EUROPEAN UNION AGENCY
FOR CYBERSECURITY


```
},
{
  "entry": [
    {
      "description": "Attacks that send requests to a system to discover weaknesses. This also includes testing processes to gather
      "expanded": "Scanning",
      "value": "scanner"
    },
    {
      "description": "Observing and recording of network traffic (wiretapping).",
      "expanded": "Sniffing",
      "value": "sniffing"
    },
    {
      "description": "Gathering information from a human being in a non-technical way (e.g. lies, tricks, bribes, or threats).",
      "expanded": "Social Engineering",
      "value": "social-engineering"
    }
  ],
  "predicate": "information-gathering"
},
{
  "entry": [
    {
      "description": "An attempt to compromise a system or to disrupt any service by exploiting vulnerabilities with a standardised
      "expanded": "Exploitation of known Vulnerabilities",
      "value": "ids-alert"
```


PROACTIVE DETECTION – MEASURES AND INFORMATION SOURCES

MAY 2020

ABOUT ENISA

The mission of the European Union Agency for Cybersecurity (ENISA) is to achieve a high common level of cybersecurity across the Union, by actively supporting Member States, Union institutions, bodies, offices and agencies in improving cybersecurity. We contribute to policy development and implementation, support capacity building and preparedness, facilitate operational cooperation at Union level, enhance the trustworthiness of ICT products, services and processes by rolling out cybersecurity certification schemes, enable knowledge sharing, research, innovation and awareness building, whilst developing cross-border communities. Our goal is to strengthen trust in the connected economy, boost resilience of the Union's infrastructure and services and keep our society cyber secure. More information about ENISA and its work can be found www.enisa.europa.eu.

AUTHORS

Piotr Białczak, Paweł Pawliński, Krzysztof Rydz, CERT Polska / NASK and Rossella Mattioli, ENISA

ACKNOWLEDGEMENTS

ENISA performed this study with the help of the contractor NASK and with the input from the members of the CSIRTs Network and other operational communities who contributed to this project. In particular we would like to thank the following persons for their input: Georgios Psykakos (CERT-EU), Marcin Dudek (CERT Polska), Michał Strzelczyk (CERT Polska). Finally, we would like to thank everyone that answered the survey: your input was crucial for this study.

LEGAL NOTICE

COPYRIGHT NOTICE

TABLE OF CONTENTS

EXECUTIVE SUMMARY

As of April 2020 there are more than 500 incident response teams in Europe[1]. These teams need every day to improve the prevention, detection and analysis of cyber threats and incidents. As envisioned by the NIS Directive[2] and in the Cybersecurity Act[3] ENISA is tasked with assisting the CSIRTs Network[4] and the Member States in improving the prevention, detection and capability to respond to cyber threats and incidents by providing them with knowledge and expertise. For these reasons ENISA aims with this study to provide an inventory of available methods, identify good practices and recommend possible areas for growth and attention to improve the proactive detection of network security incidents in EU.

In this respect, **proactive detection of incidents** is defined as **the process of discovery of malicious activity in a team's constituency through internal monitoring tools or external services that publish information about detected incidents, before the affected constituents become aware of the problem**. In 2011, ENISA published the first version of a study entitled “Proactive detection of network security incidents”[5]: The current project builds and expands on this. It aims to provide a complete inventory of all available methods, tools, activities and information sources for proactive detection of network security incidents, which are used already or potentially could be used by incident response teams in Europe nowadays.

The results of the 2019 survey and comparison with the 2011 edition have been already covered in the first document of this series already available on the ENISA website. **In the present document available methods, tools, activities and information sources** for proactive detection of network incidents were inventoried and evaluated. The inventory was based on desktop research and answers provided in the online survey, presented in the first part of the study. **Evaluation criteria** of measures for proactive detection of network incidents included **type, timeliness, accuracy, ease of use, coverage, resources, scalability, extensibility and completeness.** Information sources for proactive detection of network incidents were evaluated using criteria such as timeliness, accuracy, ease of use, data volume and completeness.

[1] ENISA CSIRTs by Country - Interactive Map https://www.enisa.europa.eu/csirts-map
[2] https://ec.europa.eu/digital-single-market/en/network-and-information-security-nis-directive
[3] https://ec.europa.eu/digital-single-market/en/eu-cybersecurity-act
[4] www.csirtsnetwork.eu/
[5] https://www.enisa.europa.eu/topics/csirt-cert-services/proactive-services/proactive-detection/proactive-detection-of-incidents

1. INTRODUCTION

In 2011, ENISA published the study entitled “Proactive detection of network security incidents” [6] and in 2019, with this study the aim is to understand what has changed in the last eight years and map the current situation among incident response teams in Europe. The objectives are to provide an inventory of available methods, identify good practices and recommend possible areas for growth and attention to improve the detection of network security incidents in EU.

Throughout this study, as in the 2011 study, **proactive detection of incidents** is defined as the **process of discovery of malicious activity in a team's constituency through internal monitoring tools or external services that publish information about detected incidents, before the affected constituents become aware of the problem.**

1.1 CONTEXT OF THE WORK

For more than fifteen years ENISA has been supporting Member States and CSIRT communities to build and advance their CSIRT capabilities. Individual teams which represent different sectors and businesses, as well as existing CSIRT communities, are indispensable elements of this shared responsibility and endeavour.

ENISA’s incident response support portfolio of work is related to setting up, running and developing capabilities of Computer Security Incident Response Teams (CSIRTs) in Europe. There are currently more than 500 CSIRTs listed in the ENISA Inventory[7]. The goal is to identify common practices across the EU to improve operational cooperation and information exchange. The primary audience are the CSIRTs Network[8] members, their leadership and the incident response community at large.

The NIS Directive[9] in Article 12 establishes the CSIRTs Network[10] “to contribute to developing confidence and trust between the Member States and to promote swift and effective operational cooperation”. The CSIRTs Network is a network composed of EU Member States’ appointed CSIRTs and CERT-EU[11] (“CSIRTs Network members”). ENISA is tasked to actively support the CSIRTs cooperation, provide the secretariat and active support for incident coordination upon request.

Moreover, with the EU Cybersecurity Act, ENISA is also mandated to increase operational cooperation at EU level and asked in Article 6 “Capacity-building” to assist Member States in their efforts to improve the prevention, detection and analysis of cyber threats and incidents and Article 7 “Operational cooperation at Union level” in advising on how to improve their capabilities to prevent, detect and respond to incidents.

In 2011, ENISA published the first version of “Proactive detection of network security incidents”[12]: The current project builds upon this study and aims to provide a complete inventory of all available methods, tools, activities and information sources (hereafter ‘measures’) for

[6] https://www.enisa.europa.eu/topics/csirt-cert-services/proactive-services/proactive-detection/proactive-detection-of-incidents
[7] https://www.enisa.europa.eu/topics/csirts-in-europe/csirt-inventory/certs-by-country-interactive-map
[8] https://csirtsnetwork.eu/
[9] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2016.194.01.0001.01.ENG&toc=OJ:L:2016:194:TOC
[10] http://www.csirtnetwork.eu/
11 CERT-EU is a Computer Emergency Response Team or CSIRT and its constituency is composed of all the EU Institutions, Agencies and Bodies. Its offices are in Brussels.
[12] https://www.enisa.europa.eu/topics/csirt-cert-services/proactive-services/proactive-detection/proactive-detection-of-incidents

proactive detection of network security incidents, which are used already or potentially could be used by incident response teams in Europe nowadays.

1.2 OBJECTIVES OF THE STUDY

The objectives of this project are to:

- provide an inventory of available methods, tools, activities and information sources for proactive detection of network incidents,
- identify good practices and recommend possible areas for growth with attention for new and already established incident response teams in Europe
- draft a list of key recommendations for policy makers in order to improve the detection of network security incidents in EU.

Figure 1: Information sources and measures covered by the study

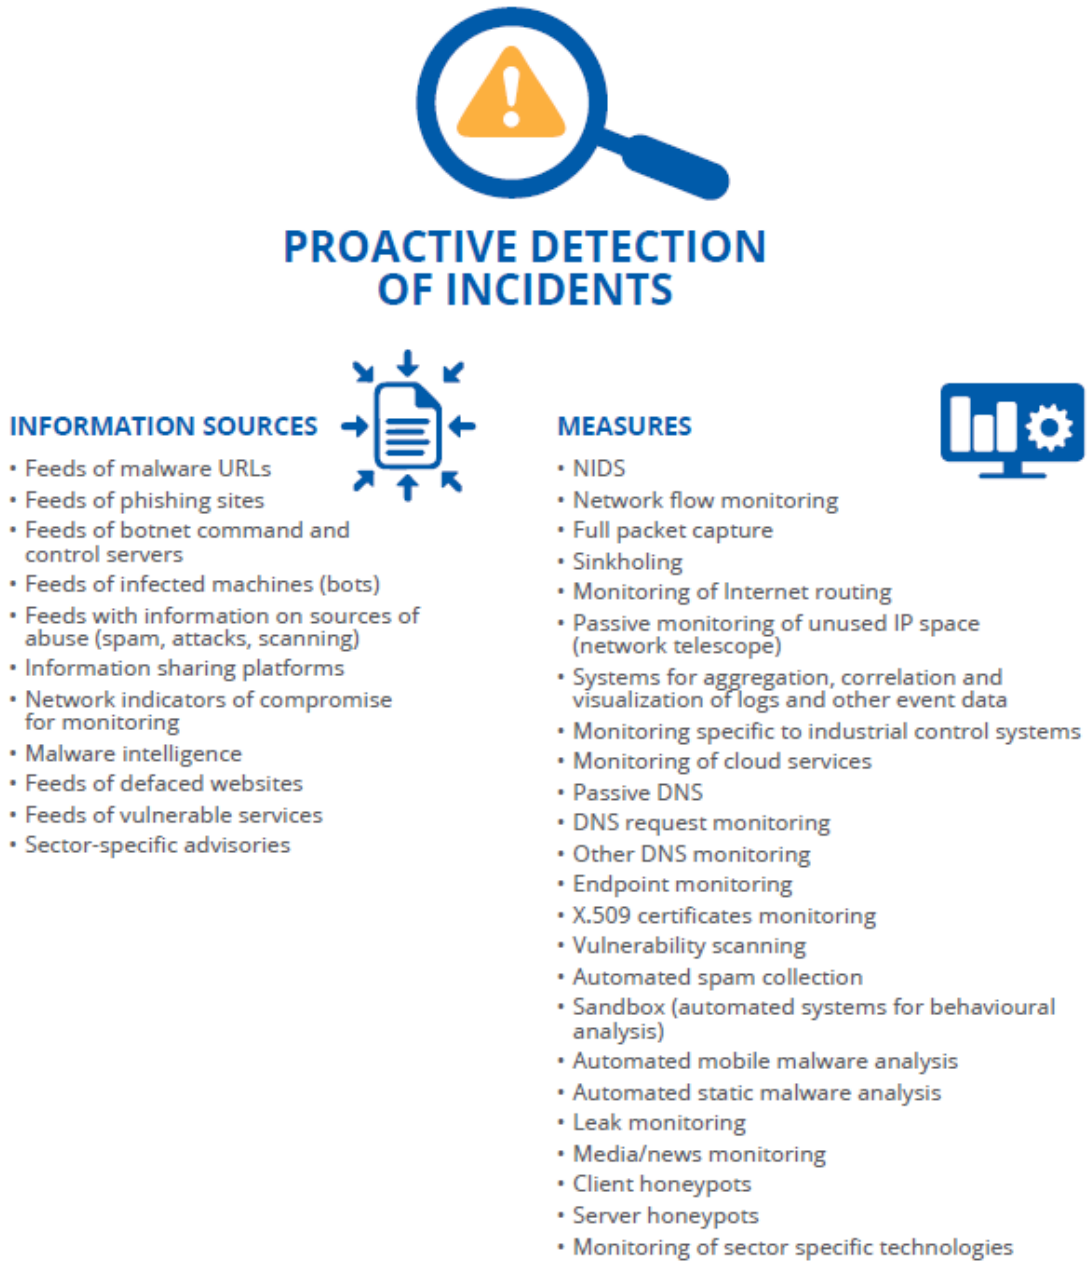

The results of this project are provided in the three parts. The **first part** contained the

- survey among incident response teams in Europe
- comparison with 2011 survey

The **second part**, the current document, covers:

- inventory of available methods, tools, activities and information sources for proactive detection of network incidents
- evaluation of identified measures and information sources

The **third part** covered:

- analysis of gathered data
- recommendations for policy makers in order to improve the detection of network security incidents in EU

Furthermore, the **current project has two formats: one is the present document which gives an overview of the findings and the other is a living document hosted on GitHub**[13] which aims to represent a point of reference to identify or reassess appropriate measures for proactive detection of incidents for new or well-established teams.

1.3 DEFINITIONS

1.3.1 Proactive versus reactive detection of incidents

As stated in the introduction and as previously used in the 2011 study, proactive detection of incidents is meant as a **process of discovery of malicious activity in a CSIRT team's constituency, before the affected constituents become aware of the problem.** On the other hand, when a CSIRT team receives an incident report, its role is only reactive - to respond accordingly to the report. In such perspective, a proactive approach can help in detection of incidents at an early stage of the attack or even before it happens.

1.3.2 Measure versus information source

In this study, "measure" is defined as a set of systems, tools and technologies deployed and used by CSIRT teams to provide information about features of a monitored network**.** Whereas "**information source**" is defined as **a source of data independent of the system producing it and consumed using its own, abstract method** as in the 2011 study. The main difference between these two categories is that tools and systems constituting measures have to be deployed and maintained in order to provide information, while the information source is provided as a service by other entity.

1.4 PREVIOUS ENISA WORK ON THE TOPIC

Since 2005, ENISA has been supporting Member States and CSIRT communities in the EU to build and advance their incident response capabilities with handbooks, online & onsite trainings and dedicated projects[14]. ENISA's portfolio of work is related to setting up, running and developing capabilities of Computer Security Incident Response Teams (CSIRTs). The goal is to identify common practices across the Union to improve operational cooperation,

[13] https://github.com/enisaeu/IRtools
[14] https://www.enisa.europa.eu/topics/csirts-in-europe

preparedness and information exchange for the next generation of cyber-attacks. More info can be found at **https://www.enisa.europa.eu/csirt-services**

Relevant ENISA deliverables and activities comprise:

- Orchestration of CSIRT Tools[15]
- Reference Security Incident Taxonomy Working Group[16]
- Exploring the opportunities and limitations of current Threat Intelligence Platforms[17]
- Actionable Information for Security Incident Response[18]
- Standards and tools for exchange and processing of actionable information[19]
- Detect Share Protect - Solutions for Improving Threat Data Exchange[20]
- Proactive Detection of Network Security Incidents – Honeypots[21]
- Proactive Detection of Network Security Incidents – Data feeds – internal and external[22]

Moreover, the following relevant trainings are also available on ENISA website:

- Proactive incident detection: handbook and VM[23]
- Automation in incident handling: handbook and VM[24]
- Honeypots: handbook and VM[25]
- Presenting, correlating and filtering various feeds: handbook and 2 VMs[26]

[15] https://www.enisa.europa.eu/topics/trainings-for-cybersecurity-specialists/online-training-material/technical-operational
[16] Reference Security Incident Taxonomy Working Group - RSIT- WG https://github.com/enisaeu/Reference-Security-Incident-Taxonomy-Task-Force
[17] ENISA, "Exploring the opportunities and limitations of current Threat Intelligence Platforms", 2018, https://www.enisa.europa.eu/publications/exploring-the-opportunities-and-limitations-of-current-threat-intelligence-platforms
[18] ENISA, "Actionable Information for Security Incident Response", 2015, https://www.enisa.europa.eu/publications/actionable-information-for-security
[19] ENISA "Standards and tools for exchange and processing of actionable information" https://www.enisa.europa.eu/publications/standards-and-tools-for-exchange-and-processing-of-actionable-information
[20] ENISA, "Detect Share Protect - Solutions for Improving Threat Data Exchange", 2013, https://www.enisa.europa.eu/publications/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs
[21] ENISA, "Proactive Detection of Network Security Incidents – Honeypots", 2012, https://www.enisa.europa.eu/publications/proactive-detection-of-security-incidents-II-honeypots
[22] ENISA, "Proactive Detection of Network Security Incidents – Data feeds", 2011, https://www.enisa.europa.eu/publications/proactive-detection-report
[23] ENISA, "Proactive incident detection training", https://www.enisa.europa.eu/topics/trainings-for-cybersecurity-specialists/online-training-material/technical-operational#proactive-incident-detection
[24] ENISA, "Automation in incident handling training", https://www.enisa.europa.eu/topics/trainings-for-cybersecurity-specialists/online-training-material/technical-operational#automation_incident
[25] ENISA, "Honeypots training", https://www.enisa.europa.eu/topics/trainings-for-cybersecurity-specialists/online-training-material/technical-operational#honeypots
[26] ENISA, "Presenting, correlating and filtering various feeds training",https://www.enisa.europa.eu/topics/trainings-for-cybersecurity-specialists/online-training-material/technical-operational#presenting--correlating-and-filtering-various-feeds

1.5 METHODOLOGY

This section describes the methodology used in different parts of the analysis.

Figure 2: Methodology

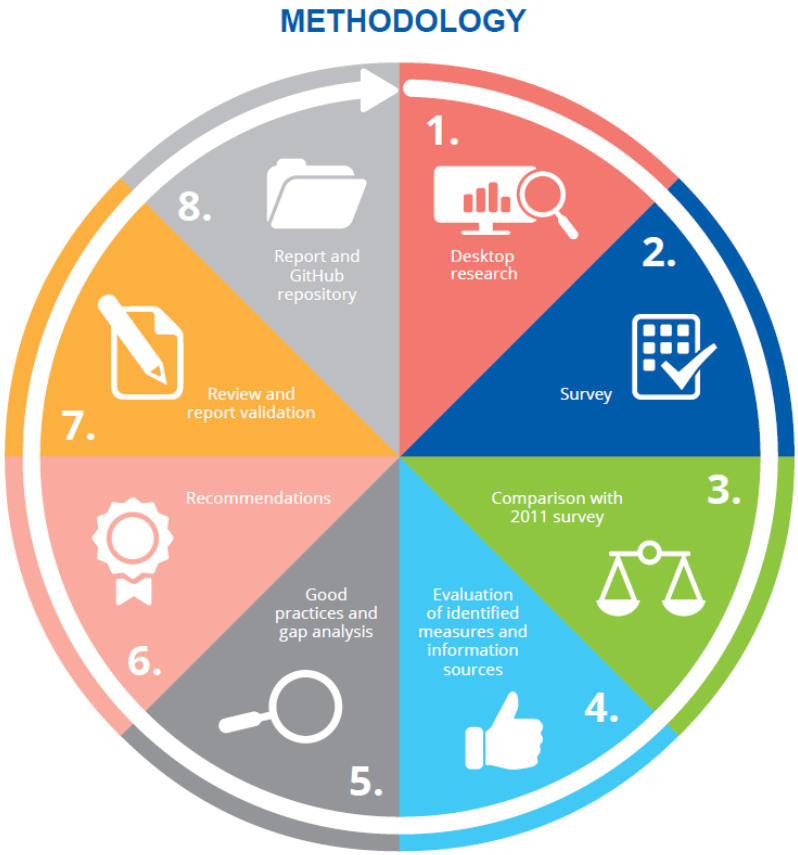

- Phase 1, 2 and 3 are detailed in “Proactive detection – Survey results”.
- Phase 4 detailed is below.
- Phase 5 and 6 are detailed in “Proactive detection - Gap analysis good practice and recommendations”.
- Phase 6 was performed collecting the input of the CSIRTs Network, the experts mentioned in the acknowledgements and via ENISA content approval workflow.
- Phase 7 is the publication on the ENISA website and GitHub repository.

1.5.1 Phase 4

In this part of the project, different knowledge sources were reviewed in order to provide an initial list of measures for proactive detection of network incidents. Particular tools and information sources were grouped into categories to give a more general overview independent of single tools. The goal was also to focus on the most crucial features, helping in proactive detection, provided by such a measures or type of tools.

For all categories open source examples have been identified and provided. The lists excluded ticketing/incident handling tools (such as Request Tracker[27]), information sharing platforms (such as MISP[28]) and forensics tools.

The above mentioned tools are covered to a different extent by other ENISA activities and deliverables mentioned in section 1.4.

1.5.1.1 Measures

The measures for proactive detection of network incidents were identified using the 2011 study, lists of tools and systems generally called “awesome lists” and other public sources. They were also extended using the practice of project team. The initial list of measures (including examples of tools) was analysed in order to reject outdated elements and to cover all aspects of proactive detection of network incidents, giving a state-of-the-art overview.

1.5.1.2 Information sources

The information sources for proactive detection of network incidents were identified using public sources of information (similarly as with the measures described above) and extended with sources used in CERT Polska’s n6 platform[29]. Particular sources were reviewed to reject outdated elements, then grouped in order to highlight similarity of their functionalities.

1.5.2 Evaluation of identified measures and information sources

1.5.2.1 Analysis of the measures

Identified measures were analysed and evaluated in the context of a typical deployment. For some measures there are several tools (primarily open source) or services provided as examples, however they were not evaluated individually. This allows for a more concise summary of expected benefits and challenges of entire measures, but also means that differences - sometimes major - between individual tools have not been described. The evaluation focused on the following nine criteria for measures and five for external information sources. The criteria were rated using a simple 4-level qualitative scale: poor, fair, good and excellent. The only exceptions to this scheme are **type** and **coverage** criteria for measures and **data volume** for information sources.

In the 2011 report, the same scale was used in the evaluation of individual sources and tools. The scope of this study does not allow for such detailed analysis and the evaluation is performed on the level of the whole measures and categories. It is important to keep in mind that measures are summarising general techniques that are implemented by multiple tools, which often have significant variance of features and capabilities. Thus, the rating is based on the characteristics of typical deployments in CSIRTs and similar entities and is meant to provide only a simplified summary of a particular aspect the given measure. The same principle applies to the evaluation of categories of information sources.

Evaluation criteria:

Type

How the measure contributes to proactive detection: does it generate alerts that can indicate an ongoing incident or a potential threat. Alternatively, does it support analysts during an investigation by providing relevant information.

[27] https://bestpractical.com/request-tracker
[28] https://www.misp-project.org/
[29] https://n6.cert.pl

Timeliness

Timeliness determines if the information collected using a particular measure is recent enough that it can be used for detecting a threat early, before the intrusion is well advanced.

- **Poor:** Information provided by the measure is not recent enough to be effective for proactive detection.
- **Excellent:**: The measure provide current information about threats with very low latency.

Accuracy

Accuracy evaluates the quality of the provided information, taking into account the number of false alarms and undetected incidents.

- **Poor:** The number of potential false positives means that all results need to be verified manually.
- **Excellent:** Analysts can be certain that the tools provide correct results.

Ease of use

Ease of use evaluates the level of resources needed for data acquisition and processing. This includes the level of analysts' experience, but also the complexity of software needed for using the measure.

- **Poor:** Analysts need to spend much time on learning how to properly use the tools.
- **Excellent:** Tools are intuitive and can be used without much specialized knowledge.

Coverage

Coverage determines the scope of the collected information and what kinds of threats can be detected.

Resources

Resources evaluates level of resources needed for operating the measure. It includes technical (hardware and software) resources, but also human resources (analysts' time and their level of experience).

- **Poor:** There is substantial cost (financial and time) in implementing and maintaining the measure.
- **Excellent:** Tools are cheap and do not require much effort to deploy.

Scalability

Scalability evaluates measure's capability to handle growing volume of data or growing network range.

- **Poor:** Measure is costly to scale to monitor more threats.
- **Excellent:** Scaling is straightforward and cheap.

Extensibility

Extensibility evaluates measure's capability to extend its basic functionality. It can be formulated by technical feasibility, but also resources needed for the extension (time needed for the extension, required level of technical experience).

- **Poor:** Tools are difficult to customize.
- **Excellent:** It is easy to adapt tools to local requirements, for example through plugins, flexible configuration, webhooks, etc.

Completeness
The level of technical detail of the information that is obtained thanks to the measure. If the measure does not provide crucial details, it means it is difficult to use for proactive detection.

- **Poor:** The measure does not provide a sufficient level of detail for effective proactive detection.
- **Excellent:** Analysts have convenient access to all relevant details to detect a threat.

Examples of tools and services for each measure are provided on the GitHub repository[30].

1.5.2.2 Analysis of the external information sources

Similarly to tools, external information sources were grouped into several categories. Each category contains sources that provide information of the same type or information that in general can support proactive detection in the same manner, for example feeds of phishing URLs.

The following five criteria were used for the characterisation of the entire categories. Individual sources were not evaluated, but operational experience, desktop research and survey results were used to describe the typical properties of sources in the category. With such an approach it is possible to summarise a large range of information providers, however individual sources can sometimes significantly differ from each other in various aspects.

In the same way as in the evaluation of measures, the criteria do not follow any quantitative or qualitative scale.

Timeliness
Whether the source provides information that is recent enough to be useful for proactive detection.

- **Poor:** Information provided by the sources is not recent enough to be effective for proactive detection.
- **Excellent:** The sources provide current information about threats with low delay.

Accuracy
Trustworthiness of the source, especially with regard to false positives. False positives in this context means that the team would see false alarms if the information from the source is used for monitoring.

- **Poor:** The number of potential false positives means that all results need to be verified manually.
- **Excellent:** In most cases, analysts can be certain that the information is correct.

Ease of use
Whether the team needs much expertise, effort or other resources to integrate the source with its internal workflows and tooling.

- **Poor:** Analysts need to spend much time to learn how to interpret and use information provided by these sources.
- **Excellent:** Using information provided by the sources for proactive detection is easy and does not require specialized knowledge.

[30] https://github.com/enisaeu/IRtools

Data volume
Expected amount of data provided, relative to other categories.

Completeness
Does the level of technical detail is sufficient for early detection purposes.

- **Poor:** The sources do not provide a sufficient level of detail for effective proactive detection.
- **Excellent:** Information provided by the sources contains all relevant details.

For each category, several examples of sources are provided on the GitHub repository[31].

[31] https://github.com/enisaeu/IRtools

2. EVALUATION OF IDENTIFIED MEASURES AND INFORMATION SOURCES

2.1 MEASURES

This chapter gathers the analysis of the measures identified during the desktop research based on the evaluation criteria . For each measure an overview is provided first, then the result of the evaluation and in the relevant GitHub paragraph some examples of available solutions best known at the time of writing. The full list of available solutions best known at the last update is available on GitHub.

2.1.1 NIDS[32]

Network intrusion detection system (NIDS) are used to monitor network traffic in order to search for evidences of malicious operations in observed networks. A basic deployment of a NIDS system consists of a probe which monitors network traffic and system for collection of logs/alerts. Depending on size of network and bandwidth, the system can depend on a single probe (small networks) or be more complex, with multiple probes and multiple collection systems. NIDS systems are fed with rules describing suspicious behaviour, for which an alert should be issued. The rules can be written by analysts, but they are also available as sets, either open sourced (free) or commercial (paid). The messages alerted by a NIDS system have to be monitored by an analyst, who can decide about the significance and priority of information provided by the system. To help in the analyst's work, logs from a NIDS system can be forwarded to aggregation, correlation and visualization systems, including SIEM (Security Information and Event Management) systems. Usage scenarios of NIDS systems by CSIRTs depend on the team's constituency. NIDS systems can be used for monitoring internal networks, but also for monitoring research/laboratory networks, for example networks of sandbox systems.

2.1.1.1 Evaluation

Type: Alerts.

Timeliness: Excellent; near real-time.

Accuracy: Poor; quality of alerts depends on the rules used; typically verification is needed due to common false positives.

Ease of use: Excellent; alerts are usually easy to interpret and tools provide a convenient way to browse the results.

Coverage: Monitoring of local infrastructure; coverage of threats depends on the rules used.

Resources: Fair; analysts need to verify alerts, which can be large in number.

Scalability: Good; scales with the number of sensors.

[32] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#nids

Extensibility: Excellent; all mainstream tools support rule-based configuration; typically tools provide plugin mechanisms.

Completeness: Fair; there is limited information about the type of threat detected and basic network information.

2.1.2 Network flow monitoring[33]

Network flow monitoring systems provide means for extraction of network flow information from network traffic. Some of the systems also help in basic analysis of network flows, including bandwidth level, protocol usage and IP addresses involved in communication. Network flows can be also used as input in specialised network anomaly detection tools.

2.1.2.1 Evaluation

Type: Alerts, support.

Timeliness: Excellent; alerts on blacklist hits are near real-time; some detection methods may work over longer time periods, for example anomaly detection.

Accuracy: Poor; alerts require verification.

Ease of use: Good; mature specialized GUI and CLI tools with query capabilities, dashboards.

Coverage: Monitoring of local infrastructure, either just external connections or inter-network traffic as well; the biggest issue is that no payloads are saved.

Resources: Fair; analysts may have a lot of alerts to investigate; additionally the flow monitoring solutions are often custom-built from ready-made components such as flow collectors, message brokers and databases with appropriate orchestration, which requires some extra resources to maintain.

Scalability: Good; scales with the number of sensors; for large networks the cost of the backend may be significant.

Extensibility: N/A; varies, tool-dependent.

Completeness: Poor; missing payloads means that further correlation is necessary; anomaly detection may provide limited context.

2.1.3 Full packet capture[34]

Full packet capture systems are used to provide means for archiving network traffic, what can be later used for precise analysis by analysts or automatic systems. Such systems can simply save network traffic to preferred file format, but some of them also provide tools for exploration and basic analysis of the network traffic.

2.1.3.1 Evaluation

Type: Support.

Timeliness: Excellent; data can be collected in real-time.

[33] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#network-flow-monitoring
[34] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#full-packet-capture

Accuracy: Good; actual contents of the network traffic; protocol analysis can fail sometimes; attackers can spoof origin of packets.

Ease of use: Good; requires some expertise.

Coverage: Monitoring of local infrastructure; the biggest problem is lack of payload for encrypted connections unless TLS inspection is deployed.

Resources: Good; systems do not require much maintenance after the initial setup

Scalability: Fair; scales with the number of sensors; for medium and large networks the amount of data collected will be a challenge and the hardware costs of the backend will be significant, this can be alleviated by keeping data for shorter periods.

Extensibility: Good; details depend on the tool.

Completeness: Good; except for encrypted or obfuscated payloads.

2.1.4 Sinkholing[35]

Sinkhole systems are be used to discover malware infections by monitoring host connections. After ceasing C&C server address, sinkhole server can be used as replacement to the original botnet's infrastructure, and track all connections made by bots. With that data, analysts can provide information about the number of infections, geographical distribution, most impacted networks etc.

2.1.4.1 Evaluation

Type: Alerts.

Timeliness: Excellent; real-time.

Accuracy: Excellent; assuming local deployment, otherwise many Internet scans will be registered, which can significantly lower the accuracy.

Ease of use: Good; output is easy to work with.

Coverage: Local deployment: monitoring of own infrastructure for known threats; global deployment: monitoring of victims of specific botnet, RAT, etc.; global sinkholing is useful primary for notifying other entities.

Resources: Excellent.

Scalability: Excellent; typical solutions for load balancing can be employed; for global botnet sinkholing hardware investments may be required for storage.

Extensibility: N/A; typically custom solutions are deployed.

Completeness: Good; can be excellent if bots can be identified, for example using HTTP headers.

[35] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#sinkholing

2.1.5 Monitoring of Internet routing[36]

Monitoring of Internet routing can provide information about status of routing paths and thus be used to detect attacks, for example BGP protocol hijacking.

2.1.5.1 Evaluation

Type: Alerts.

Timeliness: Excellent; real-time.

Accuracy: Fair; alerts need verification.

Ease of use: Fair; requires expertise to operate; COTS services are easy for general use.

Coverage: Visibility into global events; alerts are generated for the infrastructure of interest only.

Resources: Good.

Scalability: Fair; multiple sources of the BGP data can be used however there are diminishing results of adding new ones.

Extensibility: Fair; tool-dependent.

Completeness: Good; typically historical data is available.

2.1.6 Passive monitoring of unused IP space (network telescope)[37]

Passive monitoring of unused IP space (also known as network telescope) can help in identifying network attacks. As the monitored IP addresses are unassigned, no network traffic should be directed on them. From such perspective, any packets observed at these addresses are usually sent by victims of reflected Denial of Service attacks or automatic systems scanning the Internet, for example to find vulnerable hosts or to exploit vulnerable services.

2.1.6.1 Evaluation

Type: Alerts, support.

Timeliness: Good; real-time for some types of alerts; analyses can be run periodically, which may mean that the results are available only after some hours.

Accuracy: Poor; this method will susceptible to spoofing; relevant traffic is mixed with packets resulting from harmless misconfiguration or of unknown purpose.

Ease of use: Poor; except for predefined alerts generated by network telescopes, handling of the collected data is time-intensive and requires expertise.

Coverage: The traffic collected corresponds to various events world-wide; this measure is used primarily to understand global threats, not particular networks of interest.

Resources: Poor; no Commercial off-the-shelf (COTS) solutions, analysis is time-intensive.

[36] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#monitoring-of-internet-routing
[37] https://github.com/pp-/proactive-inventory/blob/master/measures_and_tools.md#passive-monitoring-of-unused-ip-space-network-telescopedarknet

Scalability: Good; the quality of the information will improve with the increase of the monitored address space, however there will be diminishing results with each IP; depending on the size of the monitored address space and how much detail is stored, the backend for processing and storage may require non-negligible investment.

Extensibility: N/A; typically custom solutions are deployed.

Completeness: Fair; for some predefined events the level of detail is satisfactory; often further research is required to understand the nature of traffic.

2.1.7 Systems for aggregation, correlation and visualization of logs and other event data[38]

Systems for aggregation, correlation and visualization of logs and other event data is an umbrella category grouping many systems. They gather big amounts of data from logging/monitoring systems and process them in order to help analysts monitor the infrastructure. Depending on type of system, they aggregate data, correlate them and visualize to present information that is most crucial to analysts. An example of a system type that provides all of these tasks is Security Information and Event Management (SIEM). SIEM can gather data from network monitoring systems like NIDS or endpoint monitoring systems, providing analysts with means for monitoring and inspection of defended infrastructure. Such measure is excellent for the purpose of fusion of all security monitoring data in the organisation.

2.1.7.1 Evaluation

Type: Alerts, support.

Timeliness: Good; data can be ingested in real-time but in the end depend on the timeliness of the input; some analyses can be run periodically (for example daily queries).

Accuracy: Good; depends on proper configuration and prepared queries.

Ease of use: Fair; mature user interfaces and APIs with a lot of functionality; it may be difficult to master these tools, however typical tasks have a moderate learning curve.

Coverage: All information relevant to the monitored infrastructure; the actual coverage depends on inputs.

Resources: Poor; to achieve good results, substantial time is needed for configuration, integration of data sources, preparing of queries and dashboards.

Scalability: Fair; commercial solutions are expensive to scale; significant hardware investment may be required for storing large amount of logs; dedicated staff can be needed to keep the systems operational at a certain scale.

Extensibility: Excellent; multiple ways to interact with the systems through APIs and various plugin mechanisms.

Completeness: Excellent; assuming sufficient data inputs have been configured.

[38] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#systems-for-aggregation-correlation-and-visualization-of-logs-and-other-event-data

2.1.8 Monitoring specific to industrial control systems[39]

Monitoring systems specific to industrial control systems (ICS/SCADA) is often similar to regular network monitoring, however industrial communication protocols are much different from a IT networks. The fact that ICS often have certification requirements and are isolated from IT networks poses some challenges with developing good monitoring capabilities. This is one of the reasons that ICS monitoring is dominated by commercial vendors and open source solutions are much less used. The evaluation below covers mainstream commercial offerings only; examples refer to selected open source tools but they have a different scope in functionality.

2.1.8.1 Evaluation

Type: Alerts, support.

Timeliness: Excellent; real-time.

Accuracy: Good; detection is often based on anomaly detection: this approach applied to ICS yield better results compared to IT networks, since OT (Operational Technology) environments typically do not undergo unplanned changes.

Ease of use: Excellent; monitoring tools are designed to be accessible by ICS engineers.

Coverage: Monitoring of local infrastructure; network traffic is typically obtained through taps; the biggest issue is presence of uncommon or custom protocols that are not supported by the monitoring tools, which means that it is not possible to inspect commands and values.

Resources: Poor; deployment of ICS monitoring is often time consuming and costly; a lot of complexity is caused by certification requirements; once installed ongoing effort for analysts should not be significant, as there is not a large number of alerts to investigate.

Scalability: Good.

Extensibility: Fair; tools support rule-based configuration; any major customisations are typically not possible without involving a vendor.

Completeness: Good; for supported protocols, full visibility into control commands and process values.

2.1.9 Monitoring of cloud services[40]

Adoption of cloud services is increasing both for enterprise and governmental institutions. Nowadays, major providers offer much more rich and complex set of services than just hosting virtual machines and having a complete understanding of the protected assets in such environment might be difficult. However, the primary concern for defence is that the infrastructure cannot be monitored directly, which makes detection of suspicious activities a challenge. Major cloud providers offer in-house specialised security tools that allow to collect and analyse logs and monitor network traffic. Capabilities of such tools vary significantly between providers: in few cases their features address teams' need for monitoring and detection. Majority of providers, especially medium and small ones, do not have sufficient offering in this regard. The evaluation below is based on commercial in-house solutions from two cloud providers and focus on monitoring of virtual machines and no other cloud services.

[39] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#monitoring-specific-to-industrial-control-systems
[40] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#monitoring-of-cloud-services

Smaller providers and standalone open source tools have not been taken into consideration, as their capabilities is too far limited and making a generalisation impossible.

2.1.9.1 Evaluation

Type: Alerts, support.

Timeliness: Good; several minutes of delay.

Accuracy: Good; for vendor-provided images and standard software there should be a small number of false positives.

Ease of use: Good; regular web-based interfaces for analysts are easy to use; command-line or other advanced tools might have a steep learning curve.

Coverage: Monitoring of network traffic and logs from VMs; endpoint monitoring typically depend on having a standard agent installed.

Resources: Poor; hosted security tools are expensive, even more if any customisations are needed; substantial time investment is needed for the initial configuration; changes to services run in the cloud impose additional maintenance burden.

Scalability: N/A.

Extensibility: Good; flexible rule-based configuration; anomaly detection possible if activity baselines are predefined.

Completeness: ~=~ IDS

2.1.10 Passive DNS[41]

Passive DNS (PDNS) systems gather information about DNS records, in particular time points in order to provide historical information about such records. The systems help in tracking changes of malicious infrastructure in time, but also provide last known IP address of a domain if the DNS record is no longer available.

2.1.10.1 Evaluation

Type: Support.

Timeliness: Excellent; can be updated in real-time.

Accuracy: Excellent; PDNS is based on the actual queries and answers so it corresponds to actual resolutions of domains at certain points in time.

Ease of use: Excellent; data can be interpreted quickly and is supported by analytical tools.

Coverage: All domains resolved within the local infrastructure can be monitored (data from local resolvers); clients using external resolvers, especially using DNS-over-HTTP (DoH) will avoid monitoring; global coverage (important for investigations) depends on the provider.

Resources: Excellent; low effort.

[41] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#passive-dns

Scalability: Good; scales with the number of sensors; monitoring more DNS traffic increase coverage but has diminishing returns.

Extensibility: N/A; often custom solutions are used.

Completeness: Fair; typically PDNS is used for correlation with other information and not standalone.

2.1.11 DNS request monitoring[42]

DNS request monitoring systems provide information about how often and when certain domain names were queried and by which addresses. Thanks to that, extended analyses can be performed, including popularity of domains, their activity lifetime, but also tracking of botnet clients when monitoring known C&C domains.

2.1.11.1 Evaluation

Type: Support.

Timeliness: Excellent; real-time.

Accuracy: Good; data is coming from actual DNS queries however without answers from authoritative name servers data can contain some noise.

Ease of use: Good; data are easy to interpret.

Coverage: All domains queried from the local infrastructure can be monitored (data from local resolvers); clients using external resolvers, especially using DNS-over-HTTP (DoH) will avoid monitoring; global coverage (important for investigations) depends on the provider but there are much fewer providers than for PDNS.

Resources: Excellent; low effort.

Scalability: Good; scales with the number of sensors; monitoring more DNS traffic increases coverage but has diminishing returns.

Extensibility: N/A; often custom solutions are used.

Completeness: Good; may reveal the profile (network or geographic distribution) of clients requesting a domain.

2.1.12 Other DNS monitoring[43]

DNS monitoring other than passive DNS and DNS request monitoring includes, for example, monitoring of new domain names in search of phishing sites or presence of domain names generated with DGA algorithms.

2.1.12.1 Evaluation

Type: Alert, support.

Timeliness: Poor; depends on the data source, but can have a delay of up to 24 hours.

[42] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#dns-request-monitoring
[43] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#other-dns-monitoring

Accuracy: Poor; typically all suspicious domains need to be manually verified.

Ease of use: Fair; depends on the tooling.

Coverage: In-house tools usually cover only a few selected top-level domains (TLDs).

Resources: Fair; analysts need to verify identified domains; often a custom solution which requires maintenance.

Scalability: N/A.

Extensibility: N/A; often custom solutions are used.

Completeness: Poor; needs further enrichment and correlation; post-GDPR important details in WHOIS are not easily accessible anymore.

2.1.13 Endpoint monitoring[44]

Endpoint monitoring systems provide means for gathering and logging information about events occurring on endpoint environments. Events can include application logs, file system monitoring or configuration monitoring. The gathered data can be then forwarded to aggregation systems such as SIEMs.

2.1.13.1 Evaluation

Type: Alerts, support.

Timeliness: Excellent; real-time.

Accuracy: Fair; actual accuracy varies, depends on the quality of rules and signatures used for identifying suspicious behaviour.

Ease of use: Fair; requires expertise to interpret and search the logs.

Coverage: Monitoring of local infrastructure; coverage depends on how widely the collection has been implemented.

Resources: Fair; can generate large amount of logs that need storage and indexing, this implies investments in the backend hardware.

Scalability: Fair; scaling can pose an IT challenge, since agents need to be deployed on a wide range and large number of endpoint devices.

Extensibility: N/A; tool-dependent.

Completeness: Good; endpoint logs can provide details that are not possible to obtain otherwise (for example from network traffic).

2.1.14 X.509 certificates monitoring[45]

X.509 certificate monitoring systems provide means for identifying network incidents through analysis of issued certificates. Monitoring can be performed with checking for certificates issued

[44] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#endpoint-monitoring
[45] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#x509-certificates-monitoring

for phishing sites or with hunting of connections to websites for which blacklisted certificates were issued.

2.1.14.1 Evaluation

Type: Support

Timeliness: Excellent; near real-time.

Accuracy: Excellent; certificate monitoring provides information about actual certificates issued or encountered in the wild.

Ease of use: Good; data can be interpreted easily.

Coverage: Certificate Transparency (CT) provides coverage of all new certificates issued by main Certificate Authorities; other certificates are primarily coming from Internet scans and their coverage vary between providers.

Resources: Good; exiting tools for using CT; exception: Internet scanning requires significant resources.

Scalability: N/A.

Extensibility: N/A; typically custom solutions.

Completeness: Fair; X.509 provide multiple details to pivot on but further correlation is usually required.

2.1.15 Vulnerability scanning[46]

Vulnerability scanning systems are used to identify any vulnerabilities in the monitored environment. Depending on tool, they can provide basic information about services, but also give extended information about identified problems. These includes for example information about identified vulnerable services on particular ports, IP addresses or Autonomous Systems.

2.1.15.1 Evaluation

Type: Alerts.

Timeliness: Poor; depends on the scanning schedule; typically days or longer.

Accuracy: Poor; reports often need verification.

Ease of use: Good; large choice of COTS and open-source tools with various levels of sophistication.

Coverage: Monitoring of local infrastructure; in practice depends on which hosts are available for scanning.

Resources: Good; analyst's time is required to verify results of the automated tools; infrastructure is not significant unless large networks or the whole Internet are scanned.

Scalability: Good.

[46] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#vulnerability-scanning

Extensibility: Good; typically with multiple ways to add new ways of checking particular vulnerabilities,

Completeness: Fair; depends on the tool, provides data on vulnerability and the scanned service.

2.1.16 Automated spam collection[47]

Spam collection systems help in gathering spam sent to the monitored environment and, when equipped with analysis systems, give good insight into current spam campaigns. Monitoring of real or decoy mailboxes in the organisation's domain can be a very relevant resource of information for detecting targeted attacks. This information is a starting for identification of network incidents giving possible indicators of compromise, including information about attachments, links, involved malware, etc. However, it also helps with prevention of network incidents by, for example, constituting the basis for notification of targeted users.

2.1.16.1 Evaluation

Type: Alerts.

Timeliness: Good; usually new campaigns can be caught with low delay.

Accuracy: Good; there is a risk of non-spam emails being received.

Ease of use: Fair; easy if tooling supports grouping and analysis; otherwise monitoring can be more time-consuming.

Coverage: Spam targeting domains and mailboxes of interest; for monitoring of opportunistic attacks, the size and diversity of the collection infrastructure determines if most of wide-scale campaigns will be detected.

Resources: Good; low effort, unless spam is analysed on a large scale.

Scalability: Good; possible use of multiple sensors, domains and mailboxes; there are diminishing results from building a large collection infrastructure; large amount of retained messages can cause non-negligible requirements for the storage.

Extensibility: N/A; tool-dependent.

Completeness: Good; header, body and attachments provide a lot of information.

2.1.17 Sandbox (automated systems for behavioural analysis)[48]

Automated systems for malware behavioural analysis (malware sandboxes) are used to provide information about the behaviour of systems after opening observed files, and as a result provide information about their maliciousness. Depending on technologies sandboxes provide information about network connections, used system libraries, modified registry items and other behavioural data. Also, the behaviour is usually analysed to provide information about malware families or variants. The sandbox systems can be operated using own infrastructure or used as a service, provided by many vendors, including free of charge bill plans.

[47] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#automated-spam-collection
[48] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#sandbox-automated-systems-for-behavioural-analysis

2.1.17.1 Evaluation

Type: Alerts, support.

Timeliness: Good; minutes or longer.

Accuracy: Poor; reports need verification; both false-positives and false-negatives are common.

Ease of use: Good; tools provide intuitive interfaces.

Coverage: Depends on the source of malware samples analysed; can vary from samples targeting a single organization to a large-scale lab with wide range of malware campaigns being analysed.

Resources: Fair; requires effort to understand and verify reports.

Scalability: Good; processing of a large number of samples requires an appropriate processing infrastructure and, more importantly, storage for behavioural reports.

Extensibility: Excellent; usually multiple ways of adapting behaviour and extending analytical capabilities.

Completeness: Good; behavioural reports contain a lot of details.

2.1.18 Automated mobile malware analysis[49]

Automated mobile malware analysis systems provide analyses similar to standard sandboxes and static analysis tools targeting desktop platforms. Most tooling focuses solely on the Android OS family. A majority of the off-the-shelf solutions that offer malware detection capability are online services. Local analysis often requires deploying a custom solution, for example based on existing emulators, and is rarely fully-automated with some human supervision involved.

2.1.18.1 Evaluation

Type: Alerts, support.

Timeliness: Good; minutes or longer.

Accuracy: Poor; with exceptions, common automated analysis tools have high false-negative rate, since malware commonly evades detection; typically manual inspection by a human analyst is required to confirm if an application is harmless or not; for online services, the mechanism where community members can vote on the maliciousness of a sample might yield results that are significantly better than any automated analyses.

Ease of use: Good; web-based interfaces for browsing results are intuitive; other functionalities like advanced search and APIs can be more challenging.

Coverage: Depends on the source of malware samples analysed; samples are mostly collected from various public sources, including app markets; filtering samples that are relevant for the constituency might require manual work.

[49] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#automated-mobile-malware-analysis

Resources: Fair; developing and deploying a sample collection mechanisms and processes can be challenging; initial triage of samples might require manual inspection; actual number of samples to fully analyse for a typical teams is usually no more than dozens per day, however their analysts time is required to interpret the results.

Scalability: N/A; a typical team will not deal with the number of samples that requires scaling of automated tools; when manual reverse engineering is needed, availability of human analysts might be a bottleneck.

Extensibility: Fair; online services often provide ability for rule-based threat hunting; more advanced modifications.

Completeness: Excellent; online tools provide detailed static and dynamic analysis results and often verdicts from AV scanners.

2.1.19 Automated static malware analysis[50]

Automated static malware analysis systems help in the analysis of malicious files without using dynamic analysis methods. The systems can operate using binaries and memory dumps in order to extract static configuration of malware. Their functionality can be extended by equipping it with the YARA signature matching.

2.1.19.1 Evaluation

Type: Alert, support.

Timeliness: Good; seconds to minutes.

Accuracy: Good; quality of results depends on the rules and methods used.

Ease of use: Fair; some of the tools can require expertise in malware analysis to operate.

Coverage: Depends on the source of malware samples analysed; can vary from samples targeting a single organization to a large-scale lab with a wide range of malware campaigns being analysed.

Resources: Good.

Scalability: Good; processing of a large number of samples requires appropriate processing infrastructure.

Extensibility: Good; tools often are rule-based, have plugin mechanisms or can be arranged in different processing workflows.

Completeness: N/A; tool-dependent; some of the tools can provide important indicators that would be impossible to obtain otherwise without manual analysis of malware.

2.1.20 Leak monitoring[51]

Leak monitoring systems provide means for detection of information leakage by scanning possible hosting sources of leaked information. This includes pastes services, but also other sources such as code repositories or cloud services.

2.1.20.1 Evaluation

[50] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#automated-static-malware-analysis
[51] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#leak-monitoring

Type: Alerts.

Timeliness: Fair; depends on data sources; can vary from seconds to days.

Accuracy: Poor; matching is typically done by regular expressions and yield many false positives.

Ease of use: Good; tools provide convenient interfaces to browse the data.

Coverage: Depends on the sources; usually new content is matched against a predefined set of rules and only data relevant to the constituency are processed.

Resources: Good; alerts need manual verification but their number is usually not very high.

Scalability: N/A

Extensibility: Fair; typically, tools support adding new data feeds to monitor.

Completeness: Poor; data dumps often come with little context and require further analysis.

2.1.21 Media/news monitoring[52]

Media/news monitoring systems help in obtaining operational information from traditional news sources, such as newspapers, but also blogs and social media services, for example Twitter. The latter can be especially helpful, as many information security researchers post current information there, before some longer forms are filled in on other platforms.

2.1.21.1 Evaluation

Type: Alerts.

Timeliness: Fair; varies; from minutes to days.

Accuracy: Fair; credibility depends on the source.

Ease of use: Excellent; news aggregators/monitors have intuitive interfaces.

Coverage: Depends on the sources monitored; can cover majority of traditional outlets, thematic blogs and Twitter.

Resources: Fair; requires time to read summaries of articles and do further research in some cases.

Scalability: N/A

Extensibility: Fair; typically, tools support adding new feeds of information and configuring how they are processed.

Completeness: Fair; depends on the source and the item; can be very ambiguous or may contain sufficient technical information.

2.1.22 Client honeypots[53]

[52] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#medianews-monitoring
[53] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#client-honeypots

Client honeypots are systems designed to mimic the behaviour of a client application in order to detect malicious behaviour of servers. In general, the honeypot interacts with a probed server, and is then analysed to uncover any malicious activity. Different applications can be monitored, however most common are web browsers. As with server honeypots, client honeypots can be high-interaction or low-interaction. The former uses environments similar to those used by standard clients in order to provide similar conditions. This is especially useful during the analysis of Exploit Kits, as such systems perform machine fingerprint and could not operate in systems with limited capabilities. The downsides of this type of honeypot are that it is harder to deploy than the low interaction ones and it requires resources for running virtual machines with the mimicked operating systems. The low interaction honeypots provide limited capabilities comparing to the high interaction type, however they are easier to deploy and require less resources per mimicked client. Usually, a honeypot simulates some basic functions of the client, then the responses from the server are analysed in search of known traces of malicious behaviour. The downside of this approach is that some unknown attacks could remain undetected, as for example with the already mentioned Exploit Kits.

2.1.22.1 Evaluation

Type: Alerts, support.

Timeliness: Poor; varied, depends on scan frequency; typically not quicker than hours.

Accuracy: Fair; depends on the detection method used by the tool: some systems (especially high-interaction ones) can be prone to false positives; low-interaction honeypots may not provide correct results at all, leading to false negatives.

Ease of use: Poor; expertise is needed both to configure and to interpret data obtained from client honeypots.

Coverage: Depends on the set of pages selected from scanning: can vary from own infrastructure, sites likely visited by the constituency, entire TLD or more.

Resources: High; alerts need verification; interpretation of output can take time.

Scalability: Good; scanning can be distributed.

Extensibility: N/A; tool-dependent.

Completeness: Fair; in principle can provide complete details of the interaction with a server; actual level of details depends on a tool, with low-interaction honeypots providing more details in general.

2.1.23 Server honeypots[54]

Server honeypots are systems mimicking servers to detect and analyse malicious behaviour of clients. A plethora of different honeypots exists, covering multiple services, including general purpose servers, web servers, FTP, databases, VoIP, SCADA etc. Generally, server honeypots open ports for popular services and analyse received data. This approach provides an opportunity to detect scanning activities, exploit attempts, as well as other behaviour. The level of interaction can differ between honeypots, dividing them into low and high interaction groups. The former provides environment to emulate behaviour of a service, thus providing limited capabilities compared to real server. However, the deployment is relatively easy and resources needed for running a single instance are lower than in a high interaction honeypot. High interaction honeypots usually are deployed based on real servers, giving detailed information

[54] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#server-honeypots

about the attacker's behaviour, which is not available in low interaction systems. However their resources requirement is higher and without appropriate mechanisms they are prone to being compromised by the attacker.

2.1.23.1 Evaluation

Type: Alerts.

Timeliness: Excellent; real-time.

Accuracy: Fair; while there should be no legitimate connections to honeypots, in practice there might be a lot of irrelevant activity from scanners or misconfigured devices; most honeypots can be identified by attackers or might have faults in service emulation, which might prevent collection of essential attack details.

Ease of use: Fair; interpretation of the output and finding relevant information requires expertise.

Coverage: Depends on the range of exposed services, deployment model (internal network or external address) and advertising (to make the honeypot discoverable).

Resources: Analysis of details of attacks can be time-intensive; no human resources needed in a fully automated setup, however honeypots provide much less value then.

Scalability: Good; scales with the number of sensors; there are diminishing results from building a large sensor infrastructure (both in internal and external deployment models).

Extensibility: N/A; tool-dependent; there exist multi-service honeypot that offer very good plugin support.

Completeness: Fair; typically complete details of a network session but at the same time fundamentally limited by the level of interaction that the honeypot can offer to attackers.

2.1.24 Monitoring of sector specific technologies[55]

Different sectors (for example aviation, health, etc.) have specific software and hardware that needs to be monitored for intrusions and other suspicious activity. However in practice, the measures used for detection fall into one of the following categories: 1) OT networks, where measures are described in "Monitoring specific to industrial control systems", or 2) IT networks, where majority of measures described above are applicable.

2.2 INFORMATION SOURCES

2.2.1 Feeds of malware URLs[56]

Blacklists consisting of websites that have been observed hosting malware or exploit kits. By correlating such lists with logs from local network monitoring, such as proxy logs, it is possible to detect connections that may have resulted in a malware infection. While all major browsers have built-in URL blacklists, their content is provided by the vendors, so using additional sources can increase coverage. This especially applies to less-common malware or information from non-public sources that may not be available in the vendor's blacklists. Additionally, once a malware URL is discovered, it is possible to search the logs retroactively to find any past connections to the malicious websites. Another common use of this data is using it for detecting

[55] https://github.com/enisaeu/IRtools/blob/master/measures_and_tools.md#monitoring-of-sector-specific-technologies
[56] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-of-malware-urls

malware hosted in the constituency, which should trigger a remediation process. Since malware is often hosted on compromised websites that are abused by criminals, typically only full URLs can be used for detection, since corresponding domains and IP addresses may host other, harmless content.

2.2.1.1 Evaluation

Timeliness: Fair; most of the data is collected via automated means by crawlers scanning websites on a large scale; opportunistic malware campaigns on popular websites will be discovered quickly, however more targeted attacks or less visited sites can take significant amount of time until they appear in such feeds, if they are detected at all.

Accuracy: Low; significant false positive ratio; by the time a report is received many of the URLs no longer serve malicious content.

Ease of use: Fair; using bulk feeds for detection or blocking can be challenging due to the volume of the data and quality issues; however, using this information for early identification of compromised machines in own constituency is very straightforward and recommended even for teams with less resources.

Data volume: Global blacklists contain a lot of entries; often a single compromised website is reported multiple times, since malware is hosted at different URLs.

Completeness: Poor; usually very limited amount of information besides URLs and detection times.

2.2.2 Feeds of phishing sites[57]

Lists containing recently reported or active phishing URLs. A significant part of this data comes from user reports however there is a large ecosystem of entities dealing with finding and taking down phishing sites. By correlating such lists with logs from local network monitoring, such as proxy logs, it is possible to detect a possible theft of credentials. While all major browsers have built-in URL blacklists, their content is provided by the vendors, so using additional sources can increase coverage. This especially applies to more targeted phishing campaigns and information from non-public sources that may not be available in the vendor's blacklists. Additionally, once a phishing URL is discovered, it is possible to search the logs retroactively to find any past connections to the malicious websites. Another common use of this data is using them for detecting malware hosted in the constituency, which should trigger a remediation process. Since phishing is often hosted on compromised websites that are abused by criminals, typically only full URLs can be used for detection, since corresponding domains and IP addresses may host other, harmless content.

2.2.2.1 Evaluation

Timeliness: Good; phishing databases often obtain their data from a community of users (both individuals and companies), which means that there is variance in reporting times; wide-scale phishing campaigns are usually discovered quickly.

Accuracy: Low; different providers have various levels of vetting reports; by the time a report is received, many of the URLs no longer serve malicious content.

Ease of use: Fair; using bulk feeds for detection or blocking can be challenging due to the volume of the data and quality issues; however, using this information for early identification of

[57] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-of-phishing-sites

compromised machines in own constituency is very straightforward and recommended even for teams with less resources.

Data volume: Medium, as most phishing reports are reported or verified manually.

Completeness: Fair; apart from the URL and timestamp, providers often share the name of the brand that is being targeted and the current status of the phishing.

2.2.3 Feeds of botnet command and control servers[58]

Data on command and control servers used by malware, usually domains or IP addresses. This information is obtained by analysing individual malware samples or tracking the infrastructure used by threat actors. Addresses of command and control servers are very good network IoCs and can be used for real-time detection and blocking, but also for identification of infected machines by correlating them with network activity logs, for example netflow.

2.2.3.1 Evaluation

Timeliness: Fair; new addresses are often added after manual analysis, which can take hours or days; some sources provide data from automated tracking of specific botnets, these information can be close to real-time.

Accuracy: Good; C&C servers are usually verified before being added to a blacklist.

Ease of use: Excellent; C&C addresses can be easily correlated with network logs using existing tools.

Data volume: Low, the number of C&C servers is much smaller than other types of malicious infrastructure.

Completeness: Fair; sufficient for detection and blocking: domains or IP addresses and malware name; some sources provide additional malware-specific details that can be used for in-depth investigations.

2.2.4 Feeds of infected machines (bots)[59]

IP addresses of machines infected with malware. The primary use of these reports is the identification of compromised machines in the constituency for remediation purposes. If the team does not have the authority to clean up infections, this data is used for notification. Notifications are implemented by sending email reports to responsible entities (common for national CSIRTs), by putting users in so-called "walled gardens" (possible for ISPs) or other by means suitable for specific environment.

2.2.4.1 Evaluation

Timeliness: Good; varies between daily to hourly or even real-time.

Accuracy: Excellent; this type of data is mostly collected through sinkholes, which have a low ratio of false-positives as they should not receive any legitimate connections; the main challenges are NAT and DHCP, which complicate identification of the actual infected machines.

Ease of use: Good; existing workflows and tools are well suited for handling these sources; this also applies for large-scale notifications (especially national teams).

[58] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-of-botnet-command-and-control-servers
[59] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-of-infected-machines-bots

Data volume: Depending on the constituency size; this can be one of the biggest data sources, especially when a large botnet is taken down and sinkholed; national teams may receive hundreds of thousands of reports per day or more; since the majority of infections affect home users, corporate networks receive relatively small number of reports; amount of data varies during the day, according to the daytime usage of computers in general.

Completeness: Fair; generally sufficient for notification (IP address, timestamp, ports); the main issue is the naming of malware families and botnets, as providers do not use a common taxonomy or even do not provide any name at all in some cases, which makes proper identification a challenge.

2.2.5 Feeds with information on sources of abuse (spam, attacks, scanning)[60]

Information on hosts that are responsible for various malicious activity on the Internet, including sending spam, performing port scans, making exploitation attempts, etc. One of the main ways of collecting this information are server honeypots listening on public IPv4 addresses. This data often come in a form of various blacklists, sometimes aggregated by entire networks. The original purpose of such blacklist was using them for filtering, however due to the ease that attackers can change infrastructure, effectiveness of such approach is questionable. Nevertheless, these sources are valuable for detection of compromised machines or bad actors in own constituency. There are also services that aggregate abuse reports from multiple sources to simplify access to the data.

2.2.5.1 Evaluation

Timeliness: Fair; most sources provide data aggregated daily, however, there are some providers with more frequent updates, even real-time.

Accuracy: Good; typically, there is a small risk of false positives.

Ease of use: Good; coordinating teams have automation tools available that can handle this type of information well; for individual organisations an major challenge might be the large number of different sources of this data, with different formats and access mechanisms.

Data volume: For non-corporate networks with a large number of users the amount of reports can be very high; there is often an overlap of reports of abuse and other malicious activity related to the same addresses, since threat actors commonly use compromised machines for further attacks, spamming, proxies etc.

Completeness: Poor; apart from classification and IP addresses, there is usually little additional detail.

2.2.6 Information sharing platforms[61]

Systems that facilitate exchange of IoCs, advisories and other threat intelligence. One of the main benefits of such platforms is the fact that they aggregate a large amount of information and provide convenient ways to access it. Depending on the platform, most of the content may come from the vendor, individual researchers, or CSIRTs. Some users provide their original findings to but also a significant part might be information obtained elsewhere that is imported into the platform for correlation and easier access. One of the common use cases for the information sharing platforms is exporting IoCs in bulk and using them for real-time detection. Analysts may also want to browse individual advisories and follow-up the relevant ones with

[60] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-with-information-on-sources-of-abuse-spam-attacks-scanning
[61] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#information-sharing-platforms

investigations to determine if the constituency has been affected by a similar attack. While almost all of the platforms are provided as an online service, MISP is an exception as it allows self-hosting and uses a federated model of sharing.

2.2.6.1 Evaluation

Timeliness: Good; varies: depends on how active the user community is; information from other sources is usually imported within hours to days.

Accuracy: N/A; depends on the particular contribution; in general there is no platform-wide verification of the data and the consumer must understand trustworthiness of different contributors.

Ease of use: Fair; features offered by the platform come with additional complexity; in general, personnel needs additional training to take full advantage of these sources.

Data volume: Data aggregation means that the total amount of entries can be high and keeping up with the new contributions can be challenging for analysts; on the other hand, the amount of data should not be a problem in the case of automated processing.

Completeness: N/A; varies: platforms allow for adding rich contextual information, however, the actual level of detail depends on the contributor; usually the amount of data is sufficient for understanding the context and using the data for proactive detection.

2.2.7 Network indicators of compromise for monitoring[62]

Feeds of indicators that describe patterns in network traffic corresponding to known attacks, botnet communication, etc. and are specifically tailored for use in NIDSes. While commercial IDS vendors provide their own feeds, it is usually possible to add custom rules. The de-facto standard for network indicators are rules compatible with Snort[63], a rule language which balances the expressive power and a design that allow the analysis of multiple gigabits of traffic per second in real-time. Most of the rules focus on characteristic elements of the payload or application-protocol headers, as these are less likely to be changed by threat actors than IP addresses. Nevertheless, IPs or domains, especially of the C&C servers, can be used for indicators as well.

2.2.7.1 Evaluation

Timeliness: Good; commercial feeds are frequently updated, typically on a daily basis; other are usually less timely.

Accuracy: Fair; varies between providers; open source feeds tend to have higher false positive rates in comparison with commercial offerings.

Ease of use: Excellent; typically these feeds can be easily imported in IDSes.

Data volume: The number of rules is low enough not to require any special consideration for importing in IDSes.

Completeness: Good; classification of detected events; references to external analyses and taxonomies (for example Common Vulnerabilities and Exposures, CVEs[64]).

[62] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#network-indicators-of-compromise-for-monitoring
[63] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#snort-community
[64] https://cve.mitre.org/

2.2.8 Malware intelligence[65]

Services that provide information from static and dynamic analysis of malware samples and other related intelligence, going beyond a sandbox service. This type of services usually offer access to a large data repository with the analysis results and extensive query capabilities to facilitate investigations, research and tracking of particular malware families. One of the common methods for finding new malware samples relevant for the constituency is using YARA[66] signatures that can be matched against newly observed samples or for historical data.

2.2.8.1 Evaluation

Timeliness: Good; information is collected continuously; results of automated analyses are available within minutes of submission; malware samples and other types of intelligence might depend on submission by the community, however for large-scale crimeware campaigns, they are usually available within hours of appearing in the wild.

Accuracy: Fair; varied automated analyses, especially sandboxes, can provide data that is not suitable as indicators; for most of the information analyst's interpretation is required.

Ease of use: Fair; personnel must have some understanding of malware analysis to use such services and interpret results; more advanced functionality requires in-depth expertise.

Data volume: N/A; services are backed by large data repositories, however teams search and access only the information that is needed for tracking particular threats (for example specific malware families) or relevant to an investigation; full datasets are not shared.

Completeness: Excellent; these type of sources can provide very detailed analysis reports, with a lot of contextual information and observables that can be used for further pivoting.

2.2.9 Feeds of defaced websites[67]

Lists of compromised websites with modified content. It is an important source for detecting defacements before the affected entity reports it to the team, so the remediation process can be triggered as soon as possible.

2.2.9.1 Evaluation

Timeliness: Good; the primary source of these reports are user submission, so the timeliness may vary, however it is usually within hours.

Accuracy: Fair; entries may be verified manually by the provider, which ensures certain trustworthiness; otherwise such reports must be treated with care as false reports, which can happen often.

Ease of use: Good; the information is straightforward to handle and does not require advanced tooling to process.

Data volume: Number of defacements relevant to the constituency is typically small, which means that each case can be verified and investigated.

Completeness: Fair; reports contain URL of the affected page; other details might include a mirrored version or some information of the threat actor.

[65] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#malware-intelligence
[66] https://github.com/Yara-Rules/rules
[67] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-of-defaced-websites

2.2.10 Feeds of vulnerable services[68]

Lists of network services that have a known vulnerability or the fact that they are exposed on a public address may pose a security risk (exposing them is against good practice). Data is obtained by large-scale of IP space (usually IPv4, however relevant subsets of IPv6 can also be scanned). This information is very valuable for teams, since early identification of vulnerable assets is key to preventing intrusions or may reveal machines that might be already compromised. Services that are not strictly vulnerable but misconfigured in a way that make them prone to abuse as DDoS reflectors are also included in this category.

2.2.10.1 Evaluation

Timeliness: Poor; depends on the scanning frequency; popular services can be scanned even daily, less common only occasionally.

Accuracy: Good; providers typically take care to avoid false positives in the results, however some degree of verification is still needed.

Ease of use: Good; existing workflows and tools are well suited for handling these sources; this also applies for large-scale notifications (especially national teams).

Data volume: Corresponding to the number of publicly accessible hosts in the constituency; for large networks the number of reports daily can be in thousands, for national teams can exceed hundreds of thousands.

Completeness: Fair; information contains scan time, IP address and name of the vulnerability; some providers offer additional details on the vulnerable service and other services running on the same host.

2.2.11 Sector-specific advisories

Information concerning entities in a particular sector, either because of specific technologies being affected (for example aviation systems) or an attack that is targeting the sector. This type of information is typically in the form of advisories, possibly with technical data such as IoCs attached. The information is usually shared between companies, regulators or other entities being part of the sector, with smaller involvement of public sources or major security vendors in comparison with other categories. There is also an important role of ISACs and both open and closed informal sharing groups for information exchange.

2.2.11.1 Evaluation

Timeliness: Fair; attacks and vulnerabilities can be disclosed with a significant delay; depends on the victim's willingness to share information or vendor's processes..

Accuracy: Good; typically, only verified IoCs and vulnerabilities are shared.

Ease of use: Fair; information often comes in multiple formats and must be processed manually.

Data volume: Low.

Completeness: Good; usually sufficient information on the threat, context and mitigation steps is provided to make these reports actionable.

[68] https://github.com/enisaeu/IRtools/blob/master/information_sources.md#feeds-of-vulnerable-services

3. GLOSSARY AND ACRONYMS

Please refer to ENISA glossaries and lists of acronyms

- https://www.enisa.europa.eu/topics/threat-risk-management/risk-management/current-risk/bcm-resilience/glossary
- https://www.enisa.europa.eu/topics/csirts-in-europe/glossary
- https://www.enisa.europa.eu/media/media-press-kits/enisa-glossary

TP-01-20-305-EN-N

ABOUT ENISA

The mission of the European Union Agency for Cybersecurity (ENISA) is to achieve a high common level of cybersecurity across the Union, by actively supporting Member States, Union institutions, bodies, offices and agencies in improving cybersecurity. We contribute to policy development and implementation, support capacity building and preparedness, facilitate operational cooperation at Union level, enhance the trustworthiness of ICT products, services and processes by rolling out cybersecurity certification schemes, enable knowledge sharing, research, innovation and awareness building, whilst developing cross-border communities. Our goal is to strengthen trust in the connected economy, boost resilience of the Union's infrastructure and services and keep our society cyber secure. More information about ENISA and its work can be found www.enisa.europa.eu.

ENISA
European Union Agency for Cybersecurity

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

Heraklion office
95 Nikolaou Plastira
700 13 Vassilika Vouton, Heraklion, Greece

enisa.europa.eu

ISBN: 978-92-9204-345-2
DOI: 10.2824/345137