter “0

Details of good practices of medium relevance".

GOOD PRACTICE	REASONING
GP5: A taxonomy should support the tagging of the event(s) leading to the incident or tagging of additional incident information that is relevant	While this is not achieved using the taxonomy directly, it is an important aspect that helps provide actionable information that can be applied in various circumstances, such as impact area, sensitivity, and many more. With this, a CSIRT can add any tag that may be relevant to their day-to-day operations.
GP6: Taxonomies should be understood by external non-CSIRT entities to effectively share information	It is crucial that professionals beyond the CSIRT community can at least understand the basics of the taxonomies, for high-level taxonomies, but not necessarily all taxonomies this includes professionals such as politicians, lawyers, law enforcement, etc. There may be low-level IoC, TTPs, etc. taxonomies that are relevant to more technical audiences.
GP7: Taxonomies should be created or updated using existing taxonomies or standards	Since the use of a taxonomy is not limited to sharing incident information or reporting, it may very well be valid to create a taxonomy based purely on the necessities of an organisation. However, if an organisation or CSIRT is mandated to use a certain style of taxonomies, then this good practice does not apply.
GP8: A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event	The essence of this good practice is that CSIRTs should be able to differentiate between what is a real event (incoming notification/report) and what is not. The "what is not" can mean different things, either it is not complete, it is a false assumption or it is incorrect. Sometimes it is not clear from the beginning whether an event is an incident or not.
GP9: A taxonomy should be able to support reporting vulnerabilities	While dealing with incidents, CSIRTs also receive vulnerability information. It may be useful to have a taxonomy that deals with these vulnerabilities, and additionally the taxonomy used for incident detection and response should be able to mark it as a vulnerability (this can be a misconfiguration or an exploitable vulnerability).

Table 7 - Identified good practices with a medium relevance to the CSIRT community.

4.1.3 Overview of other good practices with possible relevance

The following table contains an overview other good practices with "possible relevance" to the CSIRT community. Note that each good practice is described in more detail in chapter "4.4 Details of other good practices with possible relevance".

GOOD PRACTICE	REASONING TO DETERMINE RELEVANCE
GP10: An implemented taxonomy should allow an incident to be traced back to the source[28]	May not be applicable to all CSIRTs because it is likely that the information received by CSIRTs does not specify or are most likely incidents.
GP11: Taxonomies should be both human readable and machine readable	May not be applicable to all CSIRTs because most CSIRTs already use combined human and machine-readable processes as part of their operational activities.

[28] Source can be both the attribution to the threat "actor" or the "attack vector" (I.e. the vulnerability, way of entry)

Table 8 - Identified good practices with a low relevance to the CSIRT community.

4.2 Details of most relevant good practices

This chapter describes the most relevant good practices in detail.

GP1 - Simple top level categorisation

GP2 - The categories within a taxonomy should be mutually exclusive

GP3 - Taxonomies should support performance measurement

GP4 - Taxonomies should have an appropriate level of ease of use

Figure 2 - Overview of most relevant good practice.

4.2.1 Good practice 1: The top-level categorisation of a taxonomy should be simple

4.2.1.1 Problem or gap

Too many top-level categories could increase the complexity of the taxonomy. While a complex taxonomy could be useful for certain use cases (such as feeding threat intelligence, exchanging information between CSIRTs, etc.), the more complex it becomes. Too much complexity can result in the taxonomy becoming too cumbersome for day-to-day usage.

4.2.1.2 Summary

The top-level categorisation of a taxonomy should not exceed 6-7 categories. The categories should not be flat.

4.2.1.3 How does the good practice address the issue and what are the outcomes?

Having only a limited amount of top-level categorisations allows the simplification of the categorisation process, especially by humans, but also by machines. It also allows for selective filtering for sharing, feeding, reporting, etc. When the taxonomy is required to provide reporting at a less technical level, where an abstraction to the top-level categorisation is appropriate, a lower number of more general terms could increase the understanding and clarity for the target audience. If terms that are more technical are required or more complexity is desirable, then additional categories can be considered within a select top-level category, providing both an overview of the categorisation and the complexity necessary.

Additionally, if there are multiple levels of categorisation present, the recipients of the report should be able to see what they consider serious. Since what audience A considers serious is not necessarily the same as for audience B, picking the appropriate depth of categorisation per top-level categorisation allows the customisation of relevant data (relevant to the target audience).

This good practice facilitates the reporting of information to decision makers and the categorisation process. In addition, it allows for clearer mutual –exclusivity. It also enables the control of complexity more effectively. Finally, multiple levels of categorisations allow different levels of detail and technicality to be represented with the same taxonomy.

4.2.1.4 Practical example

TOP LEVEL CATEGORY	BOTTOM LEVEL CATEGORY
Malware	Virus
	Worm
	Trojan

Table 9 - A simplified example of using multiple levels of categorisation

A taxonomy used by an incident handler can also be used to make a non-technical report for management by abstracting to the top-level category. As such, instead of specifying which type of malware caused the incident, it would be sufficient to just use the category of "Malware", instead of specifying a bottom level category.

4.2.2 Good practice 2: The categories within a taxonomy should be mutually exclusive

4.2.2.1 Problem or gap

When processing an incident or the event(s) leading up to the incident, it is possible that an incident may be re-categorised or in multiple categories throughout the incident lifecycle. This may or may not be a desired outcome.

4.2.2.2 Summary

Depending on needs and requirements, the terms within a taxonomy should be as mutually exclusive as possible. In a taxonomy with multiple levels of categories, the bottom level categorisation should be mutually exclusive. Re-categorisation of incidents can influence reporting. A choice needs to be made if only the latest categorisation should apply. Reporting could also include statistics such as: "X incidents of categorisation A have eventually resulted in Y incidents of categorisation B" to keep track of categorisation trends.

4.2.2.3 How does the good practice address the issue and what are the outcomes?

If mutually exclusive categories exist, this makes the first categorisation of an incident easier, allowing for the incident to be machine readable, and reduces ambiguity in the terms.

Multiple attribution enables an incident to be categorised multiple times in different categories, which in some incidents (such as a spear phishing attack/campaign) may provide more information about the incident.

Achieving mutual exclusivity can be achieved by having clear, self-explanatory, and precisely defined terms within the taxonomy. If a taxonomy has several levels of categorisations (i.e. broad top-level categories branching down into more specific categories).

4.2.2.4 Practical examples

A CSIRT receives a set of IOCs listing C2 servers for "malware", as such the initial categorisation in "Malware" is obvious. No further information is available at this time and the Indicator of Compromise (IOCs) are placed into the IDS rules.

When hits have been found in the logs the incident is first marked as malware. Then further investigation is done into the clients, and it is discovered that the malware is "ransomware". The incident is enriched with this info, the sub-category (or tag). This is shared with the community. Enriching the IOCs with this set of information and sharing it improves the quality of the IOCs.

Further investigation is done to check how the ransomware was delivered. It is detected that a previous incident that was marked as "Spam" contains details on how the ransomware was delivered and has information on the message content. The initial incident marked as "Spam" is now set to "Malware / Spam", the malware incident can be updated with a tag "delivered via e-mail" and "Office document".

The result of the client investigation also allows to extend the incident with the ransomware family.

4.2.3 Good practice 3: Taxonomies should support performance measurement

4.2.3.1 Problem or gap

Depending on the size of the CSIRT operational team, it may not be feasible to adequately allocate all required resources to the needs of the incident, as such the prioritisation process needs to take into account the time taken for each type of incident to facilitate planning.

4.2.3.2 Summary

Introducing metrics that assess aspects such as time required to analyse the events leading to the incident, to solve an incident, the effectiveness of solving an incident can provide key information about the incident detection and prevention process. Though the taxonomy itself does not necessarily affect this, it supports the process.

4.2.3.3 How does the good practice address the issue and what are the outcomes?

With metric information, each type of incident can be assessed in terms of the time it takes to solve the incident, the effectiveness of the process used to solve it, as well as the effectiveness of employees in solving incidents. Having a measurement of how many person hours each incident would require, allows for the proper and effective allocation of resources to it, meaning that more critical incidents can be addressed more adequately.

For performance, the following measurements should be considered:

- It should support the measuring of the criticality of a resources - both in terms of complexity and importance (e.g. a simple SCADA web interface may not be complex, but can be very critical). This would assist in the triage of the incidents;
- It should measure the impact of the incident. As such, it would facilitate the prioritisation of the incident;
- It should allow each incident to be measured in terms of the time needed to resolve it and with additional measurements taken in the incident resolution process;
- It should also support in measuring the efficiency of the personnel working on incidents;
- It should measure how quickly the incident was closed.

The taxonomy should provide means or additional fields or terms that can be used for performance measurement. This can either in the form of tags that keep track of the status of an incident, with the time between opening and solving the incident recorded. Additionally, a taxonomy could also have a field to account for opening and closing time of an event or incident. With this information the above measurements could be extracted.

4.2.3.4 Practical example

CSIRTs receive a lot of phishing notices, dealing with these reports is a standard procedure. The amount of time spent on dealing with a standard phishing mail is standard. However, if one of the phishing mails contains a link to the public government website then it requires more (human) effort and time because:

- **It is part of the immediate constituency of a government or CSIRTs;**
- **It is sensitive, in the sense of potentially affecting the image or reputation of the affected entity;**
- **It has the impact that the information provided to and received from the public via the website might be intercepted or altered.**

Having a tag “Government” allows to measure how much time was spent on "gov"-incidents. If the phishing link is placed in “Government” then this will again lead to another incident and an investigation on how the website was exploited.

4.2.4 Good practice 4: Taxonomies should have an appropriate of level of ease of use

4.2.4.1 Problem or gap

If the taxonomy is too complex or has too many terms it becomes impractical and cumbersome to use in day-to-day activities, such as effectively exchanging data with other CSIRTs or feeding threat intelligence. If a taxonomy is too simple, then it may not be able to adequately explain the incident or the context of the incident.

4.2.4.2 Summary

While the definition of ease of use depends very much on who and for what purpose the taxonomy was implemented for, there are a few basic elements that are generally applicable:

- **Completeness of the taxonomy** – for a taxonomy to be implemented, used correctly and for day-to-day usage, the taxonomy should contain all the terms in the semantic vocabulary based on operational requirements. To assure completeness, the best course of action is to learn from national and international standards, or to consult other national CSIRTs on their approach. Cooperation enabled through CSIRT networks will also help assure the completeness of the taxonomy. Additionally, the experience gained from implementing and using the taxonomy on a day-to-day basis, will further uncover and requirements.
- **Ease of first categorisation (triage)** - Mutual exclusivity may not be the top priority of all taxonomies or even feasible in some cases (i.e. during the course of an incident life cycle, the incident changing categories due to further investigation).

  Regarding the use case "Recording events from different sources", it should be taken into account that some events might include information that in fact contain two incidents. For instance, someone complains that a website resource is being abused and forwards the logs, however, when the log is looked into the handler also notices other unwanted behaviour, such as , the system being used as a command and control server , which leads to the event being split into two different incidents.
- **Complexity of the taxonomy** – while the complexity depends very much on the nature of the operational activity (be it for reporting information or exchanging indicator of compromise information). In terms of complexity, the taxonomy should have more than one level of categorisation. Having multiple levels of categorisation will allow for describing an incident in greater detail, as well as to provide a high-level view if needed.
- **Simple first categorisation** – depending on the precision of the terms of the taxonomy, it is possible that by the end of the incident lifecycle, the same incident has been categorised multiple times or that the categorisation changed. Nonetheless, the first categorisation using the taxonomy should be straightforward and simple to do. This can be achieved by either clearly defining each term in the taxonomy and/or using self-explanatory terms within the taxonomy.
- **If the taxonomy is to be used to report to decision makers, the terms have to be understood by the decision makers** - this can be achieved by having multiple levels of categorisation and reverting back to the top level, as well as having either self-explanatory or clearly defined terms.

4.2.4.3 How does the good practice address the issue and what are the outcomes?

When a taxonomy is implemented, it is often used on a day-to-day basis (for incident detection). For this purpose, the taxonomy chosen must be easy to use, which would facilitate the first categorisation of an event.

4.2.4.4 Practical example

The first categorisation can define who deals with the follow-up of the incident, and the resources to be assigned. This initial process should be swift and quick.

When the incident handlers process the queue in the ticketing system, an incident that is categorised as “Copyright infringement” will get lower priority compared to an incident marked as “System compromised”. This facilitates the task of distinguishing the categorisation between these two and to prioritise them.

4.3 Details of good practices of medium relevance

This chapter describes the good practices of medium relevance in detail.

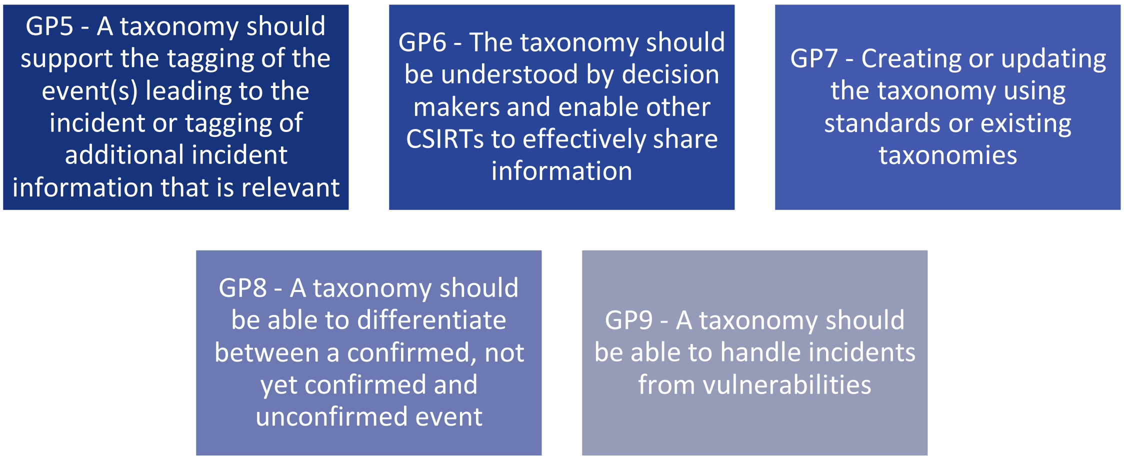

GP5 - A taxonomy should support the tagging of the event(s) leading to the incident or tagging of additional incident information that is relevant

GP6 - The taxonomy should be understood by decision makers and enable other CSIRTs to effectively share information

GP7 - Creating or updating the taxonomy using standards or existing taxonomies

GP8 - A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event

GP9 - A taxonomy should be able to handle incidents from vulnerabilities

Figure 3- Overview of good practice of medium relevance.

4.3.1 Good practice 5: A taxonomy should support the tagging of the event(s) leading to the incident or tagging of additional incident information that is relevant

4.3.1.1 Problem or gap

Incidents do not have the same impact for everyone, it may not be relevant for different parties to worry about the same event(s) leading to an incident. Furthermore, there is not just one method to add relevant information if there is no field to represent it, or critical information may not be represented appropriately.

4.3.1.2 Summary

A taxonomy should support a tagging system that can represent a varied amount of information.

4.3.1.3 How does the good practice address the issue and what are the outcomes?

While tagging may not necessarily be performed by the taxonomy itself, as framework which allows customisable tags to be added would provide a critical component that would save time and provide actionable information.

The following non-exhaustive list of tags could be considered:

- **Impact Area** - Not every incident has the same impact for every CSIRT, as such having a tag that allows the specification of the impact area may allow the facilitation of aggregation and searching of events by relevance to a particular area or sector (see use cases). Impact area can include, but is not limited to, industrial sectors, governmental sectors, public sectors, etc.
- **Constituency or victim type** - In some specific cases, such as reporting to politicians. This often entails answering statistical questions such as "how many DDoS attacks were observed" or "how many of these attacks involved the government as a victim". This can also be solved by allowing an incident handling team to set tags for “victim type” or “constituency”. It is important to note that victim type can include everything that is in the incident database, such as “all victims with the e-mail address @belgium.be”, “gov”, “sme”, “critical infrastructure”.
- **Law Enforcement applicability** (if relevant) - Law enforcement sometimes ask that the reporting includes the marking if an event or incident has led to a criminal act (sometimes also differentiating between a severe or light act). There is an ongoing debate to know what the full definition of cybercrime, therefore a use for such a tag should be foreseeable.

- **Sensitivity** - Due to the sensitive nature of some of the incidents, the event(s) leading up to the incident, the source or target of the incident that occur, a tag that accounts for the sensitivity of the incident could be included too. Alternatively, the taxonomy could be extended to include sensitive or another taxonomy (such as the Traffic Light Protocol) could be used. The TLP tags are as follows[29]:
  - "TLP: RED" - Information exclusively and directly given to (a group of) individual recipients. Sharing outside is not legitimate;
  - "TLP: AMBER" - Information exclusively given to an organisation; sharing limited within the organisation to be effectively acted upon;
  - "TLP: GREEN" - Information given to a community or a group of organisations at large. This information cannot be publicly released;
  - "TLP: WHITE" - Information can be shared publicly in accordance with the law.
- **Classification** - Depending on the content of the incident information or with whom this information is shared, it may be appropriate to use a tag to represent the classification of the incident information. For example:

For governmental or nation incidents.

With the NIS directive encouraging the sharing of information within the EU, a taxonomy such as EUCI[30] (EU classified information marking) may be applicable as use for tagging using the following:

- **"TRÈS SECRET UE/EU TOP SECRET";**
- **"SECRET UE/EU SECRET";**
- **"CONFIDENTIEL UE/EU CONFIDENTIAL";**
- **"RESTREINT UE/EU RESTRICTED".**

It is also possible to use a central management tag list as a starting basis.

Tagging may not be an effective solution for every problem. However, it does provide a convenient way to add critical and actionable information to the data or to account for any future types of incidents that could arise with relative minimal effort.

It is also important to consider that if all this information were to be implemented in the taxonomy itself and not around the taxonomy, then there would be a large number of sub-branches making it cumbersome or even unusable. Not every CSIRT will use all the branches either, some may only require "Malware" -> "Ransomware".

4.3.1.4 Practical example

If we consider ransomware, it makes sense to categorise it as "Malware" and then as "Ransomware". However, it would also make sense to know the version of the ransomware (e.g. crytpolocker version X). The following information would be relevant and provide actionable information if existing as tags for the example:

- **The delivery method, e.g. via an embedded Office document;**
- **Whether or not a decryption method is publically available;**
- **Where the incident took place, e.g. in a public administration;**
- **The impact of the loss of all the data.**

[29] https://www.circl.lu/pub/traffic-light-protocol/
[30] http://www.consilium.europa.eu/en/general-secretariat/corporate-policies/classified-information/

4.3.2 Good practice 6: Taxonomies should be understood by external non-CSIRT entities to effectively share information

4.3.2.1 Problem or gap

When reporting to decision-makers, these decision-makers will make a judgement based on their understanding of the information given to them. Therefore, to ensure that the information given to them is understood as intended, the terminology must be carefully selected or clearly defined. Therefore when reporting to the government or media (who may intend to share acquired information with the public), need to fully understand the information they will base a decision on or share to the public.

4.3.2.2 Summary

A taxonomy should allow decisions to be based on the information and wording provided.

4.3.2.3 How does the good practice address the issue and what are the outcomes?

A taxonomy should reflect that by having a simple and limited top-level categorisation (GP1) of clearly understood and defined terms. The level of simplicity should be defined based on operational needs or intended target audience.

4.3.2.4 Practical examples

Reporting towards politicians, law enforcement, lawyers, judges, etc. I.e. entities that may be required to make decisions based on incident information received.

4.3.3 Good practice 7: A taxonomy should be created or updated using standards or existing taxonomies

4.3.3.1 Problem or gap

Taxonomies are usually created to serve a specific organisational purpose. If the taxonomies cannot easily be mapped, this complicates the exchange of incident information with other CSIRTs..

4.3.3.2 Summary

Taxonomies should be designed, implemented, or revised using established standards or taxonomies in use by other CSIRTs.

The following are standards, which provide relevant information for creating, implementing, and exchanging information using taxonomies:

- RFC 2350 [31]- Best Current Practice; Expectations for Computer Security Incident Response;
- RFC 2828 [32]- Internet Security Glossary;
- RFC 5070[33] - The Incident Object Description Exchange Format;
- NIST SP 800-61 revision 2[34] - Computer Security Incident Handling Guide;
- NIST SP 800-150[35] - Guide to Cyber Threat Information Sharing.

Many CSIRTs already use taxonomies to handle incidents or report on them, having already conducted research into taxonomies themselves, or learned from using the taxonomy what worked and what did not work. Therefore, it is advisable to contact the respective national CSIRT before creating a completely new and untested taxonomy. .

[31] RFC 2350 - https://tools.ietf.org/html/rfc2350
[32] RFC 2828 - https://tools.ietf.org/html/rfc2828
[33] RFC 5070 - https://tools.ietf.org/html/rfc5070
[34] NIST SP 800-61 r2 - http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf
[35] NIST SP 800-150 http://csrc.nist.gov/publications/drafts/800-150/sp800_150_second_draft.pdf

4.3.3.3 How does the good practice address the issue and what are the outcomes?

Contacting other CSIRTs to gain more information can save a lot of effort and resources in terms of taxonomy creation and implementation, as well as usage.

4.3.3.4 Practical examples

Many CSIRTs will exchange threat information via MISP. Information that is stored in MISP often is already categorized by other participants. By using a taxonomy that also is supported in the tagging of information in MISP taxonomies, then the category information put there by others remains in the events when synchronizing them to local instances.

4.3.4 Good practice 8: A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event

4.3.4.1 Problem or gap

Taxonomies do not all account for contextual difference with regards to intrusions vs. misuse, or a false positive for a vulnerability. For instance, if an account is compromised on a website, it is not necessarily the website that has suffered an intrusion, it could be merely the credentials of a user being misused. This distinction can provide key information in an investigation.

4.3.4.2 Summary

A taxonomy should have the necessary fields to be able to distinguish between whether an incident occurred (e.g. an intrusion or vulnerability), or an unconfirmed event (e.g. the misuse of the credentials of a user, but not an actual breach).

As such, the following status terms could be applied:

- “Confirmed”;
- “No yet confirmed”;
- “Unconfirmed”;
- “Pending”.

In addition, the taxonomy should be accompanied with a document that clearly defines, in terms of organisational requirement, for the CSIRT each of the above statuses. This may enable external parties to gain understanding into the statuses.

4.3.4.3 How does the good practice address the issue and what are the outcomes?

For this, a field can be added upon the taxonomy to distinguish such detail in relevant cases. This helps in the investigation and proper categorisation of an incident.

4.3.4.4 Practical examples

The following example provides a high level illustration of the life cycle of an event:

1. **A report arrives from a governmental institute that one of the other governmental institutes that provides them data (via a web interface) has become unavailable. The incident status is “Pending".**

2. **Initial investigation shows that the whole network is indeed unreachable.**

3. **Incident responders try to contact the other .gov network but their SIP telephone infrastructure is not reachable Initial categorisation can be a DDoS attack. The incident status is “Not yet confirmed".**

4. **When contacting the NOC of the ISP that provides access to both organisations it turns out there has been a human error causing a routing problem. The incident is an "Unconfirmed event".**

4.3.5 Good practice 9: A taxonomy should be able to support reporting vulnerabilities

4.3.5.1 Problem or gap

Not every incident that is reported stems from a security incident. CSIRTs also receive vulnerability information. A taxonomy should be able to handle this information (if applicable). Note that a vulnerability can be a misconfiguration or something exploitable.

4.3.5.2 Summary

A taxonomy should be able to handle incidents and vulnerabilities if applicable.

4.3.5.3 How does the good practice address the issue and what are the outcomes?

If the taxonomy can handle the sharing and reporting of vulnerability data, then those issues can be resolved more effectively if the proper authorities are alerted. This can be achieved by adding the required categories to deal with vulnerabilities, as well as any other incident type to the taxonomy.

4.3.5.4 Practical examples

A DNS server open to a DNS amplification attack needs to be reported just as any other incident would be.

4.4 Details of other good practices with possible relevance

This chapter describes the good practices with possible relevance in detail.

GP10: An implemented taxonomy should allow an incident to be traced back to the source

GP11: Taxonomies should be both human readable and machine readable

4.4.1 Good practice 10: An implemented taxonomy should allow an incident to be traced back to the source

4.4.1.1 Problem or gap

Current taxonomies do not always allow the tracing of incidents back to the source.

4.4.1.2 Summary

The implemented taxonomy should provide sufficient information based on the threat event and the threat actor that allows the tracing back of the events to the source (can be the source of entry (vulnerability, attack vector) or the source the attribution. However, when sharing sources and possibly doing attribution, different privacy related considerations are involved. For instance, to the press it is not appropriate to disclose the exact nature of the victim(s), while it might be indispensable when dealing with another CSIRT .

4.4.1.3 How does the good practice address the issue and what are the outcomes?

Implementing additional controls around the taxonomy allows the assurance of having the full picture of any given event. This can be achieved by adding some additional metadata or terminology to track the source of the incident.

4.4.1.4 Practical example

A report of a security incident obtained via a public feed is handled with a different priority than a report received from a local IDS that's been setup with your own provided set of rules.

Good practice 11: A taxonomy should be able to differentiate between an attempt and a successful incident

4.4.1.5 Problem or gap

Depending on the criteria of what constitutes an incident, when automatically aggregating incidents, there is not always a clear way do differentiate between what is an attempt and successful incident. If there is no indicator to differentiate both cases, the attempts could lead to a false positive. Therefore, it would not be useful as actionable information.

4.4.1.6 Summary

A taxonomy should be able to differentiate between attempts and successful incidents, and differentiate between malware and legitimate software wherever possible. Adding an additional field to track whether it was a successful incident, an attempt, or a false positive.

4.4.1.7 How does the good practice address the issue and what are the outcomes?

It provides a more comprehensive picture of the incidents. Additionally, it would facilitate the searching of incidents if required by adding an additional search option.

4.4.1.8 Practical example

Sites that run WordPress often see scans for wp-login. The web request for wp-login is something was handled by automated scanning and is nothing more than just a login attempt. A successful request for the /admin/ pages after the login, however, means a successful incident.

4.4.2 Good practice 11: Taxonomies should be both human-readable and machine readable

4.4.2.1 Problem or gap

When automatically processing incident information, it is important that the incident handlers understand the information that is being processed.

4.4.2.2 Summary

This can be achieved by clearly defining the terminology of the categories.

4.4.2.3 How does the good practice address the issue and what are the outcomes?

Clear definitions allow the incident to become machine-readable in addition to facilitating a first categorisation of an incident. This can be achieved with clear definitions for each term within the taxonomy. This can be achieved by keeping additional documentation with each of the terms defined.

4.4.2.4 Practical example

Examples of this are the JSON files that are included in the MISP taxonomy list. The https://github.com/MISP/misp-taxonomies/blob/master/enisa/machinetag.json

4.5 Linking the good practices to the use cases

This section summarises which good practices of high and medium relevance (as defined in chapters 4.2 and 4.3) apply to which use cases (as defined in chapter "3.2 Linking taxonomies to use cases based on defined requirements and metrics") according to the CSIRTs that have been involved in this study.

4.5.1 UC.TI.101 - Recording events from different sources

The table below lists which good practices of high and medium relevance support achieving use case "UC.TI.101 – Recording events from different sources". The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP1 - Simple top level categorisation	For this use case, having multiple levels of categorisation is also beneficial, as it allows adjusting the level of complexity.
GP2 - Mutual exclusivity	Mutual exclusivity enables a more effective recording of events. As each distinct source may require specific terminology or present a different context.
GP3 - Performance measurements	As this is concerned with different sources, this practice would allow the analysis of each source and determine statistics such as time taken per source or person hours required to process an incident per source.
GP4 - Ease of use	The level of use required depends on the types of source from which events are recorded.
GP5 - Tagging	With the use of tagging, the source can be used as a tag.
GP8 - A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event	If the source of the event is, for example, a website with authentication, then knowing how the event occurred is important. Furthermore, knowing the initial results of the event may also provide crucial information.
GP9 - Vulnerability incidents	As the incidents stem from different sources, it is highly likely that there will be a vulnerability among the source. As such, the taxonomy needs to be able to handle such information.

4.5.2 UC.TI.002 - Automatic deduplication

The table below lists which good practices of high and medium relevance support achieving use case "UC.TI.002 - Automatic deduplication". The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP2 - Mutual exclusivity	If the categories are not shared and mutually exclusive, this makes removing duplicates easier. If there are two identical incidents, which are effectively identical but categorized differently, then a duplicate could remain.

APPLICABLE GOOD PRACTICE	REASONING
GP4 - Ease of use	The level of detail affects the effectiveness of deduplication. In this case, a lower level of detail and complexity should be implemented. Since more terms may cause identical incidents to be categorised differently, and therefor duplicate.

4.5.3 UC.TI.003 - Ability to export in other taxonomies

The table below lists which good practices of high and medium relevance support achieving use case "UC.TI.003 - Ability to export in other taxonomies". The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP2 - Mutual exclusivity	If there is mutual exclusion, then there should be clear divide in the concepts represented by the taxonomy, as such choosing to which category another should be exported/mapped to make it simpler.
GP4 - Appropriate ease of use	If a taxonomy has a higher level of completeness is terms of incident types, then exporting into other taxonomies can be facilitated. This is especially true if the information is exported into a more general taxonomy, where categories require merging.

4.5.4 UC.TI.104 - Ability to aggregate and search events in the data

The table below lists which good practices of high and medium relevance support achieving use case "UC.TI.104 - Ability to aggregate and search events in the data". The "reasoning" column describes how each good practice supports the corresponding use case. The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP1 - Simple top level categorisation	Simple top-level categories may simplify the aggregation and search process, as it would allow for aggregation and search at different levels of categorisation.
GP2 - Mutual exclusivity	If a taxonomy has a higher level of completeness is terms of incident types, then exporting into other taxonomies can be facilitated. This is especially true if the information is exported into a more general taxonomy, where categories require merging.
GP4 - Ease of use	If a taxonomy has a higher level of completeness in terms of incident types, then exporting into other taxonomies can be facilitated. This is especially true if the information is exported into a more general taxonomy, where categories require merging.
GP5 - Tagging	The use of tags would reduce the time it would take to aggregate and to search event data, by giving more options to aggregate and search for.
GP7 - Using standards and existing taxonomies	Aggregation of events is facilitated if the terms are already well defined and proven to work if already implemented elsewhere.
GP8 - A taxonomy should be able to differentiate between a confirmed, not yet	This helps with both the aggregation and searching of events, as it allows another option to be used as a filter.

APPLICABLE GOOD PRACTICE	REASONING
confirmed and unconfirmed event	
GP9 - Vulnerability Incidents	For aggregation purposes, being able to differentiate an incident stemming from a vulnerability from other sources may provide valuable information.
GP11 - A taxonomy should be able to differentiate between an attempt and a successful incident	This would simplify the aggregation and searching of events, i.e. provide additional information that can be used to refine the search or aggregation category

4.5.5 UC.TI.105 - Ability to exchange data with other CSIRTs

The table below lists which good practices of high and medium relevance support achieving use case "UC.TI.105 - Ability to exchange data with other CSIRTs". The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP2 - Mutual exclusivity	If there are distinct fields, then the data can be shared effectively at greater ease.
GP4 - Ease of use	This practice depends on the ease of use of the taxonomy used to exchange within the CSIRT network. If all the CSIRT members use the same taxonomy, then the ease of use needs to be adequate to provide a maximum amount of information.
GP6 - Understanding of decision makers	Information can also be exchanged between LEAs and politicians, i.e. law enforcement and policy makers. Therefore, it is important that the information is shared in such a manner that leaves no doubt about the meaning of it.
GP8 - A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event	The exchanged information should contain information regarding whether the event was confirmed, not yet confirmed and unconfirmed.
GP9 - Vulnerability Incidents	To share the maximum amount of relevant information, the type of the information should be included.
GP 12 - Definition and the broadness of a concept within the taxonomy should be precise and self-explanatory	Helps automated sharing of data with other CSIRTs.

4.5.6 UC.IM.006 - Feeding threat intelligence

The table below lists the good practices of high and medium relevance that support achieving use case "UC.IM.006 - Feeding threat intelligence". The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP1 - Simple top level categorisation	Makes the selection process for which types of global level incidents are shared more effective.
GP2 - Mutual exclusivity	This ensures the clear distinction between categories, allowing the network to be fed automatically more effectively, as there are no ambiguities.
GP3 - Performance measurements	Complexity is an important factor for incident handling. As such complexity needs to be measured appropriately to allow the estimation of required resources for day-to-day activity.
GP4 - Appropriate ease of use	This will affect the level of detail the intelligence that is fed possesses, and whether the system that is being fed can handle the level of detail and completeness of the terms used.
GP5 - Tagging	With a tagging system, the relevant incident or event information can be picked out quicker by providing a term to group incident, or even a tag "For threat intelligence"
GP8 - A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event	Depending on who intends to use the intelligence after it is fed into threat intelligence, it is important for proper and clear definitions to be in place that are also shared over the threat intelligence feed if possible.
GP9 - Vulnerability Incidents	To share the maximum amount of relevant information, the type of the information should be included.
GP11 - A taxonomy should be able to differentiate between an attempt and a successful incident	This is important information to share for threat intelligence, as it provides key information that can help determine the severity of an event.
GP 12 - Definition and the broadness of a concept within the taxonomy should be precise and self-explanatory	Having clearly explain concepts in place allows the sharing to be more effective. I.e. alleviate any confusion that could arise from terms.

4.5.7 UC.IM.007 - Incident report management

The table below lists which good practices of high and medium relevance support achieving use case "UC.IM.007 - Incident report management". The "reasoning" column describes how each good practice supports the corresponding use case.

APPLICABLE GOOD PRACTICE	REASONING
GP1 - Simple top level categorisation	Depending on the required complexity, having a simple top level category that can be abstracted to for reporting in high level, or going to bottom level detail for technical level.

APPLICABLE GOOD PRACTICE	REASONING
GP2 - Mutual exclusivity	When reporting, it is important to have a clear understanding of the concepts presented, regardless of their complexity. As such being mutually exclusive would enable the target audience to have a very clear distinction between the terms and associated concepts.
GP4 - Appropriate ease of use	This again highly depends on the intended target audience. Including required complexity, completeness, etc. Regardless of ease of use, it is an important practice to consider.
GP6 - Understanding of decision makers	Clearly defining a taxonomy may be very important depending on the audience of the report. If the audience is a politician with the power to make decisions that potentially affect others, the each high level term must be clearly defined.
GP8 - A taxonomy should be able to differentiate between a confirmed, not yet confirmed and unconfirmed event	Being able to report separately on confirmed, not yet confirmed and unconfirmed event is important.
GP11 - A taxonomy should be able to differentiate between an attempt and a successful incident	Being able to report separately on attempted and successful incidents is important.

Table 10 - Linking of use cases with identified good practices

5. Case studies demonstrating the use of taxonomies

Case studies provide relevant information gathered from the CSIRT community to illustrate each of the cases specified in this chapter. The objective of these case studies is to illustrate the use of taxonomies in CSIRT operational activity while taking into account the use cases established in previous chapters.

This chapter contains three case studies. Each case study is introduced in terms of the situation presented and the involvement of a taxonomy. Further, the case is analysed taking into account CSIRT operational activities and a brief conclusion illustrates what can be drawn from the case and what the impact is.

The following case studies will be explored in the sections below:

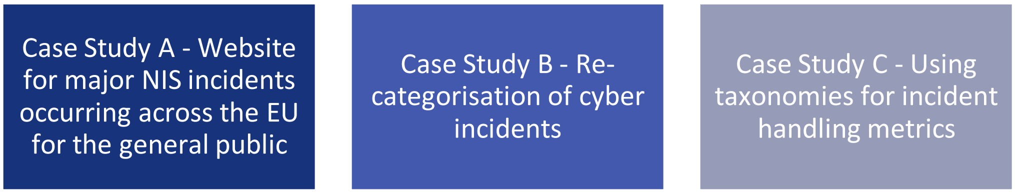

Case Study A - Website for major NIS incidents occurring across the EU for the general public

Case Study B - Re-categorisation of cyber incidents

Case Study C - Using taxonomies for incident handling metrics

Figure 4 - Overview of the case studies in this chapter.

The figure below is the outline of the case studies, the case studies in this chapter will adhere to this structure:

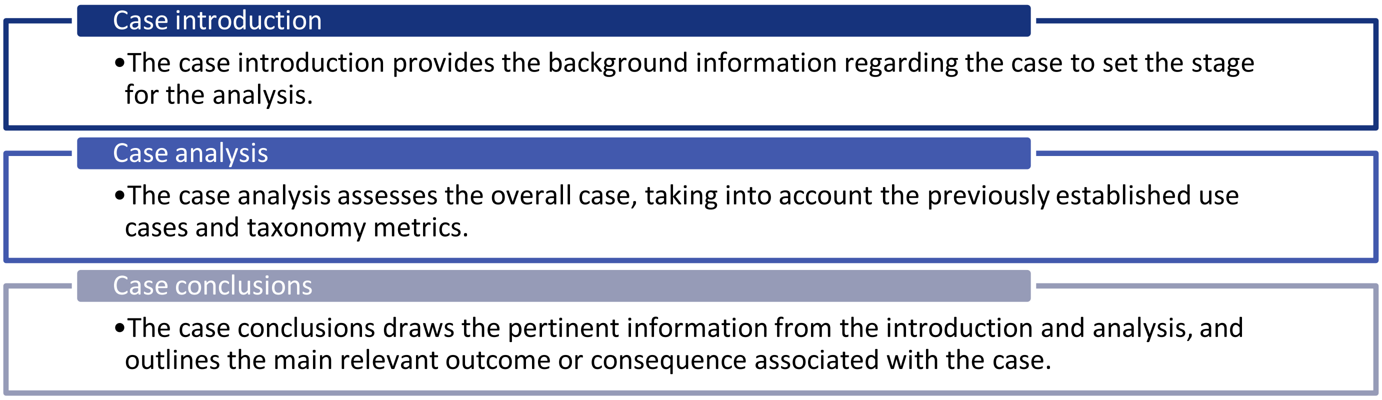

Case introduction

- The case introduction provides the background information regarding the case to set the stage for the analysis.

Case analysis

- The case analysis assesses the overall case, taking into account the previously established use cases and taxonomy metrics.

Case conclusions

- The case conclusions draws the pertinent information from the introduction and analysis, and outlines the main relevant outcome or consequence associated with the case.

Figure 5 - Structure of the case studies in this chapter.

5.1 Case Study A: Website or page dedicated for major NIS incidents occurring across the EU for the general public to support the use of taxonomies

5.1.1 Case introduction

Information about NIS (Network and Information Security) incidents is increasingly valuable to the public and businesses, particularly small and medium-sized businesses as their livelihood could depend on the integrity and confidentiality of their information. In some cases, such information is already provided via websites at the national level, in the language of a specific country and focusing mainly on incidents and occurrences with a national dimension. Given that businesses increasingly operate cross-border and citizens use online services, information on incidents should be provided in an aggregated form at EU level.

According to recital 40 of the NIS Directive[36], the secretariat of the CSIRTs Network is encouraged to maintain a website or host a dedicated page on an existing website where general information on major NIS incidents occurring across the Union is put at the disposal of the public, with a specific focus on the interests and needs of businesses. CSIRTs participating in the CSIRTs Network are encouraged to provide on a voluntary basis the information to be published in this website. However, such a website is not supposed to include confidential or sensitive information.

5.1.2 Case study

To ensure a sufficient level of incident information contained in a publically accessible location, the taxonomy should have the right attributes. As such, the following are some of the requirements that may be applicable for such public information:

- **Easy to understand** - While CSIRTs can contend with complex and technical terms, to ensure public understanding of each of the terms used, the terms should remain simple and general, or at least clearly defined somewhere;
- **Both human and machine-readable** - Due to the number of incidents that could potentially be submitted to this public portal, it may be cumbersome to process them all manually, as such automating the process for machine use may be beneficial. However, since the target audience is not likely to have the same experience level as a CSIRT expert, it should to be understood by humans too. If incident information is received from a member CSIRT or constituency and if they followed the taxonomy properly, the system can feed and automatically understand the incident information. Effectively easing the integration of incident information. It is also worth noting that machine readable information may be more technical than strictly human readable information;
- **Have more than one detail level** - If a member of the public has more knowledge in the area of a specific incident, it may be useful for that person to be able to get a more detailed explanation of the incident.

For the information to be understood by the public, but also provide more detail should it be required, a taxonomy in this context should have at least one level of categorisation, but no more than three, where the top layer should be as simple as possible and the next layers down providing more information about the incident specifics.

For example the table below, the public may want to know that a major incident involved "information gathering", but may not be interested or may not have the expertise for concepts such as "sniffing". In that case knowing that an incident is of the top level category is important, and the bottom level category provides additional useful details

TOP LEVEL CATEGORY	BOTTOM LEVEL CATEGORY
Information gathering	Scanning
	Sniffing
	Social engineering

Table 11 - Example of 2 levels of categorisation potentially allowing the defining of detail level.

In addition, top-level complexity should neither be too high, i.e.it should be easy to go through the taxonomy, nor it should not be too flat with too many categories . Additionally, to be clearly understood by the public, the terms

[36] http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2016.194.01.0001.01.ENG&toc=OJ:L:2016:194:TOC

within a taxonomy should be straightforward and easy to explain, i.e. avoiding the use of acronyms such as the term "DoS", instead using the term "Denial of Service".

When incidents occurs at CSIRT level and incident information is published or shared upwards (see information path below), the taxonomies implemented in this specific case as defined by the body who shares incident information should be capable of providing context about the incident information. If all n/g CSIRTs have a clear taxonomy defined for their respective communities, then the information has a natural path:

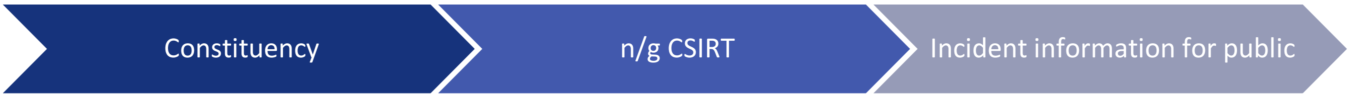
Constituency

n/g CSIRT

Incident information for public

Figure 6 - Information path.

Where each arrow would represent the sharing of incident information, either by sharing the information directly following the format of the sharing medium or mapping to another taxonomy.

For instance, mapping the term "Bot" in eCSIRT.net MkII to "Application compromise" in the ENISA taxonomy, while these terms may not be entirely similar the original concept is not entirely lost.

5.1.3 Case conclusions

While various CSIRTs may not be using the same taxonomy, the adoption of the NIS Directive and the work to be performed in the CSIRTs Network could serve as a catalyst to normalise the taxonomies used. For instance, to push the use of taxonomies with similar terms throughout the CSIRTs community increasing the potential for cooperation and allowing the process of detecting and preventing incidents to become more effective.

5.2 Case Study B: Re-categorisation of cyber incidents

5.2.1 Case introduction

One of the typical operational activities that is part of the responsibility of the CSIRT community is the categorisation and handling of events or incidents. A taxonomy is a crucial component when it comes to categorisation of incidents, not only must the incident handler correctly interpret the incident, but the handler must then also correctly interpret the terminology used by the taxonomy to correctly categories the incident or event. In situations where there is a high turnover of incident handlers or where very large teams operate, the oversight, differing interpretations by different people or difference of experience may lead to the re-categorisation of an incident. Furthermore, some incidents may be part of a bigger picture, requiring the re-categorisation of all involved incidents for correctness.

5.2.2 Case Study

The taxonomy plays a vital role in this situation, the incident handler needs to make a decision for first categorisation (or triage). The incident will be classified on the complexity of the terms of the taxonomy, their mutual exclusivity, the experience of the handler, and the nature of the incident or event leading to the incident. More mutual exclusivity may reduce ambiguities that cause the re-categorisation of incidents. Depending on operational requirements, it may even be the case that an incident or event has multiple categorisations.

This ambiguity may yield additional consequences. For instance, when reporting the decision must be made whether or not to use only the latest categorisation of an incident and ignore all the iterations before that. However, even after the report is published and then a categorisation is prompted, then the report needs to be updated too to reflect such a change. However, an incident can also be re-categorised to provide more clarity on the nature of the incident, for instance:

If an incident is classified as "Abusive content", changing the classification to "Spam" may provide more insight into the event or incident.

Furthermore, an incident may be re-categorised depending on the target audience, as some may be interested in the source of an event, and some on the impact of the event or incident.

Some of the incidents that cause the most ambiguity are those involving spam, malware distribution, and the ones related to fraud. In those cases, if the taxonomy does not have a sufficient level of mutual exclusion, the triage itself becomes more ambiguous.

For example, in the case of a taxonomy containing the following:

INCIDENT CLASS	INCIDENT TYPE
Malware	Infection
Intrusion	Account compromise
Information security	Unauthorised access

Table 12 - Example taxonomy where one incident falls into multiple incident types and classes.

If there is a malware that causes a compromised account or yields unauthorised access where data is ex-filtrated, then the first categorisation could potentially fall into either three categories.

This is an example where the terms contained within the taxonomy are not mutually exclusive enough, causing difficulty in the triage process where the incident handler must rely on experience to aptly categorise the incident or vulnerability (depending on the source). Additionally, re-categorisation can be triggered when more information is available from the incident, not necessarily from internal organisational source, but also when an external source provides new information about an incident.

According to some CSIRTs it is estimated that approximately 20% of incidents get re-categorised into a different category at least once during their lifecycle. This number is affected by the turnover of security event analysts and differences in experience yielding more re-categorisations. Suspicious network traffic incidents have the highest chance to be re-categorised into a more specific category, e.g. a C&C communication detected being re-categorised into targeted malicious email upon investigation.

The chain of events below show an example of an incident being reclassified:

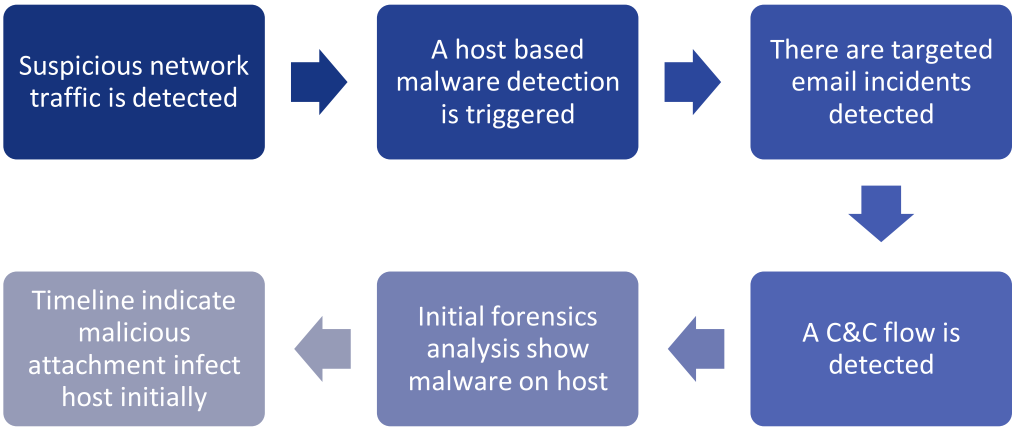
Suspicious network traffic is detected
A host based malware detection is triggered
There are targeted email incidents detected
A C&C flow is detected
Initial forensics analysis show malware on host
Timeline indicate malicious attachment infect host initially

Figure 7 - Chain of events leading to re-classification.

In the chain of events above the incident could potentially be classified and reclassified multiple times depending on the taxonomy used. If the taxonomy contains many terms that are mutually exclusive, then the chance of re-categorisation using the previous case is high. The same however is applicable to a taxonomy with fewer terms, but does not contain mutually exclusive categories.

Another approach that could deal with this instead of re-categorising the incident or event, is with the usage of tags. This would help in the case where an incident may fit into multiple categories. As such, the use of a primary tag could be the main category.

Re-categorisation of an event can affect reporting, especially to customers with potential impact on the SLA (Service Level Agreement). This can also potentially impact the prioritisation of which incidents to treat first. This also has additional implications when considering the confidence of an organisation, there needs to be a repeatable way to gain confidence in report information.

5.2.3 Case conclusion

It may be possible to reduce the re-categorisation of events or incidents, but it may not be possible to remove the re-categorisation entirely due to the chain of events, although the mutual exclusivity of the terms of the taxonomy aid in the reduction, as triage may be performed before the full discovery of the incident life, then re-categorisation becomes unavoidable. It may be feasible for an event to be able to have a primary category, and additional meta-categories or -tags to further enrich the information about the event.

5.3 Case Study C: Using taxonomy for incident handling metrics

5.3.1 Case introduction

A CSIRT has to generate regular reports to their local lawmakers on their recent CSIRT activities. Some of these reports are done regularly (monthly or yearly) to provide an update of the operations. Some reports are prepared by special request from the lawmakers. This is often the result of an issue that was raised at the political level, due to the potential high impact or sensitivity of an event or incident. These issues are not always related to immediate threats, but can deal with topics that have been covered by the press and for which a politician received questions (i.e. parliament or the public).

Similar to the reports for politicians, it is expected that n/g CSIRTs provide a regular overview of the "security state" in their constituency. These quasi-public reports are provided to the press and to other stakeholders.

5.3.2 Case analysis

Although the two type of reports have a different target audience they have a similar construct. They need to downsize the approach of CSIRTs to report based on "technology" to a report aimed towards impact.

A typical line of questioning asked by politicians are or in regards to statistics, such as:

"How many DDoS attacks were there against our government"?

This might introduce a need in the taxonomy for "confirmed" (as being handled via a report) or "not yet confirmed", to have an accurate number of attacks and not divulge speculative data.

These types of questions also immediately make room for a taxonomy that supports identifying the target (or victim). Identification should not be per-victim, but per sector (government, ISPs, financial, etc.). For such a taxonomy the focus would not be on the incident itself, but on the affected entities.

A taxonomy for this purpose may contain terms such as "Energy" or "Transport" (if referring to critical infrastructure operating).

Sometimes it is request that CSIRTs provide the number of serious incidents (as defined by the CSIRT being asked). This can be dealt with by the taxonomy with impact.

It is important to note that the concept of severity can differ between different institutions. For an incident to be considered serious by the government, a real impact might be required, whereas an incident considered serious for a LEA may require there to be a legal follow-up (whether or not it is a criminal case). This difference can be handled by using different types of "impact". Not every CSIRT will have to report this to their local LEA's but as time moves it is expected that LEA at least need a basic reporting on what is going on. This does not describe reporting incidents on a 1-1 basis, it is merely reporting on what has been observed (i.e. the general trends).

There may also be requests on knowing how many sensitive incidents or events were observed. The sensitivity also highly depends on who the entity in question is, or what the audience of their reports consider sensitive, there may be different types of sensitivity as well. For example:

The defacement of a system that host governmental website may be considered as high impact but low sensitivity, whereas, the defacement of a system that handles security clearances may have a relatively low impact but highly sensitive based on the handled data.

This may also lead to the re-categorisation of the incident or event based on the impact or sensitivity.

The definition of impact may also not be static for all entities. Some may be more interested in defining impact as how much work is required to fix an incident, whereas others may define impact as how much is at stake. An example of this would be:

A popular sites containing usernames and passwords is compromised, if this data were to leak, users compromised with a government login may require more effort to be spent on than regular users. I.e. instead of sending general information email to explain the data breach, certain user groups may need to be contacted privately.

A taxonomy in this situation would need to handle such information, either with a specified taxonomy that deals with impact and sensitivity, with terms like "highly sensitive" or "confidential", or with categories within the taxonomy that address those issues (number of terms and number of category levels).

Another aspect to consider is privacy, for some types of incidents at specific sources it may be preferred to provide less specific and more general information, particularly for governmental agencies. To provide the flexibility for a taxonomy to be able to handle aspects such as sensitivity, impact, or privacy adequately, allowing multiple levels of categorisation (number of category levels) would enable that, for example:

Consider a taxonomy with 3 classification levels as defined below, where the type refers to the infrastructure in question, where the sector refers to the sector of the infrastructure and the provider type.

TYPE	SECTOR	PROVIDER TYPE
Critical infrastructure	Energy	Nuclear power plant
		Coal power plant
	Water	Reservoir

Table 13 - Example infrastructure taxonomy

Here the following aspects could be defined in the follow way:

Impact

- An incident at a coal power plant would likely have a lower impact than an incident at a nuclear power plant

Sensitivity

- An incident at a coal power plant would likely have a lower sensitivity than an incident at a nuclear power plant

Privacy

- Depending on the required privacy there are different ways to report on the incident source
  - If high privacy is required, then the incident location can be categorised as “Critical infrastructure”
  - If less privacy is required, but specific details should not be divulged, then the incident source can be categorised as “Energy sector” under “Critical infrastructure”
  - If privacy is not a consideration, then the category can be “Nuclear power plant”

As politicians also often decide on the budget it is important to report all the different operational tasks done by a CSIRT. This involves listing the time spent on things that do not immediately concern a security incident (according to CSIRT experts it is almost impossible to avoid receiving “non-incident” queries via the regular incident reporting communication flow). Typically these questions can be of the following format “I have system X, how do I secure it?”

The taxonomy should support categorising these questions. This could be solved with the use of extra tagging. If multiple questions are received with a similar tag of type (e.g. home routers, WordPress, etc.) this can provide some insight on whether it is a real attack or press involvement (i.e. the press reporting on a certain type of incidents to raise awareness, may also increase the number of reports received for that incident).

If the taxonomy feature terminology that could be used to tracking the current state of an incident, i.e. whether it is solved or still open, may provide some valuable performance measurement metrics if the time between an open incident and closed incident is kept track of. It should also be considered that each team may have different expectation of what constitutes a solved incident.

5.3.3 Case conclusion

There are numerous aspects that should be considered when reporting, for the taxonomy to be suitable, a lot of different metrics and aspects need to be considered. The following are some of the taxonomy considerations that could be made: target audience, source of incident, nature of incident; impact, sensitivity, and privacy considerations

6. Conclusions

This chapter serves to state relevant conclusions based on the previous chapters and their reasoning. It is noteworthy to keep in mind that not every CSIRT will perform their respective operational activities in an identical manner as stated in the previous chapter, as such even the conclusions may not apply to every single CSIRT or constituency. The following conclusions will be detailed below:

Conclusion 1: There is no consensus on concepts and definitions related to taxonomies

Conclusion 2: Ease of use of taxonomies should be evaluated depending on the use case

Conclusion 3: Taxonomies currently lack fields to properly handle the impact of an incident, incidents with no malice intended, explicit fields for ransomware, whether the incident is confirmed, and the differentiation between intrusion attempts and intrusions

Conclusion 4: For reporting incident or event information, it is more effective to have a taxonomy with multiple levels of categorisation used by incident handlers instead of having a taxonomy specifically for reporting

Conclusion 5: The identified areas for potential improvement of existing taxonomies are related to the complexity, contextual information, mutual exclusivity or ambiguity, performance measurement, impact, sensitivity, confidentiality, and purpose of taxonomies

6.1 Conclusion 1: There is no consensus on concepts and definitions related to taxonomies

As mentioned before, there are various terms, which have been used in different ways in various studies. These include: ”information exchange standard”, “ontology”, “taxonomy”, “data type”, “data/field format standard”, “data/Field representation format”, “classification”, “semantic vocabulary”, “field” and “knowledge map”. In addition, it seems that there is no clear consensus on the exact interpretation of each one of these terms in the operational environment. In addition, the definition of taxonomy can vary of the associated use case.

6.2 Conclusion 2: Ease of use of taxonomies should be evaluated depending on the use case

There are several considerations when evaluating ease of use, most importantly, what is considered easy to use by one CSIRT, may not apply to other CSIRTs, as such ease of use should be evaluated from one of the following point of views.

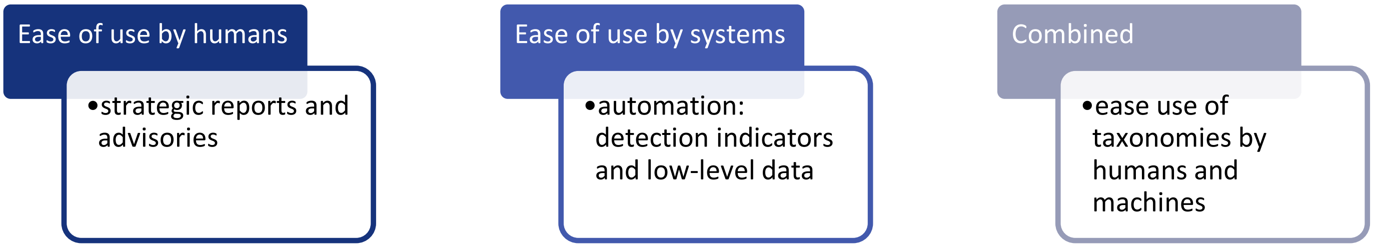

Ease of use by humans
- strategic reports and advisories

Ease of use by systems
- automation: detection indicators and low-level data

Combined
- ease use of taxonomies by humans and machines

6.2.1 Ease of use by humans (for strategic reports and advisories)

There are a number of CSIRTs that use custom taxonomies. Those CSIRTs that are associated with CSIRT networks or are part of a constituency are encouraged or obligated to use the same or similar taxonomies. In some countries, CSIRTs have mandates that oblige them to share specific information with a common agreed taxonomy to effectively share and report upon information.

An observation made during the course of the interviews with the CSIRTs, was that although there were a number of differences in some of the taxonomies, there were similarities in the style of taxonomy. It was observed that a number of CSIRTs use taxonomies based off or modified versions of other taxonomies.

6.2.2 Ease of use by systems (for automation of detection of indicators and low-level data)

A lot of taxonomies are implemented in frameworks to automate various processes, such as the sharing of Indicators of Compromise, the acquisition of incident events, the aggregation and searching of events, etc. These tools enable the automatic sharing of incident information effectively as most of these are community driven. Since they can automate the sharing, it gives more incentive for others to share data and receive shared data on the same framework, encouraging cooperation. The taxonomies used by these frameworks tend to be more technical and complex than those intended for manual use, as the automation makes the process of recording or sharing incident information more effective. Additionally, the use of ticketing systems to keep track of the status of incidents is frequently used.

6.2.3 Combined use of taxonomies by humans and machines

As of the moment of detection of the first event, to the incident being marked as solved, an incident is often treated by both humans and machines using a combination of taxonomies and systems to track or handle the incidents to make the process more effective. Some taxonomies used within the MISP taxonomies are also used outside the framework itself (e.g. such as eCSIRT or ENISA taxonomy). If a machine has the taxonomy completely and automatically implemented, it helps to integrate IoCs. If an IoC is received and the implemented taxonomy was properly followed, then the system can feed and understand automatically, allowing the human counterpart to focus on the incident resolution based on the fed information.

6.3 Conclusion 3: Taxonomies currently lack fields to properly handle the impact of an incident, incidents with no malice intended, explicit fields for ransomware, whether the incident is confirmed, and the differentiation between intrusion attempts and intrusions

From the qualitative assessment of taxonomy fields and terms it the following was concluded that:

- There are currently no adequate fields to define the impact of an incident or event, “Impact” can refer to either political impact or technology impact (See “**Error! Reference source not found.**”);
- Many taxonomies address issues regarding malicious activity (such as “Malware” or “Sabotage”), but very few taxonomies contain fields to deal with accidental or none malicious incidents or events. While this may not represent actionable information for every CSIRT in the community, it may be considered as such by others;
- The majority of the current taxonomies do not currently explicitly categorise ransomware as its own category, however due to the rising likelihood and the general high impact it can have, it should perhaps be considered as a new field;
- Current taxonomies generally do not have fields to account for whether or not the incident, event, or vulnerability is confirmed or not (e.g. “Status” -> “Confirmed”), such information would affect statistics and reporting and may provide useful actionable information;
- The different between an intrusion attempt and an intrusion is not explicitly covered in the majority of the taxonomies assessed. This is however actionable information that may be important to some CSIRTs or constituencies.

6.4 Conclusion 4: For reporting incident or event information, it is more effective to have a taxonomy with multiple levels of categorisation used by incident handlers instead of having a taxonomy specifically for reporting

Taxonomies created are usually (if not always) created by a specific entity or group of entities to fulfil a purpose (i.e. incident handling, classification of information, etc.). As such to fulfil the requirements and aims of said entity or entities, the taxonomy ostensibly is adapted to their needs. This may have for consequence that the sharing or reporting in this taxonomy requires some special consideration (such as converting between taxonomies, which may yield to some loss of contextual information).

When reporting to decision makers, those decision makers will make a judgement based on their understanding of the information given to them. To ensure that the information given to them is understood as intended, the terminology should be carefully selected or clearly defined. Therefore, when reporting to the government or media (who may intend to share acquired information with the public), need to fully understand the information they will base a decision on or share to the public.

The audience may also have a major impact on the taxonomy layout and terms. As the reports made usually cater to specific audiences or the audience expect a specific format and terminology, the taxonomy needs to be adapted to the envisaged purpose (see case study C). In order to accurately cater to different audiences, there should not be a taxonomy dedicated to reporting, instead the level of detail can be selected by the use of multiple category levels.

When sharing incident or event information (or reporting that information to other CSIRTs) there are considerations that need to be accounted for which may provide crucial or actionable information for the entity the data is being shared with:

- Not every incident that is reported stems from a security incident. CSIRTs also receive vulnerability information, as such taxonomy should be able to handle this information (if applicable). Note that a vulnerability can be a misconfiguration or something exploitable;
- The difference between an intrusion and an intrusion attempt may also provide actionable information;
- When automatically processing incident information, it is important that the incident handlers understand the information that is being process.

One of the major areas of improvements, is that there is currently no existing scheme or recommendations to increase the commonalities in the terms within the taxonomies of all the CSIRTs. While some CSIRTs or constituencies may be mandated to record and share incident information within their own national or constituency network in a certain way, there is a lack common and clear accepted definitions for terms.

6.5 Conclusion 5: The identified areas for potential improvement of existing taxonomies are related to the complexity, contextual information, mutual exclusivity or ambiguity, performance measurement, impact, sensitivity, confidentiality, and purpose of taxonomies

This section lists the various areas of improvement that may be possible based on interviews with the CSIRT community, and desk research into the taxonomies and past research. Note that even within these areas of improvement there are things that cannot effectively be addressed, this is mentioned wherever applicable. The areas of improvements are the major areas of improvements identified and generalised as to be applicable to a maximum number of CSIRTs within the CSIRT community.

6.5.1 The complexity of the taxonomy

Too many top-level categories or number of terms could increase the complexity of the taxonomy. While a complex taxonomy could be useful for certain use cases (such as feeding threat intelligence, exchanging information between CSIRTs, etc.). I.e. there may be benefits if all CSIRTs made use of CAPEC, as it has around 504 attack patterns and in some cases 4 or more levels of categorisation, providing a very high level of available detail. The day-to-day usage may be too cumbersome however and the number of terms and categories may increase the time it takes to categorise an incident.

There should be a reasonable middle ground. There may be a difference of expectation between producers and consumers of taxonomies. More data may be better for the consumers of the taxonomy (i.e. more terms to categorise an incident or event), but this can put additional strain on the producer of the taxonomy to keep the taxonomy up to date and ensure it is used appropriately (especially if the producer is the consumer as well). The overall level of complexity depends on the operational requirements and capabilities of CSIRTs. If a taxonomy is too simple however, then it may not be able to adequately describe the incident or the context of the incident.

6.5.2 Contextual information of the incident or event

Taxonomies do not account for the "big-picture" perspective. For instance, if a user had their credentials hacked and account on a website compromised as a direct consequence, it was not the website that was compromised, but the user.

Depending on the criteria of what constitutes an incident, in some situations such as automatically aggregating incidents, there may not always be a clear way do differentiate between what is an attempt and successful incident in every taxonomy. As such, if there is no indicator to differentiate both cases, then the attempts could lead to a false positive. Therefore, it would not be useful as actionable information.

Different taxonomies may not share the same semantic vocabulary, either using the same term for different concepts or using different terms for the same concept. For example:

A taxonomy may use the term "Worm" to describe the actual malware, while another taxonomy may refer to "Worm" to describe the malware action (SANS malware taxonomy compared to Veris respectively). Or a taxonomy may refer to "All incidents which don't fit in one of the given categories should be put into this category" while other taxonomies may use the term "Other" instead, while the former provides a clear instruction, the latter is also self-explanatory.

6.5.3 Re-categorisation of events caused due to lack of mutual exclusivity or ambiguity of terms

When processing an incident or the event(s) leading up to the incident, it is possible that an incident may be re-categorised or in multiple categorisation categories throughout the incident lifecycle. This may or may not be a desired outcome.

Note that the possibility of re-categorisation based on gaining additional information (from internal or external sources) is still present as outlined in case study B.

6.5.4 Measurement of taxonomy and process performance

Depending on the size of the CSIRT operational team, it may not be feasible to adequately allocate all required resources to the needs of the incident. In most cases, taxonomies do not inherently support the measurement of performances of the solving process for each incident. See case study C.

Being able to measure the performance has several operational benefits:

- It may provide crucial information regarding the efficiency to solve an incident for each incident type within the taxonomy;
- It may allow the measurement of the performance of the analyst or handler in charge of the incident by assessing the time taken for solving the incident;
- It may help in the identification of any "bottle necks" in the process that could help smooth the workflow.

6.5.5 Impact, sensitivity and information classification

There are several considerations (see case study C). First, the impact and sensitivity of an incident or an event may not be the same for all constituencies or governments. Second, if the data relates to information that is or should be classified, then there needs to be a way of effectively keeping track of that.

6.5.6 The purpose of the taxonomy

One of the main aspects that needs to be very clear to adequately make, modify, or use a taxonomy is the purpose of the taxonomy. One possible limitation of a taxonomy is about the type of things the terms describe (i.e. whether they describe an action, the source of the incident, critical infrastructure, etc.).Although it is possible to map an action to a threat actor (or have both of them categorised in one taxonomy), it may not be viable to include all domains in a single taxonomy. As such, there is no single taxonomy that can deal with all possible domains currently in existence. This depends on the domain that is to be represented in the taxonomy. Refer to the qualitative assessment chapter for more detailed information on some of the possible domains of a taxonomy.

Furthermore, it is also possible that not all existing domains relevant to the CSIRT community may be currently existing as a taxonomy, there are currently very little examples of taxonomies that address abuse domains.

7. Recommendations

This chapter contains recommendations that can be used related to taxonomies based on the established good practices. These recommendations have been detailed while taking into account the gaps and problems related to taxonomies as discussed with the CSIRTs both during the 11th CSIRT ENISA Workshop, as well as the validation call with CSIRTs on 28 September 2016. Some of the recommendations were included having in mind the tasks to be performed by the CSIRTs network as described in the NIS Directive. Each of the following recommendations outlines an area of improvement, a recommendation on how it could be improved, followed by the advantages of the recommendation:

Recommendation 1: A centralized repository for hosting all relevant taxonomies along with their versions should be set up

Recommendation 2: A small set of common taxonomies for specific use cases should be agreed on EU level

Recommendation 3: An “Other” or “Unknown”, “Tag” field should be used as an indicator to revise taxonomies, if there is an increase in that category with incidents or events of the same type

Recommendation 4: A roadmap towards standardised exchange formats in the CSIRTs community should be established on EU level

Each of these recommendations are detailed in the chapters below.

7.1 Recommendation 1: A centralized repository for hosting all relevant taxonomies along with their versions should be set up by ENISA

Due to the high volume of available taxonomies, the abundance of sources and modified versions of those taxonomies available, there is no clear indication to suggest if the version found is the latest version (note this does not always apply, as national CSIRTs often keep an up-to-date version of their taxonomy on their website).

To deal with this problem, it would be convenient to have a centralized repository (such as GitHub[37]) or website hosted by a central entity where all taxonomies and their versions are visible. Additionally if this included references such as RFCs, and national and international standards, it would enable anyone to adequately revise, update, or create a taxonomy based on a wealth of information. This repository or website should be updated regularly, but only by one team.

This would be a great benefit to the CSIRTs community as not only would it allow the selection of appropriate taxonomies for specific use cases, but it may also provide a general overview of what taxonomies or variations thereof are used by CSIRTs, which may be particularly useful in keeping statistics. This would also allow a way of measuring the usage of a taxonomy.

[37] https://github.com/

7.2 Recommendation 2: A small set of common taxonomies for specific use cases should be agreed upon by CSIRTs at EU level

The high numbers of taxonomies does not provide a clear indication on which taxonomy should be used for what use case. This increases the effort it takes to conduct research into taxonomies if a CSIRT wishes to change, update, or implement a taxonomy.

While a common taxonomy may not be viable under the current taxonomy landscape, if there were a set of five taxonomies for five recurring use cases, it may be a first step towards a common taxonomy or at least towards increased cooperation and exchange within the CSIRT community. Additionally, if the terms within those taxonomies are agreed upon by the CSIRTs Network, then it may provide the benefit that the descriptions for the terms may be incorporated into pre-existing taxonomies. A set of common taxonomies would also ease the automatic and manual exchange of incident or event information.

This would provide example taxonomies based on the requirements of the CSIRTs network, which can then be either implemented or be used to implement a modified version of the taxonomy, saving time and effort that would be spent into researching taxonomies. If all relevant taxonomies were centralised, the statistics gathered could also be used to find common fields, enabling other CSIRTs that do not have the field implemented to consider implementing it, further normalising the taxonomy landscape.

There is not always a clear indication as to the purpose of a taxonomy, or an indication of what a taxonomy describes in terms of the event or incident (i.e. whether it refers to the event, the source of the event, the action that caused the event, etc.).

7.3 Recommendation 3: An “Other” or “Unknown”, “Tag” field should be used by the owners of taxonomies as an indicator to revise taxonomies, if there is an increase in that category with incidents or events of the same type

Due to the fluid nature of incidents, with a growing number of different events leading to incidents and incidents themselves, the taxonomy should be pliant enough to enable future additions to it. While most taxonomies do already contain an “Other” field, the other field should be considered as an indicator to adding new fields to the taxonomy should there be a surge of incidents or events that fall into the other category but are of the same or similar type.

A “tag” field (not the same type of tagging than in MISP), would allow specific information of an event or incident to be recorded. For example:

In a case involving ransomware, it is relevant that it should be categorised ransomware, but also the type of ransomware (such as crypto locker, etc.), if the same tag repeatedly used then it might also indicate the need for a new field.

One important aspect to consider is reporting, as the notification of an increase of a specific incident or event will only happen if it is reported properly.

7.4 Recommendation 4: A roadmap towards standardised exchange formats in the CSIRTs community should be established by the CSIRTs network on EU level

A practical roadmap should be established to allow the CSIRTs community in the EU to jointly work together (for example in the context of the CSIRTs Network and the CEF Cybersecurity Digital Service Infrastructure[38]) to reach

[38] https://ec.europa.eu/digital-single-market/en/news/connecting-europe-facilities-cybersecurity-digital-service-infrastructure

consensus on a number of key success factors to increase the cooperation and information exchange between them. The listing below is an indicative roadmap that could be considered by the CSIRTs community.

1. Agree on use cases that would benefit from taxonomies and exchange standards;
2. Agree on definitions and concepts from an operational point of view for each use case;
3. Perform quantitative assessment on the taxonomies used in each Member State of the EU;
4. Agree on centralised repository in line with Recommendation 1 of this study;
5. Agree on a list of tags/values that can apply across taxonomies;
6. Establish a search engine within the centralised repository to looks for such values across all taxonomies;
7. Agree on a set of common taxonomies per use case to put forward for the CSIRTs community in the EU.

8. Annex

8.1 Preliminary discussion on concepts and definitions related to taxonomies

As mentioned in recommendation 4 of this study, it would be beneficial to have the CSIRTs community build a common interpretation/understanding of key concepts to be used consistently and establish the relation between these key concepts. During the course of this study various discussions have been held with a number of CSIRTs. Although this did not result in conclusive results, the sections below have been added to this study as food for thought and as possible departing point for future work to be performed by CSIRTs. This section summarises various definitions found by the project team on the key concepts. In addition, we summarise some considerations to be made in relation to this study.

8.1.1 Information exchange standard

8.1.1.1 Possible definition

An information exchange standard is a common (non-proprietary) language constraining the information transferred between activities performed by devices, software, or humans, whose goal is to enable effective encoding and decoding of information to successfully perform the required tasks[39].

8.1.1.2 Considerations made in relation to this study

An information exchange standard avoids misunderstanding and ensure a proper documentation and example. For example, an Information Exchange Standard could say that the structure of information should respect a specific taxonomy, all dates should respect a specific data format standard, all chain of character should contains a maximum of 4000 characters etc.

8.1.2 Ontology

8.1.2.1 Possible definition

In computer science and information science, an ontology formally represents knowledge as a set of concepts within a domain, and the relationships between those concepts[40].

8.1.2.2 Considerations made in relation to this study

Ontologies are considered as 3-dimensional although by using a taxonomy in an ontology[41]. The 3rd dimension is "relationship between concepts". In addition, ontology is highlighting the inheritance and differentiation exposed by a taxonomy.

8.1.3 Data representation

8.1.3.1 Possible definition

By data representation is meant, in general, any convention for the arrangement of things in the physical world in such a way as to enable information to be **encoded** and later **decoded** by suitable automatic systems[42].

[39] Information Modeling for Interoperable Dimensional Metrology by Y Zhao, T Kramer, Robert Brown, Xun Xu
[40] https://www.enisa.europa.eu/activities/Resilience-and-CIIP/Incidents-reporting/metrics/ontology/ontology_taxonomies
[41] https://www.enisa.europa.eu/media/news-items/getting-the-right-concept-by-using-the-right-words-ontology-taxonomies-for-critical-infrastructures
[42] "Data Representation", C. M. Sperberg-McQueen, Black Mesa Technology, David Dubin, University of Illinois, Urbana-Champaign

8.1.3.2 Considerations made in relation to this study

Data representation has to be strictly defined and respected in order to keep data accurate. In addition, the data representation will be the way a data will be internally processed and/or stored.

8.1.3.3 Example

The image of a triangle ("data") can be represented as a matrix of dots or as combination of three vectors. Another example is the representation of a date, which can be either a list of characters ("January, 1st 2000") or a simple integer like 946684800.

Data representation is mainly used for internal (within the system implementing taxonomy) representation and automatic exchange. For example: As depicted in the figure below, an email server will use a specific common data representation format to be able to transfer an email to another email server; but the last email server may store the email using its data representation (external). Any tool, like an anti-spam filter having the same data representation may read the email and process it (internal). In addition, a user may connect directly to the email server to see the stored email, it would require the email server to use a data representation format so that it is showing in a human friendly way**.**

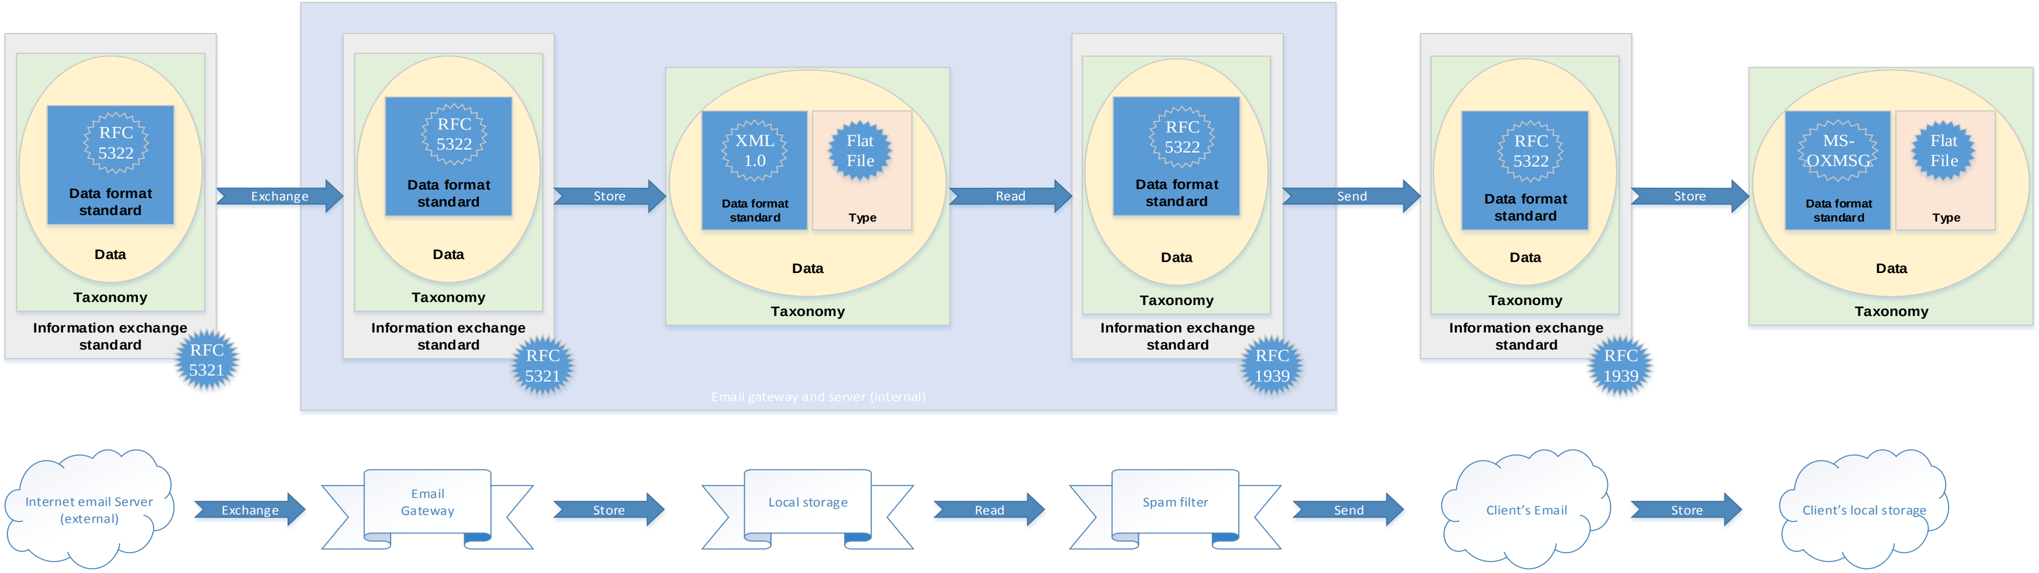
RFC 5322
Data format standard
Data
Taxonomy
Information exchange standard
RFC 5321
Exchange
RFC 5322
Data format standard
Data
Taxonomy
Information exchange standard
RFC 5321
Store
XML 1.0
Data format standard
Flat File
Type
Data
Taxonomy
Read
RFC 5322
Data format standard
Data
Taxonomy
Information exchange standard
RFC 1939
Email gateway and server (internal)
Send
RFC 5322
Data format standard
Data
Taxonomy
Information exchange standard
RFC 1939
Store
MS-OXMSG
Data format standard
Flat File
Type
Data
Taxonomy
Internet email Server (external)
Exchange
Email Gateway
Store
Local storage
Read
Spam filter
Send
Client's Email
Store
Client's local storage

8.1.4 Data type

8.1.4.1 Possible definition

The data type tells what kind of data a value can have. Data types define particular characteristics of data used in software programs. A data type can be Integer, Date, Text, etc.

8.1.4.2 Considerations made in relation to this study

The data type helps to identify the content of a data without defining any representation or format. Data type may be part of a semantic vocabulary.

8.1.4.3 Conclusion

For example, the data type can be "date", it is possible to give a sense of the data but it does not say if the data was stored as a text, integer and it does not say if the data will be written like "January, 1st 2000" or "01/01/2000".

8.1.5 Data format standard

8.1.5.1 Possible definition

A data format standard defines how particular information elements are represented in files or in communications by describing the syntax of a description language and, the semantics associated with those descriptions[43].

8.1.5.2 Considerations made in relation to this study

As data representation format may be included in the semantic vocabulary of taxonomies, its form as a "standard" can also be used in further steps.

8.1.5.3 Examples

For example, the data format standard ISO8601 states that the representation format of the date "January, 1st 2000" should be 2000-01-01

8.1.6 Data representation format

8.1.6.1 Possible definition

Data representation format is the format used for data representation. A format is a way of presenting data, such as centring text or adding dollar signs to represent currency; or the order how to represent a date (DD/MM/YYYY, YYYY-MM-DD, etc.)[44] .

8.1.6.2 Examples

For example, the representation format of a month can be a two-digit integer between 1 and 12. Another representation format could be a value inside a defined vocabulary ("January", "February", etc.)

8.1.6.3 Considerations made in relation to this study

- The data format indicates how incoming and outgoing data should look like
- It gives rules to validate inputs and outputs

[43] https://www.enisa.europa.eu/activities/cert/support/actionable-information/standards-and-tools-for-exchange-and-processing-of-actionable-information
[44] http://mlp12.com/task/114

- The data representation format is particularly useful to allow data exchange and may be included in the semantic vocabulary
- Data representation format is mainly used to present data externally and in most of cases to humans or to external automatic processing. Refer to “Data representation” example.

8.1.7 Classification

8.1.7.1 Possible definition

Classification is designed to group related things together and to define the relationship these things have to each other[45]. In addition, classification is the repartition of events and incidents into classes, not to be confused with the level of classification of a document.[46]

8.1.7.2 Considerations made in relation to this study

One can group related items on different levels. It is possible to classify data types (scalar, enumeration, etc.), as well as data representation formats (long format, short format, etc.) and events to classify taxonomies/ontologies. This is similar to the concept of “attributes” in MISP, which can be put on each level of data.

The NATO level of classification can be seen as an attribute. It is not a classification by itself even if this attribute can be used to do a classification (like the colour, size or brand of candies).[47]

Classification is an action that can be performed a taxonomy at every step and level.

It has been noted that in some cases there is a need for a “default fall-back” as a lot of information is not classified. Therefore, a lot of actionable information starts with a default classification. This is often a problem when data is received for example in JSON format from a CSIRT and it is classified a TLP AMBER. However, it is not always clear whether the JSON file is classified as a whole or just the content.

8.1.7.3 Examples

For example, it is possible to classify candies by colour, size, quantity of sugar, brand, etc. It is also possible to classify people eating candies by age (adult, children, etc.) or by sex (male, female). Another example of data in light of different classifications: a breach of data in a hospital. All data related to non-medical information (how the breach happened, exploited vulnerability, location of the data, type of storage, logs) will be classified as "information leak" by a CSIRT, whereas for the same incident the specific medical data (records, content ...) will be classified differently by a privacy regulator.

8.1.8 Semantic vocabulary

8.1.8.1 Possible definition

A vocabulary to describe knowledge and information assets. The vocabulary must be controlled to ensure that each entry is unambiguous and to also ensure that alternate or less precise terms are excluded[48].

8.1.8.2 Considerations made in relation to this study

As well as classification, semantic vocabulary is part of a taxonomy and therefore has to be used at every further step.

[45] https://www.enisa.europa.eu/activities/Resilience-and-CIIP/Incidents-reporting/metrics/ontology/ontology_taxonomies
[46] Information sharing and common taxonomies between CSIRTs and Law Enforcement
[47] https://github.com/MISP/misp-taxonomies
[48] https://www.enisa.europa.eu/activities/Resilience-and-CIIP/Incidents-reporting/metrics/ontology/ontology_taxonomies

8.1.8.3 Examples

On example of various semantics can be TLP-RED, TLP:RED, TLPRED, TLP RED.

Another example concerns the description of IP. In some cases you can have ambiguity, in some data types you can have IP addresses where some are source IP and others are destination IP. However, sometimes one just needs IP information without caring whether it is source or destination. For example, Net Flow records often mix source and destination IPs. Therefore, in the context of "semantic vocabulary", when using an IP source, some will just interpret it as an IP. Therefore, some data cannot be fully interpreted due to lack of level of detail available.

8.2 Clarifications related to the available definitions of key concepts

8.2.1 "Taxonomy" vs "Ontology"

A graphical example of a taxonomy versus an ontology: if a taxonomy can be represented as a tree then the ontology would be a forest. Another possible example of a taxonomy is phylogenetic tree (animals, mammals, etc.). Yet, another example in a human family, a taxonomy will define the relation between a child and his parents where an ontology will also define the marriage relation between a child of a family and another child from a distinct family.

The namespace of TLP could be considered to be a taxonomy. On the other hand, for example TLP:RED is often used with espionage topic, the relationship thereof is rather an ontology.

8.2.2 "Data representation" vs "data representation format":

Data representation could be seen as an "internal**"** concept (storing, internal exchanges, etc.) while data representation format is more an "external**"** concept (human presentation, serialisation, etc.). For example, a computer represent the first of January as 946684800 but format it as "January 1st 2000".

Another example is an MD5 hash which binary wise is 128 bits of data, which is easy to process by system. Representation of MD5 is usually hexadecimal which in term is used by humans. This example is relevant when exchanging malware samples

8.2.3 "Information classification" (NATO[49]/EU) vs "Classification"

Level of information classification (NATO/EU) is an attribute of a document, this attribute can be used for data classification but it is not a classification by itself. Initially MISP separated this concepts. NATO in the past considered "classification" to be something different than "tagging" in MISP. Although on paper this might sound different, when implementing it seems to be the same.

8.3 Relationships of key concepts

The figure below is a visual representation of our understanding of the relationships of key concepts.

[49] formal categorization and marking of material by level of sensitivity

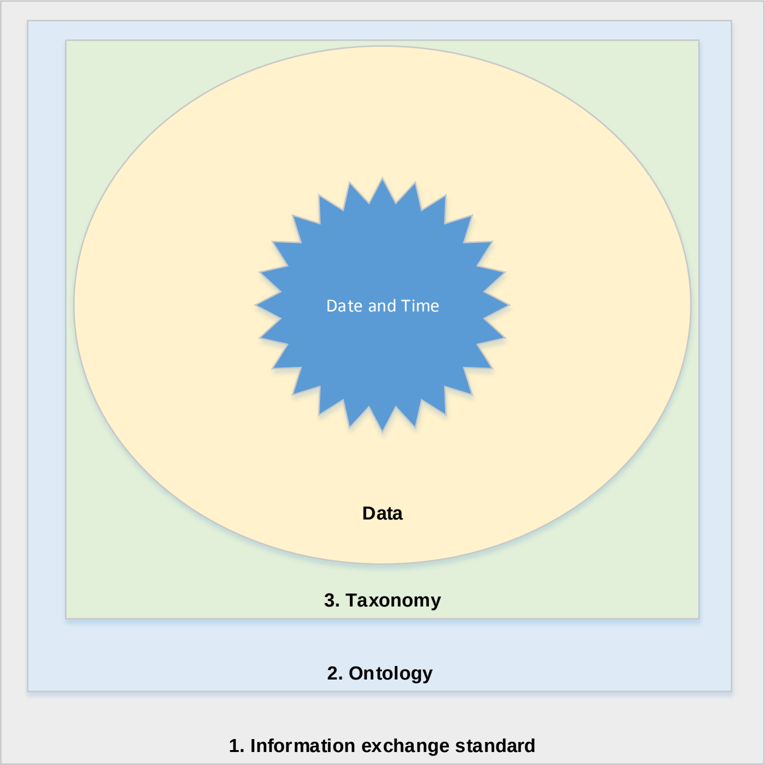

Figure 1 – Relations between Information Exchange Standard, Ontology, Taxonomy and Data/Fields

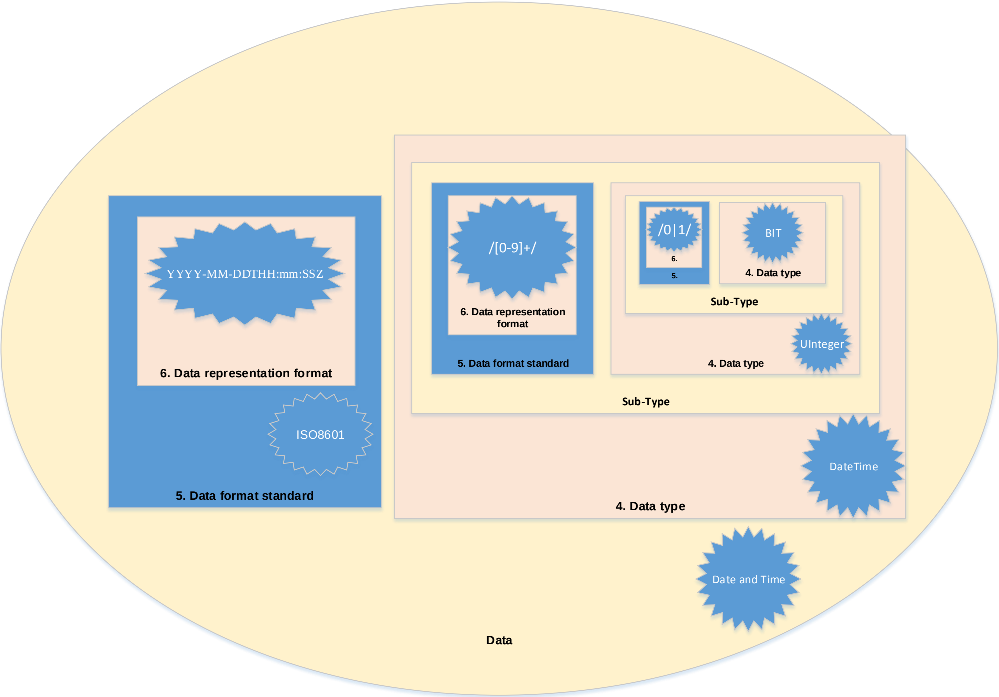

Figure 2 – Close up view of how a data field is structured including Data, Data Type, Data Representation Format, Data Format Standard and Data Representation as a sub-type

Note that in the figure above the different sub types are laid out hierarchically compared to the data type. This is correct but not always the case: one data type can have both multiple data types (horizontally, on the same level) and hierarchically.

8.4.1 A practical case reflecting a possible current understanding of key concepts

The table below shows examples of key concepts applied to storing a date and time in a database:

CONCEPT	CONCEPTUAL EXAMPLE	PRACTICAL EXAMPLE
4. “Data type”	Date and time	N/A
5. “Data format standard”	“ISO8601” which is a **data format standard**	20160407 140648
6. “Data representation format”	On 4 digits followed by month on 2 digits and day on 2 digit, a space and the time with hour on 2 digits, minutes on 2 digits and second on 2 digits	YYYY-MM-DDTHH:mm:SSZ
7. data representation	A number representing all seconds since midnight January, 1st 1970 (GMT)	1460038008

8.5 Statistics gathered by a MISP instance maintained by CIRCL

This annex contains statistical data that provides information about taxonomy usage in a real MISP instance maintained by CIRCL. It allows the collection of statistical conclusions such as the relevance of the terms used by CSIRTs during day-to-day incident handling.

8.5.1 Context

The following section is statistical information provided by CIRCL regarding a single MISP instance and used by a single community at CIRCL that is composed of around 600 organisation that have read access to the data and around 40-50 organisations that contribute to the instance (i.e. creating events)

While these statistics are composed of data from only one community, meaning they may not be representative of all CSIRTs, they illustrate the applicability of some of the conclusions made in this chapter (e.g. “Malware” and “Phishing” are common terms shared within taxonomies due to the large number malware or phishing events overall).

8.5.2 Statistics

The following are statistics regarding the taxonomies used for sharing based on a single community at CIRCL:

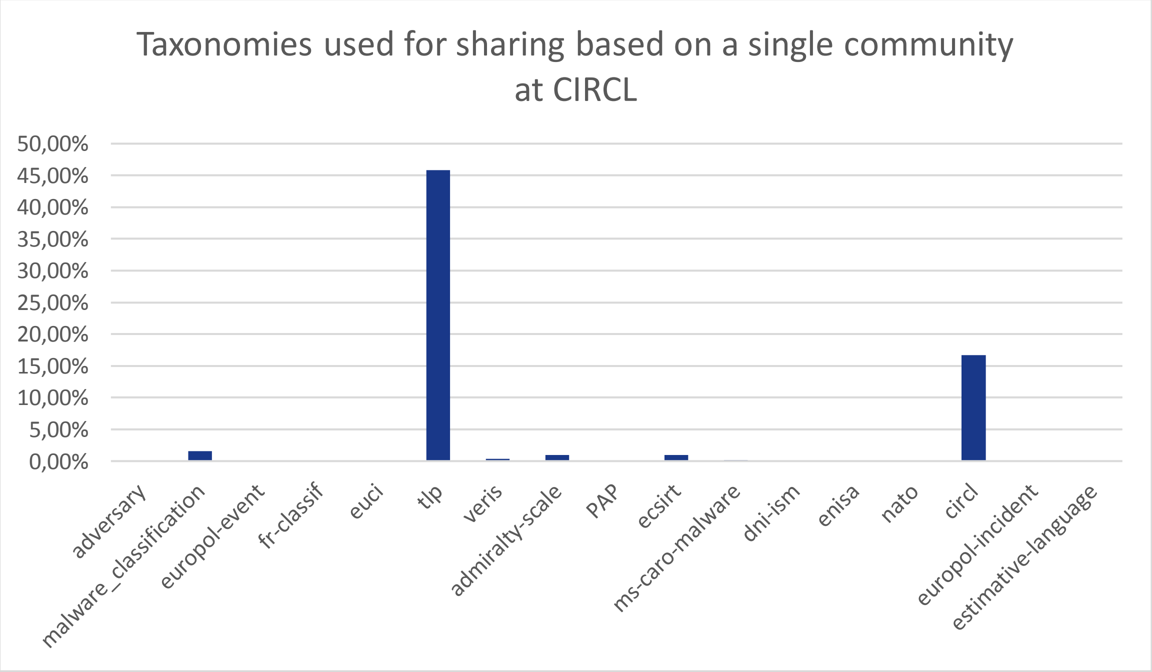

Figure 8 - Statistics of taxonomies used for sharing of information based on a single community at CIRCL.

From the above figure, it is observed that this one community uses parts of five taxonomies to share information. The following taxonomies and their uses are used:

- Using CIRCL and eCSIRT to share incident type and event information;
- The traffic light protocol (TLP) for categorising sensitivity;
- The admiralty scale for categorising the reliability and credibility of the information;
- Malware classification for malware categorisation.

The following are statistics regarding the tags used by the CIRCL taxonomy for sharing based on a single community at CIRCL (Note, due to the use of other taxonomies, the statistics are on the entire dataset, therefore the tags used will only account for 16.72% of the whole dataset):

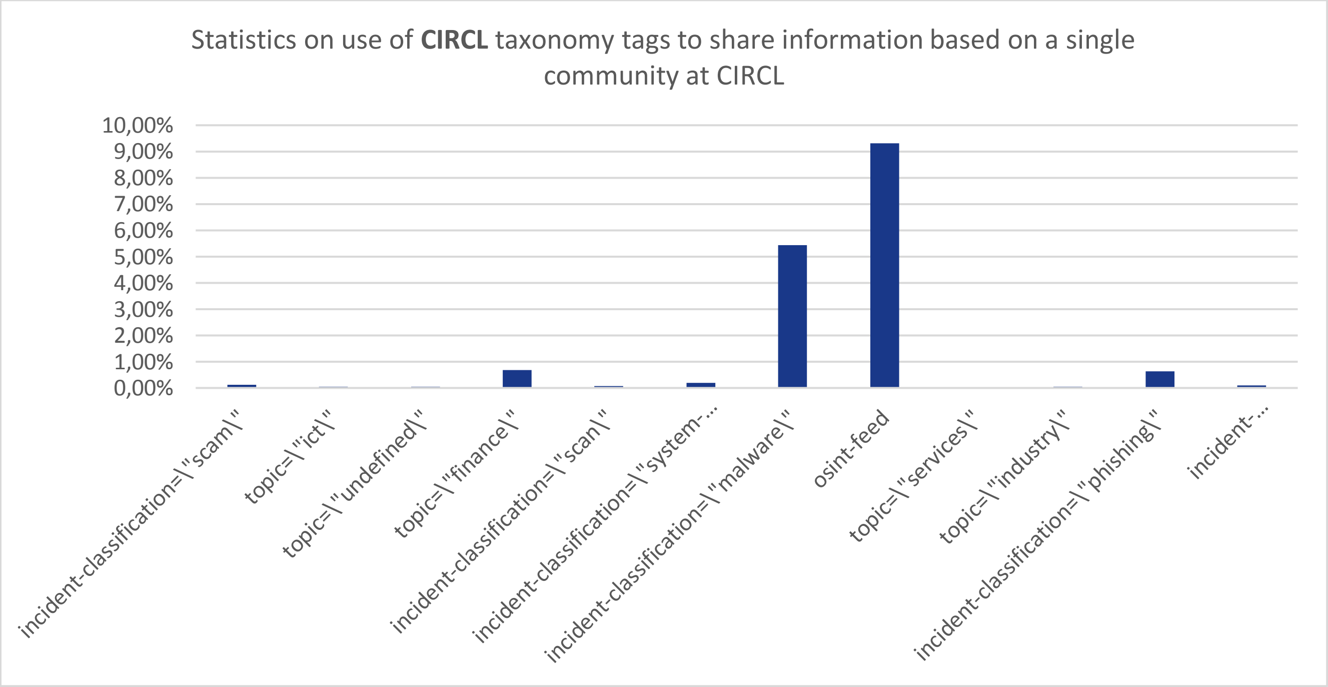

Figure 9- Statistics of the CIRCL taxonomy tags used for sharing of information based on a single community at CIRCL (accounts for around 16.72% of the entire dataset).

From the above figure it is observed that:

- In the CIRCL taxonomy the incidents “Malware” and “Phishing” are the most prevalent, and the “Finance” topic the most used.
- The “osint-feed” tag is highly used.

The following are statistics regarding the tags used by the **TLP** taxonomy for sharing based on a single community at CIRCL (accounts for 45.79% of the whole dataset):

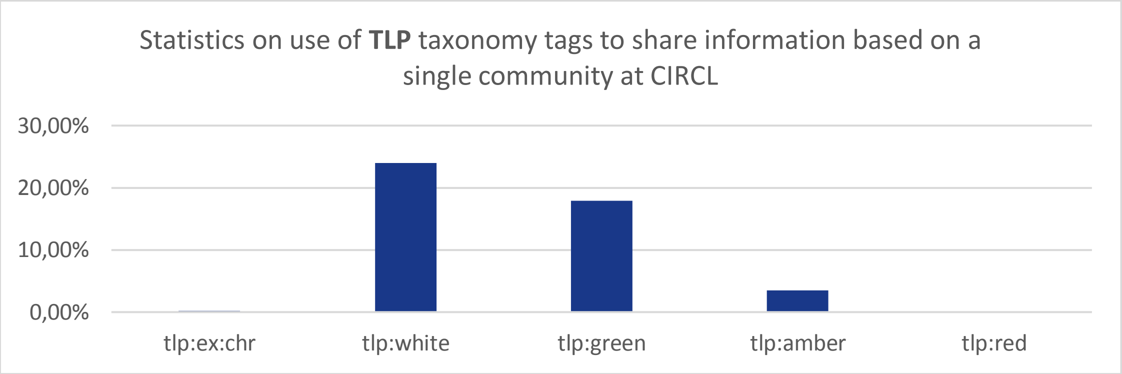

Figure 10 - Statistics of the TLP taxonomy tags used for sharing of information based on a single community at CIRCL (accounts for around 45.79% of the entire dataset).

From the above figure it is observed that most incidents or events shared have a low sensitivity.

The following are statistics regarding the tags used by the **eCSIRT** taxonomy for sharing based on a single community at CIRCL (accounts for 1.01% of the whole dataset):

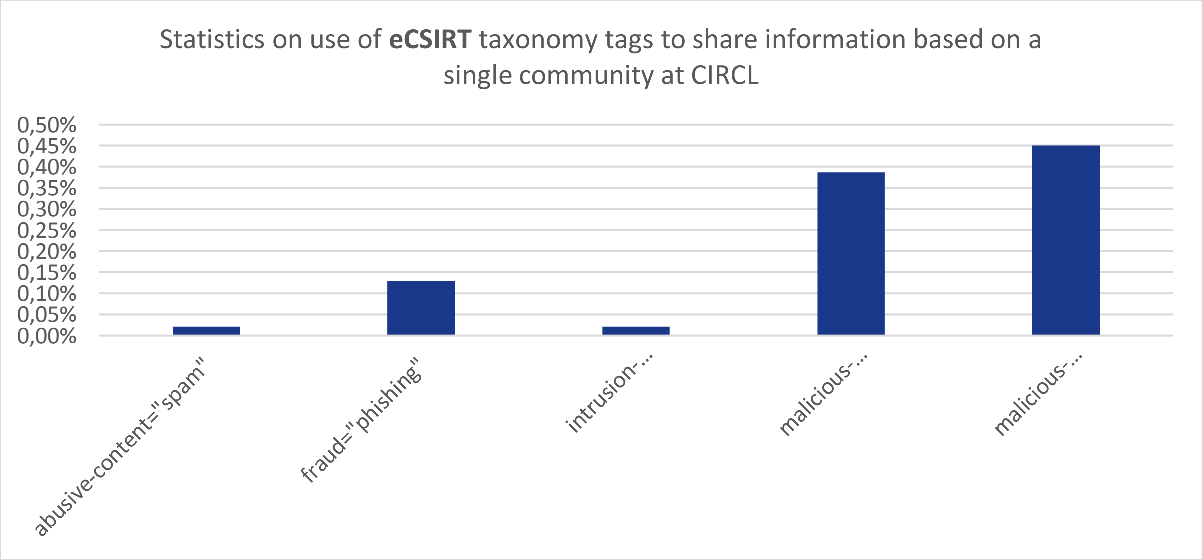

Figure 11 - Statistics of the eCSIRT taxonomy tags used for sharing of information based on a single community at CIRCL (accounts for around 1.01% of the entire dataset).

From the above figure it is observed that the terms “Spam”, “Phishing”, “Exploit”, “Malware”, and “Ransomware” are the most used.

The following are statistics regarding the tags used by the **Malware Classification (based on SANS)** taxonomy for sharing based on a single community at CIRCL (accounts for 1.59% of the whole dataset):

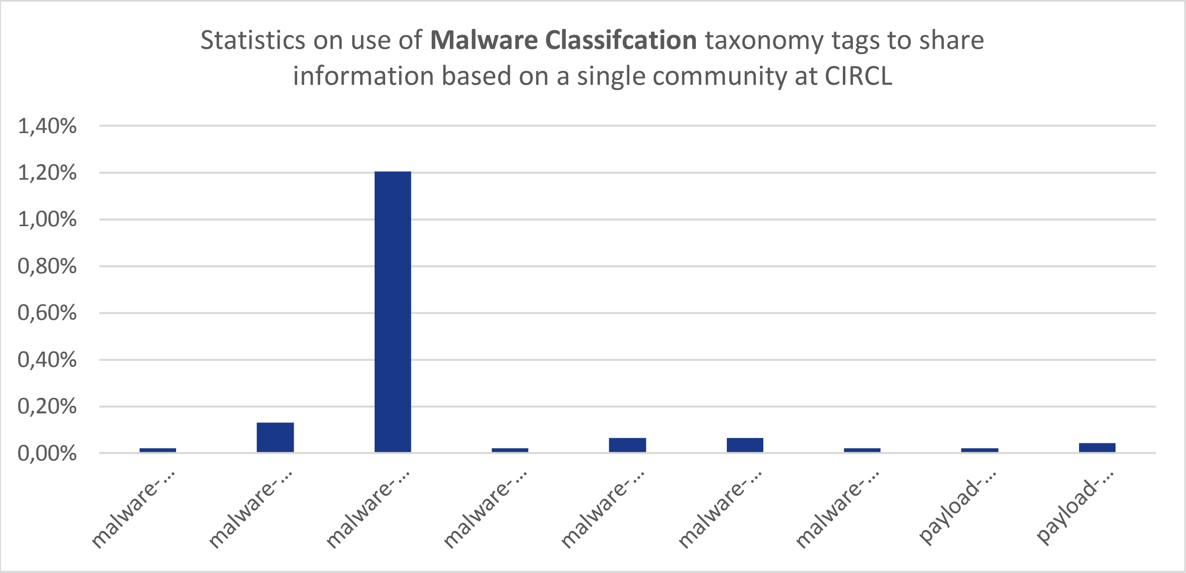

Figure 12 - Statistics of the Malware Classification (based on SANS) taxonomy tags used for sharing of information based on a single community at CIRCL (accounts for around 1.59% of the entire dataset).

From the above figure it is observed that the term “Ransomware” is the most used.

The following are statistics regarding the tags used by the **Veris** taxonomy for sharing based on a single community at CIRCL (accounts for 0.32% of the whole dataset):

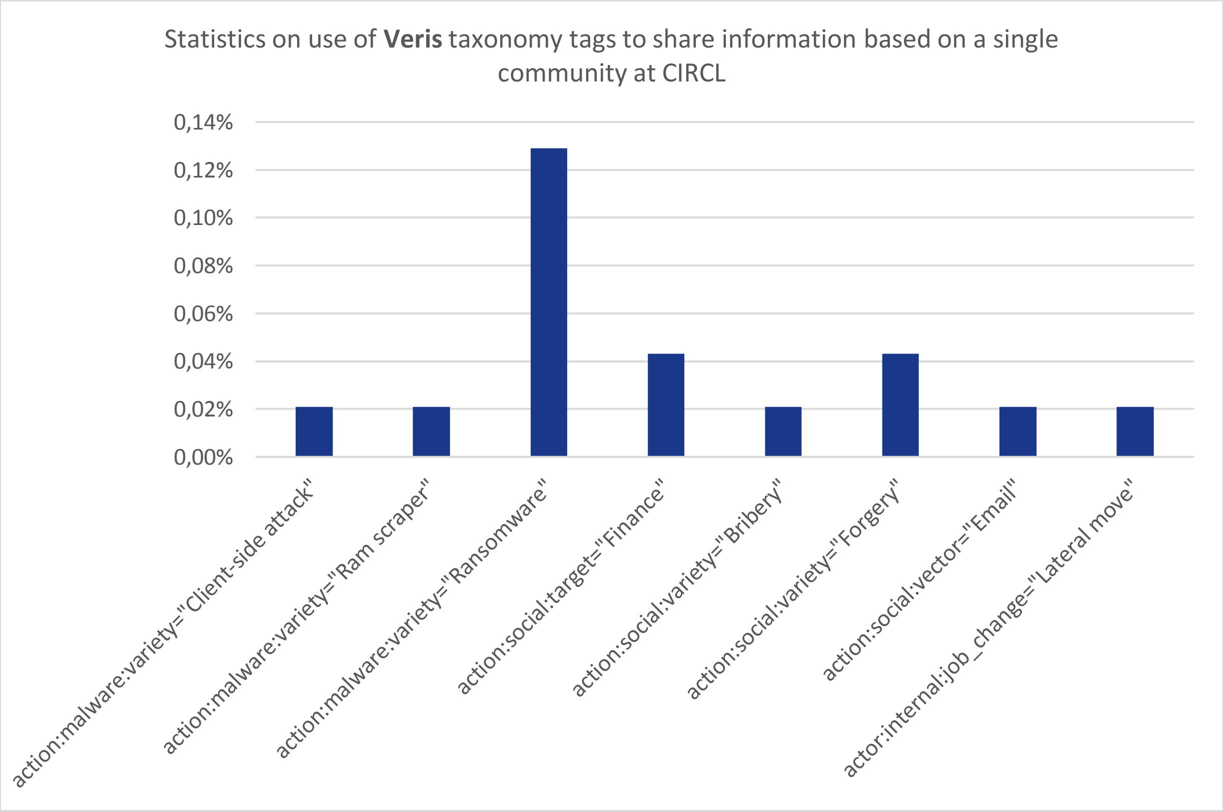

Figure 13 - Statistics of the Veris taxonomy tags used for sharing of information based on a single community at CIRCL (accounts for around 0.32% of the entire dataset).

From the above figure it is observed that the term “Ransomware” is the most used.

The following are statistics regarding the top 5 tags (accounts for 81.22% of the whole dataset):

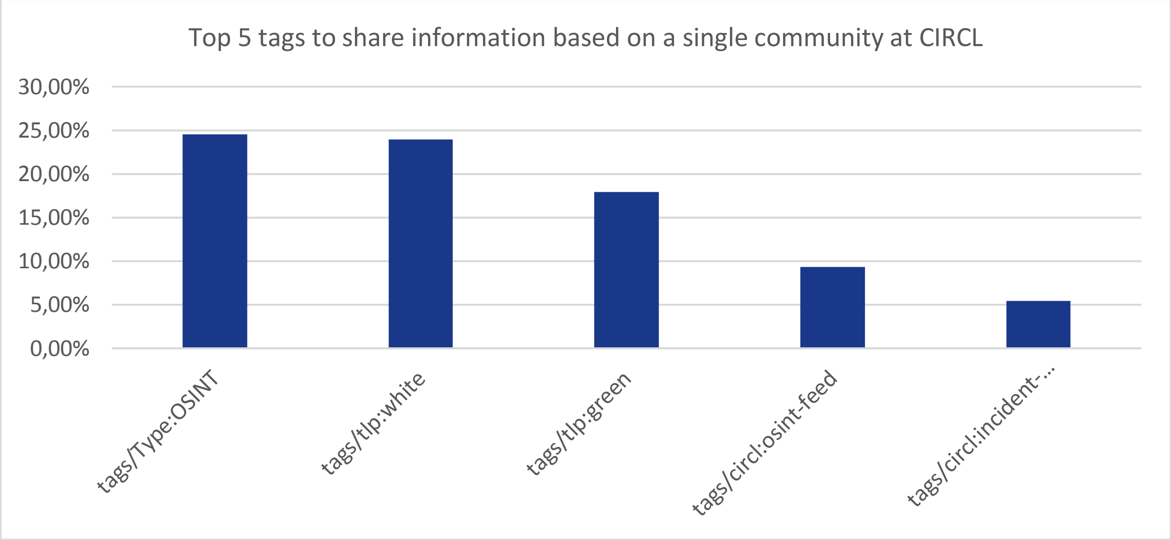

Figure 14 - Top 5 tags used for sharing of information based on a single community at CIRCL (accounts for around 81.22% of the entire dataset).

From the above figure it is observed that:

- TLP accounts for the majority of the dataset, due to the ease of use and applicability to almost every event or incident;
- The "Malware" tag is frequently used;
- The "OSINT" or Open-source intelligence is often used, indicating that information was gathered from publically available sources.

The following are statistics about the top 10 most common artefacts used for sharing:

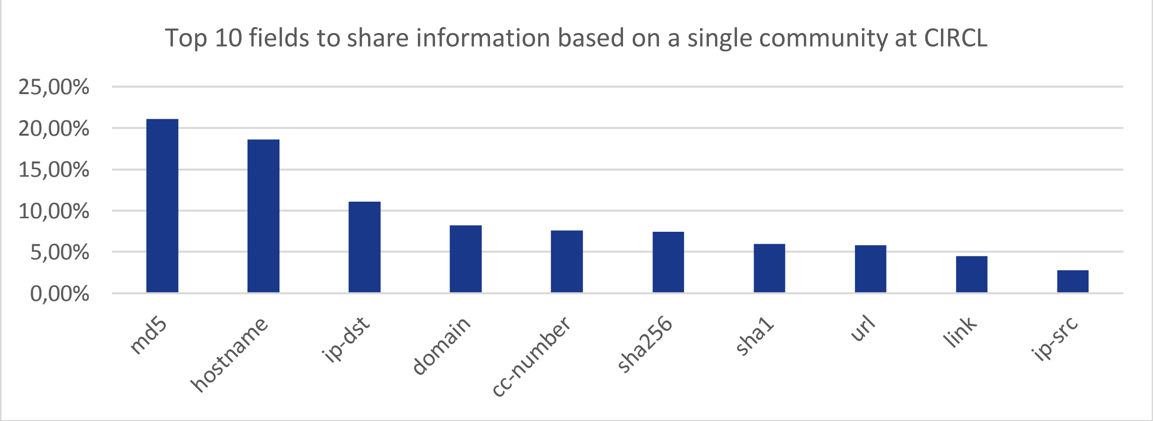

Figure 15 - Top 10 fields used for sharing of information based on a single community at CIRCL (accounts for around 93.03% of the entire dataset).

From the above figure it is observed that:

- Hashes are very commonly shared
- IPs, hostnames, domains, URLs, and links are very commonly shared
- Credit card numbers are commonly shared

The following are statistics about attribute types shared in the community:

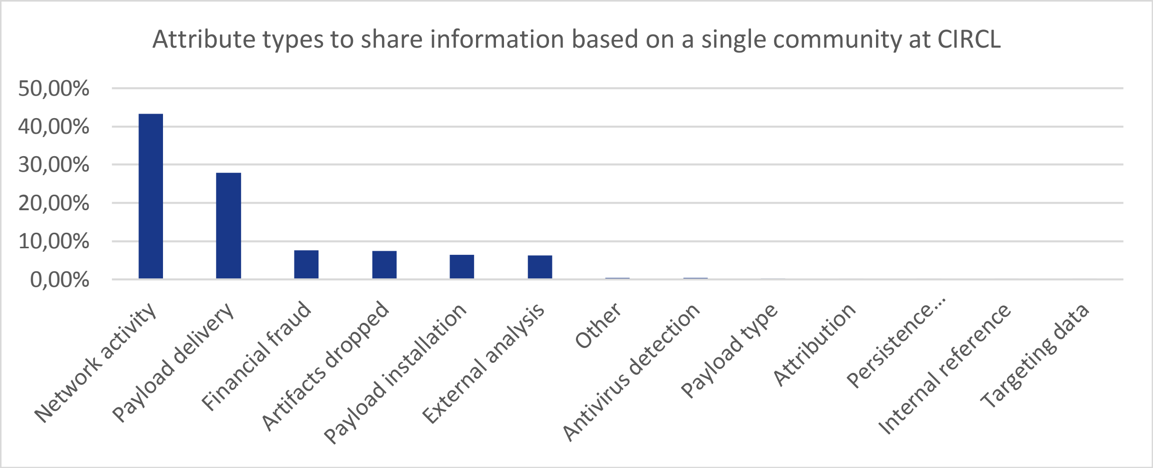

Figure 16 - Attribute types used for sharing of information based on a single community at CIRCL (accounts for around 100% of the entire dataset).

From the above figure it is observed that “Network activity” and “Payload delivery” are the most common

ENISA
European Union Agency for Network
and Information Security
Science and Technology Park of Crete (ITE)
Vassilika Vouton, 700 13, Heraklion, Greece

Athens Office
1 Vasilissis Sofias
Marousi 151 24, Athens, Greece

TP-06-16-306-EN-N

PO Box 1309, 710 01 Heraklion, Greece
Tel: +30 28 14 40 9710
info@enisa.europa.eu
www.enisa.europa.eu

ISBN: 978-92-9204-194-6
DOI: 10.2824/780536