



Proactive Detection of Network Security

Incidents

[Deliverable – 2011-12-07]



Contributors to this report

The report production was commissioned to CERT Polska / NASK.

- Authors: Katarzyna Gorzelak, Tomasz Grudziecki, Paweł Jacewicz, Przemysław Jaroszewski, Łukasz Juszczyk, Piotr Kijewski (CERT Polska / NASK)
- Editor/contributor: Agris Belasovs (ENISA)

Acknowledgements

We wish to thank all the 45 CERTs that responded to the initial survey and the following members of the expert group for providing additional input:

Name	Affiliation
Brian Honan	IRISS
Chris Camacho	World Bank
Corrado Leita	Symantec Research
Francisco Monserrat	IRIS-CERT
Gavin Reid	CISCO CSIRT
Harri Sylvander	FUNET CERT
Jan Goebel	Siemens
Jaroslav Petrovsky	CSIRT.SK
Jeff Carpenter	CERT/CC
Jose Nazario	Arbor Networks
Luis Morais	CERT.PT
Maarten van Horenbeeck	Microsoft
Marco Ho	HKCERT
Petra Hochmannova	CSIRT.SK
Roman Danyliw	CERT/CC
Shehzad Ahmad	DK-CERT
Till Dorges	PRE-CERT/PRESENSE
Torsten Voss	DFN-CERT
Udo Schweigert	Siemens-CERT
Wim Biemolt	SURFcert

About ENISA

The European Network and Information Security Agency (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 details

For contacting ENISA for general enquiries on CERT-related matters or this report specifically, please use the following details:

- Email: [cert-relations \(at\) enisa.europa.eu](mailto:cert-relations@enisa.europa.eu)
- Internet: <http://www.enisa.europa.eu>

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 ENISA Regulation (EC) No 460/2004 as lastly amended by Regulation (EU) No 580/2011. 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.

Reproduction is authorised provided the source is acknowledged.

© European Network and Information Security Agency (ENISA), 2011

Contents

1	Executive summary.....	1
2	Introduction and background.....	3
2.1	Objectives of the study and description of work	3
2.2	Intended target audience.....	4
2.3	Proactive versus reactive detection	4
2.4	Internal monitoring and external services: different approaches for proactive detection	4
2.5	What kind of network incident data are available from external data feeds?	5
3	Methodology used.....	6
3.1	Desktop research.....	6
3.2	Survey of CERTs in Europe.....	6
3.3	Establishing an expert group, initiating and moderating discussions.....	7
3.4	The workshop.....	7
3.5	Analysis of the measures identified	7
4	Analysis of the survey results	8
4.1	Respondents' organisation profile	8
4.2	Methods of data acquisition	9
4.3	External sources of information.....	12
4.4	Tools for internal monitoring	14
4.5	Resources available	19
5	Inventory and description of identified services and tools.....	22
5.1	Evaluation criteria	22
5.1.1	Timeliness.....	22
5.1.2	Accuracy of results	23
5.1.3	Ease of use.....	24
5.1.4	Coverage.....	24
5.1.5	Resources required.....	25
5.1.6	Scalability.....	26
5.1.7	Extensibility	26
5.2	Services for the proactive detection of network security incidents	27
5.2.1	DNS-BH Malware Domain Blocklist.....	27
5.2.2	MalwareURL	28
5.2.3	Dshield.....	30
5.2.4	Google Safe Browsing Alerts	32

5.2.5	HoneySpider Network	33
5.2.6	AusCERT	35
5.2.7	Cert.br Distributed Honeypot Project	37
5.2.8	FIRE (FInding Rogue nEtworks).....	38
5.2.9	Team Cymru – TC Console.....	40
5.2.10	EXPOSURE.....	42
5.2.11	Zeus/SpyEye Tracker	43
5.2.12	AMaDa	46
5.2.13	Malware Domain List.....	48
5.2.14	The Spamhaus Project (Spamhaus DNSBL Datafeed).....	49
5.2.15	Shadowserver Foundation	51
5.2.16	SGNET / Leurre.com HoneyNet Project	54
5.2.17	ARAKIS	55
5.2.18	Malc0de database	57
5.2.19	ParetoLogic URL Clearing House / malwareblacklist.com.....	58
5.2.20	SpamCop.....	60
5.2.21	Arbor ATLAS.....	61
5.2.22	Composite Blocking List.....	63
5.2.23	Team Cymru’s CSIRT Assistance Program	64
5.2.24	CERT.BR Spampots	65
5.2.25	Project Honeypot.....	67
5.2.26	Malware Threat Center	68
5.2.27	Smart Network Data Services.....	69
5.2.28	Malware Patrol	71
5.2.29	Zone-H	72
5.2.30	Cisco IronPort SenderBase Security Network	73
5.3	Tools/mechanisms for the proactive detection of network security incidents	75
5.3.1	Client honeypots	75
5.3.2	Server honeypot	78
5.3.3	Sandboxes.....	81
5.3.4	Firewall	83
5.3.5	IDS/IPS	85
5.3.6	NetFlow	87

5.3.7	Darknet.....	89
5.3.8	Passive DNS monitoring	91
5.3.9	Antivirus programs	93
5.3.10	Spamtrap	95
5.3.11	Web Application Firewall	96
5.3.12	Application logs	98
5.4	Summary of the evaluation of services and tools/mechanisms	99
6	Services and tools recommended for proactive detection by CERTs.....	101
6.1	Explanation behind the rationale for selection.....	101
6.2	Top 5 services (data feeds) for detection of network security incidents.....	102
6.2.1	Shadowserver Foundation	102
6.2.2	Zeus/SpyEye Tracker	103
6.2.3	Google Safe Browsing Alerts	103
6.2.4	Malware Domain List.....	103
6.2.5	Team Cymru’s CSIRT Assistance Program	103
6.3	Top ‘must have’ tools (and mechanisms) for proactive detection of network security incidents	104
6.3.1	Standard tools and mechanisms	105
6.3.2	Advanced tools and mechanisms.....	106
6.3.3	Upcoming tools and mechanisms	107
7	Identification of shortcomings in the proactive detection of incidents.....	108
7.1	Technical issues	108
7.1.1	Data quality and reliability concerns.....	108
7.1.2	Correlation is still limited	111
7.1.3	Lack of automation.....	113
7.1.4	Lack of common formats.....	114
7.1.5	Lack of own monitoring (sensor networks).....	116
7.1.6	Lack of client honeypot deployments and sandbox analysis capabilities.....	117
7.1.7	Visualisation still underutilised.....	118
7.1.8	Lack of incident reporting tools integrated with users’ desktop software	119
7.1.9	Lack of long-term trend analysis of incident data	120
7.1.10	Targeted attacks underreported	121
7.1.11	DDoS attacks underreported.....	121
7.1.12	Passive DNS monitoring underused	122

7.1.13	Lack of services for data leak reporting (data repatriation).....	123
7.2	Legal and Organisational issues.....	124
7.2.1	Legal issues impede data sharing.....	124
7.2.2	Lack of human resources.....	126
7.2.3	Obstacles in reaching closed groups.....	126
8	Summary of recommendations.....	128
8.1	Recommendations for data providers.....	128
8.1.1	Identification and vetting of data consumers.....	128
8.1.2	Data format and distribution.....	128
8.1.3	Data source quality enrichment.....	129
8.2	Recommendations for data consumers.....	129
8.2.1	Acquiring access to datasets.....	130
8.2.2	Integration of feeds with internal incident handling systems.....	130
8.2.3	Verification of quality of data feeds.....	131
8.2.4	Deployment of rising technologies.....	131
8.3	Recommendations for further activities on the EU and national level.....	131
9	Conclusions.....	133
10	Annex I: Abbreviations.....	134
11	Annex II: CERT survey analysis.....	136

1 Executive summary

This document is the final report of the ‘Proactive Detection of Network Security Incidents’ study. The goal of the study was to investigate ways in which CERTs – national and governmental ones in particular – proactively detect incidents concerning their constituencies, identify good practice and recommended measures for new and already established CERTs, analyse problems they face and offer recommendations to relevant stakeholders on what can be done to further this process. It is important to note that the results of the study are largely community driven. That is, they are based not just on research and the experience of the experts who conducted the study, but to a large extent on the results of a survey carried out amongst 105 different CERTs (which resulted in 45 responses overall) and external expert group input. The outcome is thus a work by the community for the CERT community.

Proactive detection of incidents is the process of discovery of malicious activity in a CERT’s constituency through internal monitoring tools or external services that publish information about detected incidents, **before the affected constituents become aware of the problem**. It can be viewed as a form of early warning service from the constituents’ perspective. Effective proactive detection of network security incidents is one of the cornerstones of an efficient CERT service portfolio capability. It can greatly enhance a CERT’s operations, improve its situational awareness and enable it to handle incidents more efficiently, thus strengthening the CERT’s incident handling capability, which is one of the core services of national / governmental CERTs.¹

The report covers 30 external services identified that can be used by CERTs to obtain information about their constituency, often in an automated manner. Some are public and some have restricted access. Most of them are free. In many cases, national and governmental CERTs can gain access to data covering an entire country. Additionally, 12 different categories of internal tools were identified (with specific tools as examples) that can be used by a CERT to detect incidents. Both external services and tools were rated according to several criteria defined during the study, and priorities for their implementation suggested.

The study has identified that CERTs are currently not fully utilising all possible external sources at their disposal – despite their wide availability and relative ease of use, and despite the fact that many CERTs declare their readiness to adopt new sources of information. Similarly, a large number of CERTs do not collect incident data about other constituencies. Even those that do, often do not share this data with other CERTs. This is an area of concern as exchange of such information is key to the effective combating of malware and malicious activities and is extremely important in a cross-border environment.

These and other shortcomings (16 overall) in the process of detection of incidents are examined in more depth, both on a technical and legal/organisational level. The most important technical gaps identified include problems with data quality (such as the existence of false positives in reports, poor

¹ ENISA, Baseline Capabilities of National / Governmental CERTs, Part 2:
<http://www.enisa.europa.eu/act/cert/support/files/baseline-capabilities-of-national-governmental-certs-policy-recommendations>

timeliness of delivery, lack of contextual information, no validity indicators, unclear data aging policies), and lack of automation and correlation, partly due to data quality issues, but also due to the lack of standard formats, tools or simply resources and skills. The most important legal problem involves privacy regulations and personal data protection laws that often hinder the exchange of information – an obstacle faced by CERTs but unfortunately not by miscreants responsible for network attacks.

For each identified shortcoming, one or more recommendations are formulated as part of the study. They are aimed at a) data providers, b) data consumers and c) organisations at the EU or national level. For data providers key recommendations focus on suggestions on how to better reach out to CERTs, better data format and distribution approaches as well as data quality improvement and enrichment. For data consumers a guide on how to acquire access to datasets is given, suggestions on better integration of external feeds with internal monitoring systems put forward, additional activities that can be performed by a CERT to verify quality of data feeds enumerated, and specific deployments of new technologies recommended. Finally, at the EU or national level, activities are pointed out that are aimed at achieving a balance between privacy protection and security provision needs, the encouragement of the adoption of common formats and underused technologies and the integration of statistical incident data on a wider scale. Research is also suggested into the area of data leakage reporting.

The summary of recommendations made per stakeholder group is provided in Section 8 (Summary of recommendations). Detailed shortcomings identified, together with extended recommendations to mitigate them, are provided in Section 7 (Identification of shortcomings in the proactive detection of incidents).

It is hoped that the results published here will encourage national/government CERT managers of both new and established CERTs to obtain access to many identified external sources of incident information as well as to consider additional internal tools to collect such information that they can deploy at their organisation. Enhancing their own network incident detection infrastructure enables them not only to get better at proactively detecting incidents in their own constituency but also to detect incidents that concern others. This fosters cooperation and data sharing between CERTs, which helps to resolve the incidents and improve the security of the Internet.

2 Introduction and background

This document is the final report of the 'Proactive Detection of Network Security Incidents' study conducted between April 2011 and September 2011. The study, commissioned by ENISA, is aimed at identifying and improving ways for CERTs to proactively detect network incidents. The document is structured as follows:

- Chapter 2 *Introduction and background* explains in more detail the research objectives of the study, intended target audience, the concept of proactive and reactive detection of incidents, as well as different approaches to detection.
- Chapter 3 *Methodology used* is a description of work carried out as part of the study, together with an analysis of a survey carried out amongst CERTs primarily in the European Union Member States and the setting up of an expert group.
- Chapter 4 *Analysis of the survey results* presents an overview of the survey results, which drives much of this study.
- Chapter 5 *Inventory and description of identified services and tools* is an inventory of existing services and tools that can be used by CERTs for proactive detection, along with subjective ratings.
- Chapter 6 *Services and tools recommended for proactive detection by CERTs* provides a list of prioritised services and tools for CERTs and explains the rationale for their selection.
- Chapter 7 *Identification of shortcomings in the proactive detection of incidents* identifies shortcomings and gaps in the area of proactive detection and gives recommendations on their mitigation
- Chapter 8 *Summary of recommendations* provides a summary of recommendations and best practice for incident data providers, data receivers and other relevant stakeholders

2.1 Objectives of the study and description of work

The objectives of this study are:

- to provide an inventory of available methods, activities and information sources (hereafter 'measures') for proactive detection of network security incidents, which are used already or potentially could be used by national/governmental and other CERTs
- to identify good practice and recommend measures for new and already established CERTs
- to outline possible further activities in order to mitigate the common shortcomings identified including tasks and roles of different stakeholders.

To achieve these objectives following activities were performed:

- Desktop research
- Survey among CERTs (mainly European)
- Analysis of the measures identified
- Discussions in the expert group meeting and mailing list

More details on the performed activities are provided in Chapter 3 *Methodology used*.

2.2 Intended target audience

The intended target audience for this report are the managers and technical staff of national / governmental CERTs. However, the report can be used by any other CERT or abuse team as well. It is aimed at both new and existing CERTs. New CERTs can use the report to quickly bootstrap their basic operations, while existing CERTs can identify missing data feeds as well as use the suggestions in the report to mitigate shortcomings in their detection and incident handling process. Data providers, including CERTs, may find suggestions on how to improve their data feed service to others.

2.3 Proactive versus reactive detection

For the purpose of the study, proactive detection of incidents is the process of discovery of malicious activity in a CERT's constituency through internal monitoring tools or external services that publish information about detected incidents, before the affected constituents become aware of the problem. Note that this is different from obtaining information through incoming incident reports from the constituency – in that case constituents are the first to know that an incident took place, and a CERT's role is entirely reactive. Hence proactive detection of incidents can be viewed as a form of an early warning service from the constituents' point of view. Effective proactive detection of network security incidents can greatly enhance a CERT's operational capabilities, improve its situational awareness and enable it to handle incidents more efficiently.

2.4 Internal monitoring and external services: different approaches for proactive detection

Based on the nature of the origin two distinct categories of sources of information about network security incidents can be distinguished:

- The first are tools that can be **deployed by a CERT** to internally monitor events in its constituency. For example, they may be firewalls, antivirus systems or honeypots. In some cases these will be deployed only on the CERT's network itself. In other cases they will span an entire enterprise that the CERT team is responsible for. In the most complex cases they will be part of a wider sensor network deployed across a constituency made up of independent entities that operates on a national scale.
- The second category collects the **services that have been made available on the Internet** which provide information about detected network security incidents to affected parties. These can be public, closed or commercial. Some are subscription based, where a recipient has to specify netblocks or autonomous systems that they are interested in and subsequently receive a data feed in various forms. Others provide all information that they can find regarding a detected incident, leaving it up to the recipient to parse and extract relevant information about a network under their jurisdiction. Public sources generally can be used by

anyone. Closed sources require some form of vetting of the recipient. In some cases their very existence is a secret and information on how to subscribe not published anywhere. Some sources come and go. Others pop up when new forms of attacks or means for their detection appear. These data feeds are run by various security organisations, projects, vendors, universities, CERTs or non-profit initiatives, or even enthusiastic individuals.

All of the above sources are of use to CERTs, which can subscribe to their services. In the case of national/government CERTs quite often netblocks not under the direct control of the CERT can be provided, essentially allowing for the submission of all netblocks belonging to a constituency, which can even be a whole country. This is of great value for CERTs that are unable to collect network incident data in a direct manner from constituent networks, which is ultimately the case for CERTs that have entire countries or regions as their responsibility.

2.5 What kind of network incident data are available from external data feeds?

Most external services offer incident data in the form of IP addresses, URLs, domains or malware associated with a particular malicious activity, such as a bot, C&C server, malicious URL or scanning. Sometimes more sensitive data are offered – such as stolen user credentials or credit card data. Some services may also offer alerts in more abstract forms depending on detection models used internally in the service.

3 Methodology used

This section describes in more detail the methodology used in this study and creation of the final report.

3.1 Desktop research

In this activity, information was gathered about the methods, activities, and internal or external (public, closed, commercial, etc.) information sources, already used or which can be used by a CERT for the proactive detection of network security incidents. We included experiences of the CERT Polska team in honeypot and client honeypot design and deployment, management of network early warning systems (such as ARAKIS – see 5.2.17), and results of analysis of data of such systems. Additional desktop research carried out on the Internet and academic publications (to identify potential new measures) was also performed by team members. The individual expertise and experience of team members helped to provide added value in this research.

The research was conducted in two phases:

1. During the first phase a list of services and tools which can be used by incident handling teams was created. The list of tools was divided into categories and the list of services was completed, with the most promising ones investigated in more detail.
2. The second phase of desktop research aimed at extracting the most important characteristics among categories of tools and services themselves. The research team used criteria such as timeliness, accuracy of results, ease of use, coverage, required resources, scalability and extensibility to provide as accurate as possible descriptions of important key features that can directly impact the proactive detection and incident handling processes.

The outcome of the research is an inventory of tools and services which, together with recommendations on activities and methods, can be used as a basis to improve or extend the operations of a CERT.

3.2 Survey of CERTs in Europe

A survey of CERTs in Europe was carried out to gather information on actual measures dealing with the area of proactive detection, currently used or planned to be implemented. The survey was also sent to selected non-European CERTs that were identified as playing an important role in the worldwide CERT community. The survey was carried out in the form of a web and email survey. It included an initial list of criteria for evaluation, leaving a field for suggestions from the survey participants. The relevant stakeholder groups included governmental CERTs, CERTs with national responsibility (based on the list as identified by the CERT/CC amended by our knowledge), ISP CERTs (or ‘abuse-teams’), academic CERTs, vendor CERTs, and CERTs from other institutions.

The survey was sent to 105 potential respondents from the established list. A total of 51 addresses of specific contact persons at the CERTs were used; in the remaining 54 cases generic CERT addresses

were used, as there was not a known person to get in touch with in every case. Three follow-up rounds of reminders were also carried out. Despite sending multiple rounds of reminders, no response was received from 11 individuals and refusals were received from five individuals. In all, 45 filled-in surveys were obtained.

The survey was carried out during April and May 2011. The survey included 96 questions – the full version of the survey can be found in Annex II: CERT survey analysis.

3.3 Establishing an expert group, initiating and moderating discussions

As part of this task an expert group was established. A Terms of Reference document for the work of the expert group was created to better explain the vision and goals of the study to facilitate better interaction within the expert group. All the participants in the survey were invited to take part in the expert group. The list of experts was also extended to include security specialists outside the CERT community but directly involved in the area of security threat monitoring and intrusion detection. To facilitate the exchange of information, an email discussion list was established, hosted by CERT Polska. The list was created after the survey was conducted. Experts on the list were asked to take part in the discussion of interim results: survey results and initial results of the desktop research.

3.4 The workshop

An expert group meeting of the ENISA Proactive Detection of Network Security Incidents study was held in Vienna on 14 June 2011, co-located with the FIRST² Conference, well attended by the CERT community. There were 18 experts at the workshop, including representatives of the CERT community and vendors. The goal of the meeting was to present and discuss the interim results of the study, in particular the results of the survey carried out by CERT Polska and ENISA concerning methods and tools used by CERTs to proactively detect incidents. The outcome of the discussion was a compilation of additional questions that warrant investigation and valuable remarks made by experts.

3.5 Analysis of the measures identified

This phase included analysis of all the information gathered during the study. The benefits and shortcomings of the measures identified during the previous tasks were elaborated. Good practices and recommended measures for new and already established national, governmental and other CERTs were proposed. For every shortcoming identified, the analysis included proposed ways of mitigating the shortcoming, including possible tasks and roles of different stakeholders. The analysis was based on an extended set of criteria for evaluation proposed initially. This led to the creation of the draft of the final report. The draft was then sent to ENISA and the expert group for comment. Feedback from ENISA and the experts was then incorporated in the final report.

² Forum for Incident Response and Security Teams: <http://www.first.org/>

4 Analysis of the survey results

The analysis of survey results was an important driver in the process of identifying shortcomings in the way CERTs proactively detect incidents. Without it, getting a full understanding of the problems CERTs faced is difficult. Interestingly, many of the problems CERTs focused on were not directly linked to the problem of detecting an incident but in its subsequent processing. The full presentation of the survey results can be found in Annex II: CERT survey analysis.

4.1 Respondents' organisation profile

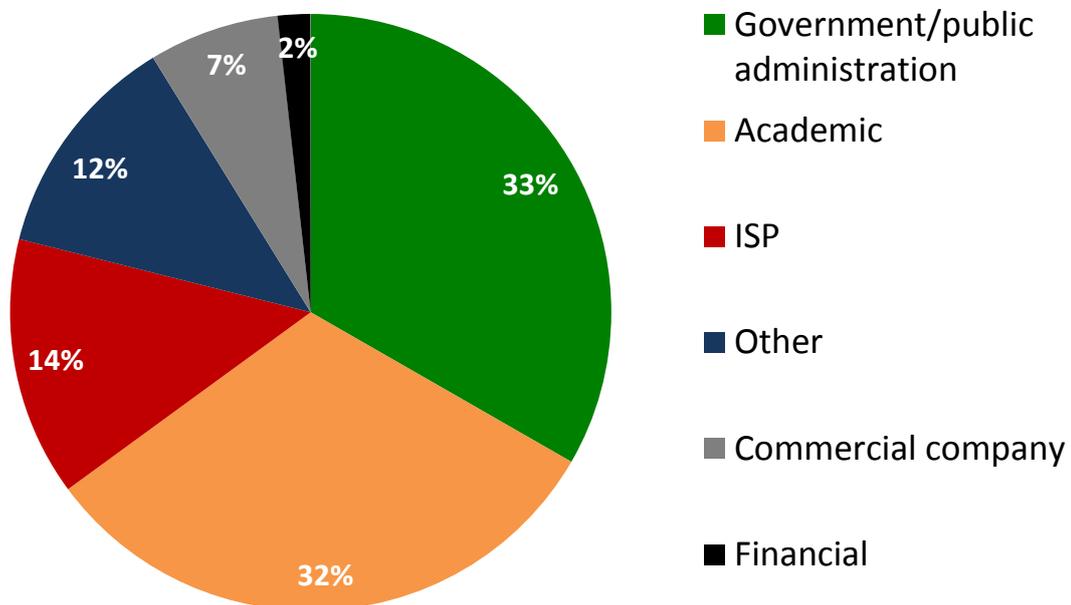


Figure 1: Respondents' organisation profile

The majority of the survey respondents were from government/public administration (33%) and academic organisations (32%), and a smaller but significant group of participants in the survey were *ISP* representatives (14%). There were additional single answers for *foundation* or *non-profit organisation*.

The group of investigated organisations was relatively varied. This is illustrated by a wide range in the number of incidents handled per year – stated as ranging from 10 to 2 million (although this may be due to differences in defining the term incident) – and number of full-time equivalent employees of the surveyed CERTs – from 0.5 to 41, giving an average of 9 employees.

4.2 Methods of data acquisition

Participants in the survey were asked several general questions about methods by which they acquire data, as well as what methods they see particularly fit for this purpose and what they find missing or otherwise problematic. A summary of the most significant questions and answers is provided in Figure 2.

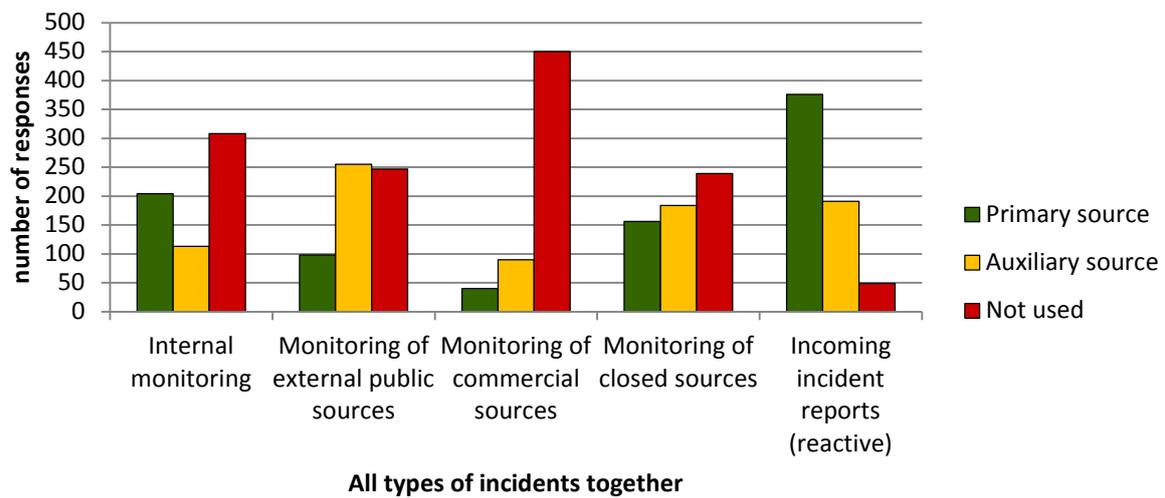


Figure 2: How do you obtain incident-related data about your constituency?

Note that in the chart, the number of responses is higher than the number of CERTs answering the survey (45) because the respondents were asked how they obtain information related to 16 incident types, with multiple choice being possible. These are all added together.

As the survey showed, most of the surveyed CERTs use *incoming incident reports* and *internal monitoring* as two basic ways of obtaining incident data. The high position of incoming incident reports is probably due to the fact that this is the basic way the CERTs obtain the information. When a new CERT is established it usually starts obtaining information by offering its constituents the option of reporting incidents directly, because at the very beginning of operation it is the easiest and most cost effective way for the CERT to gain information. In the course of time the CERTs evolve and develop more sophisticated ways of obtaining information, but they never stop using incoming incident reports, as it is the fundamental service of CERTs.

Worth noting is the relative unpopularity of *monitoring of commercial sources*. The reasons for this may be the financial effort and/or workload required (more formal agreements, legal concerns) to be allowed to gain access. Commercial companies that share their information usually demand to be paid for access to their source. It is sometimes possible to get this information without a fee but it depends on very good relations between the CERT and the entity, and individual agreements. This option is more often available for CERTs with a good, established position in the community.

Monitoring of closed data sources is an important category. Going deeper into the survey dataset it can be seen that information related to malicious URLs, botnets (especially C&C servers) and malware in general is mostly obtained through closed sources.

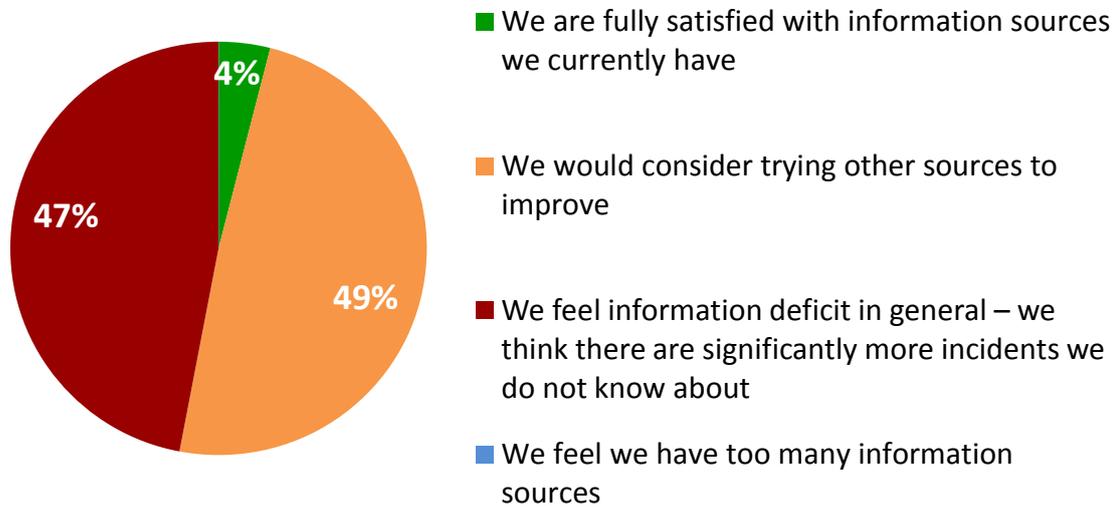


Figure 3: General feelings regarding information sources in respondent's constituency

Many additional answers in the survey and lively discussions on this question during the workshop showed that this is one of the crucial questions asked. Basically, respondents are divided into two groups: those who would be willing to consider trying new sources of information and those who feel that there is a general information deficit.

Not one respondent stated that they had too many information sources, which points to a unanimous need for more information. More specific questions as well as discussion in the experts group showed that the real concern is not so much the number of sources available as the quality of data provided, ease of use and type of incidents which are reported – and specifically those which are not reported.

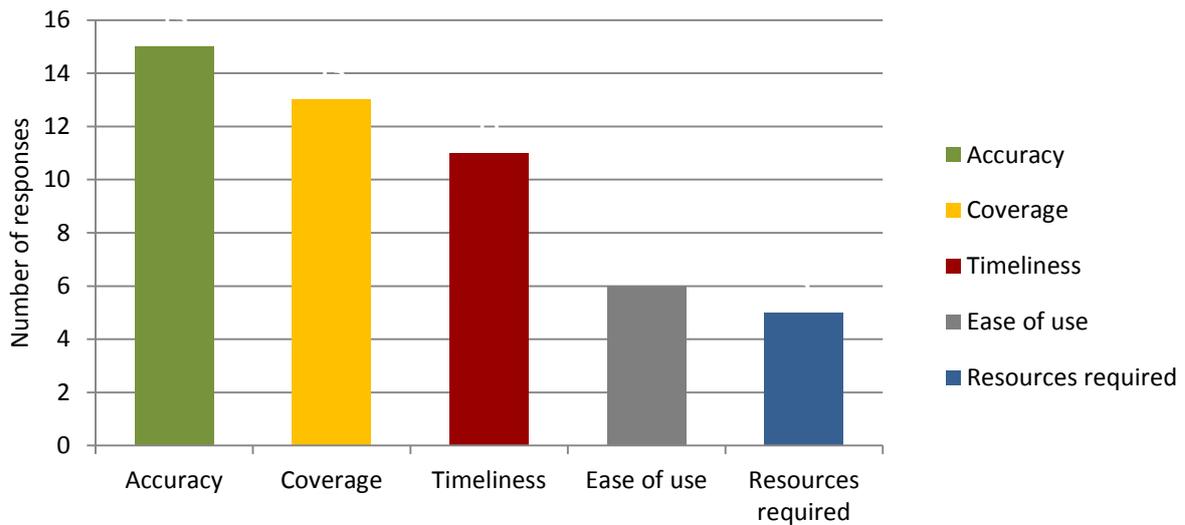


Figure 4: What would you like to improve when trying new sources in obtaining information about security incidents in your constituency? (see 5.1 for definitions of the above terms)

Taking timeliness, accuracy and coverage as elements of data quality, answers to this question clearly point to issues in this area. The quality of data diminishes significantly when any of these three factors is subpar. Consider timeliness. Even if the information is good, reliable and exhaustive, a time of delivery that is too long means that it may become out-of-date. Many members of the expert group during the workshop discussion also noted that in their opinion it is not more data that are needed (in particular for mature CERTs), but data of better quality. Fewer answers highlighted the 'ease of use' and 'resources required' factors, even though they were considered by many as areas to be improved as well. These factors are associated with the overhead required to add and receive data from external sources.

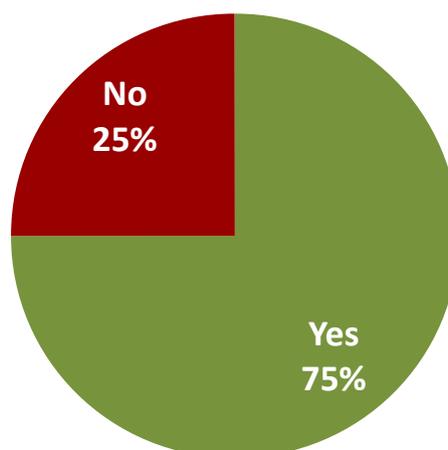


Figure 5: Do you think that some incidents are underreported?

Three-quarters of the respondents felt that some incidents are underreported. Respondents often pointed to specific attack categories that they felt were being missed; this included targeted attacks, DoS/DDoS (including failed DoS-attacks) and data leakage. Respondents also mentioned that incidents that for one organisation may seem unimportant might be important for another. Moreover, in cases when revealing information about an incident can influence the image of the organisation it is often left undisclosed.

4.3 External sources of information

The survey included a large section with questions about specific external data sources. Sources included in the survey were identified by the desktop research as most frequently used by CERTs to proactively acquire data about network security incidents. Respondents were also able to enumerate and evaluate other external sources they were using. We also included questions about non-public sources that cannot be named and disclosed and how they compare with public ones.

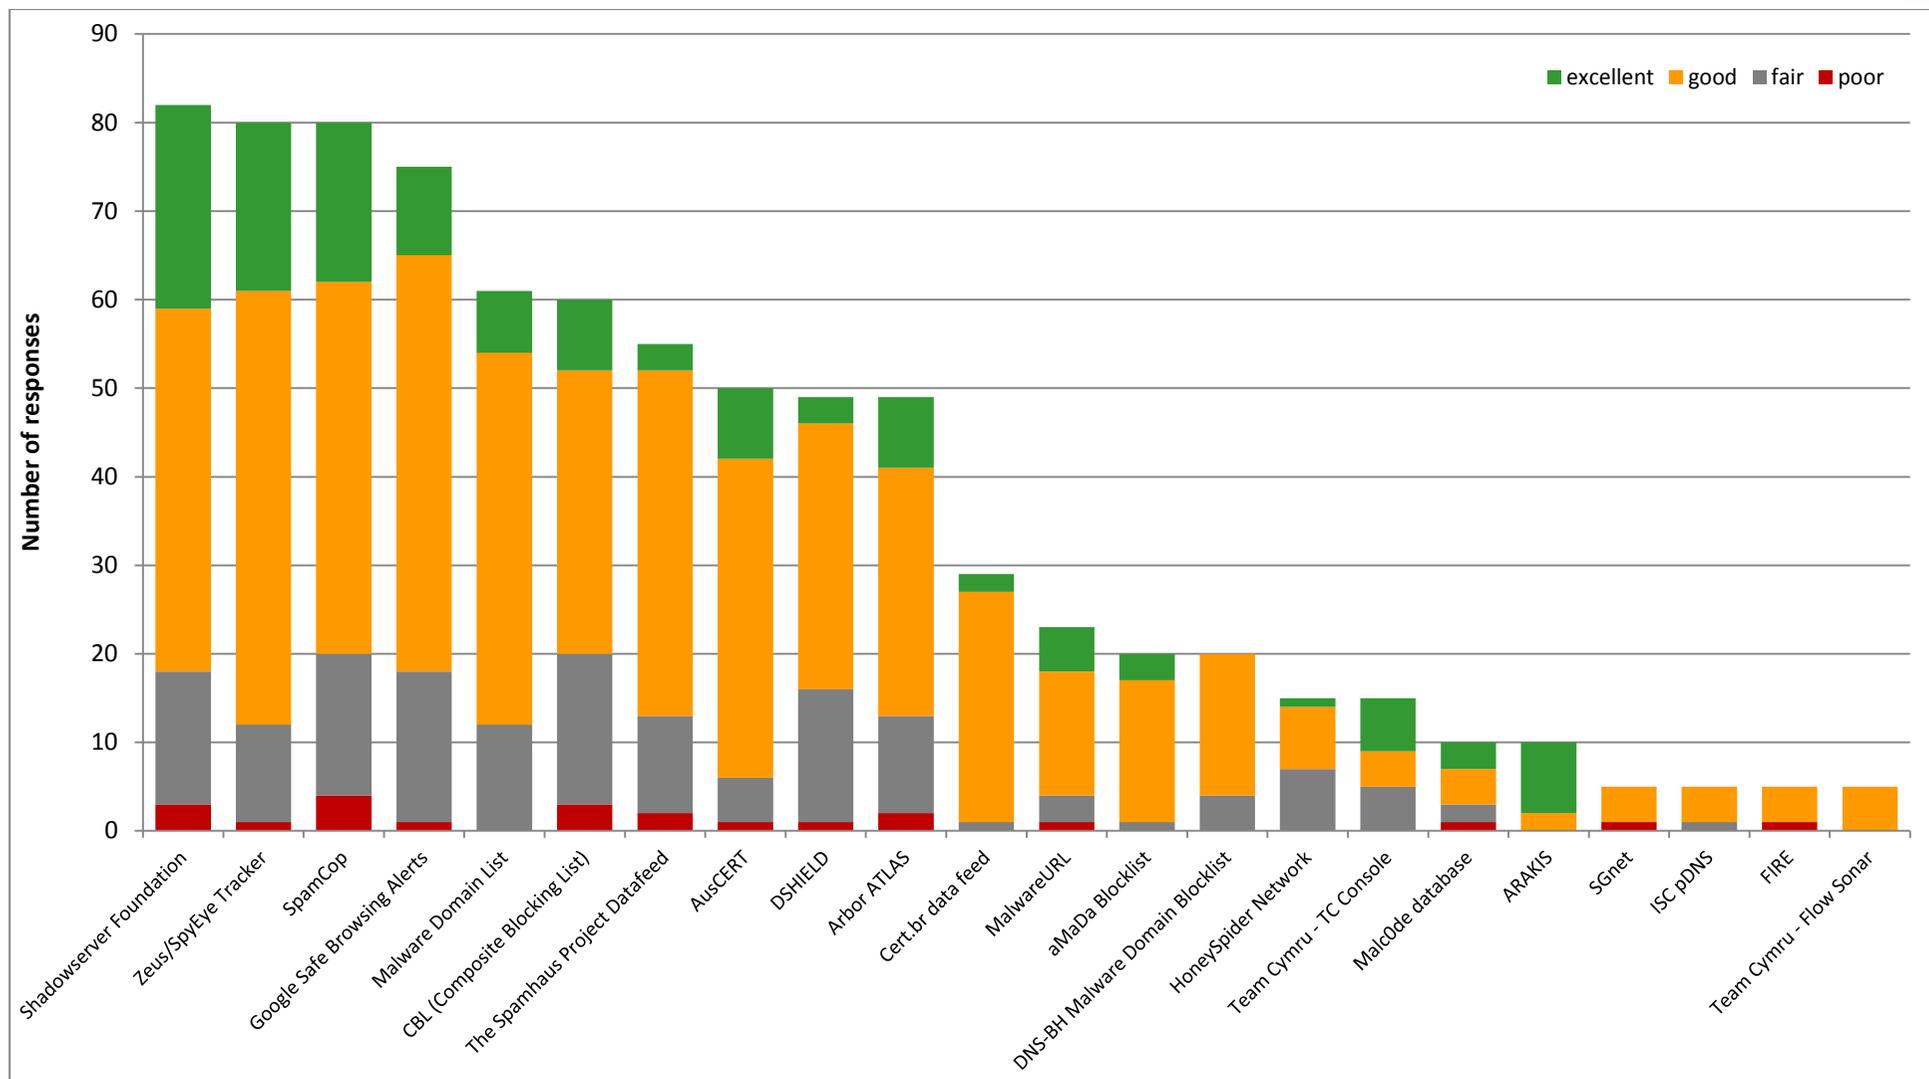


Figure 6: Evaluation of sources providing information on malicious or problematic URLs, IP addresses or domains

Figure 6 presents an evaluation of the most widely used sources of information. The initial list was created for purposes of the survey. However, respondents were able to suggest their own sources. Not surprisingly, the largest number of *excellent* and *good* grades went to the most used sources. Respondents were asked to judge criteria such as *timeliness*, *accuracy of results*, *ease of use*, *coverage* and *resources required* of every source that they use in a scale from *poor*, *fair*, *good* to *excellent*. Answers to this question were presented in the following manner: one bar represents the total number of responses (responses for *timeliness*, *accuracy of results*, *ease of use*, *coverage* and *resources required* are all summed up) broken down per grade. That is, every service was rated a maximum of five times, one for each category. Respondents graded only the sources that they are using. Note that many CERTs pointed to closed information sources as well, which cannot be named.

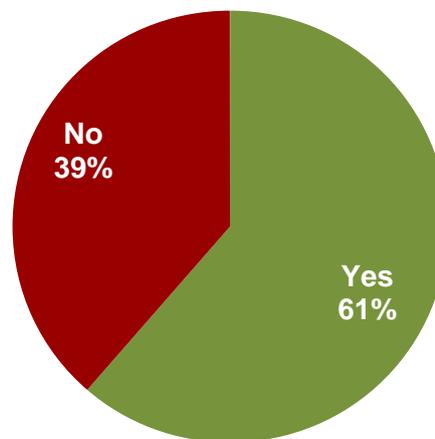


Figure 7: Do you use any closed sources of information you cannot disclose?

Sixty-one percent of CERTs use closed information sources to proactively detect incidents in their constituency. Respondents stressed in the comments that the accuracy, timeliness, reliability and quality of information gathered from closed sources is better in comparison to public sources. They state that they obtain information before it becomes available publicly.

4.4 Tools for internal monitoring

Internal monitoring tools; that is, tools deployed in networks that are directly under the control of a CERT or in networks that are part of the constituency, can be a valuable source of proactive information about incidents not only originating in own constituency, but also from other networks. These data can be exchanged with other CERTs and researchers. Therefore, we have included several questions evaluating application of different network monitoring tools for this purpose as well as asking about data sharing between CERTs.

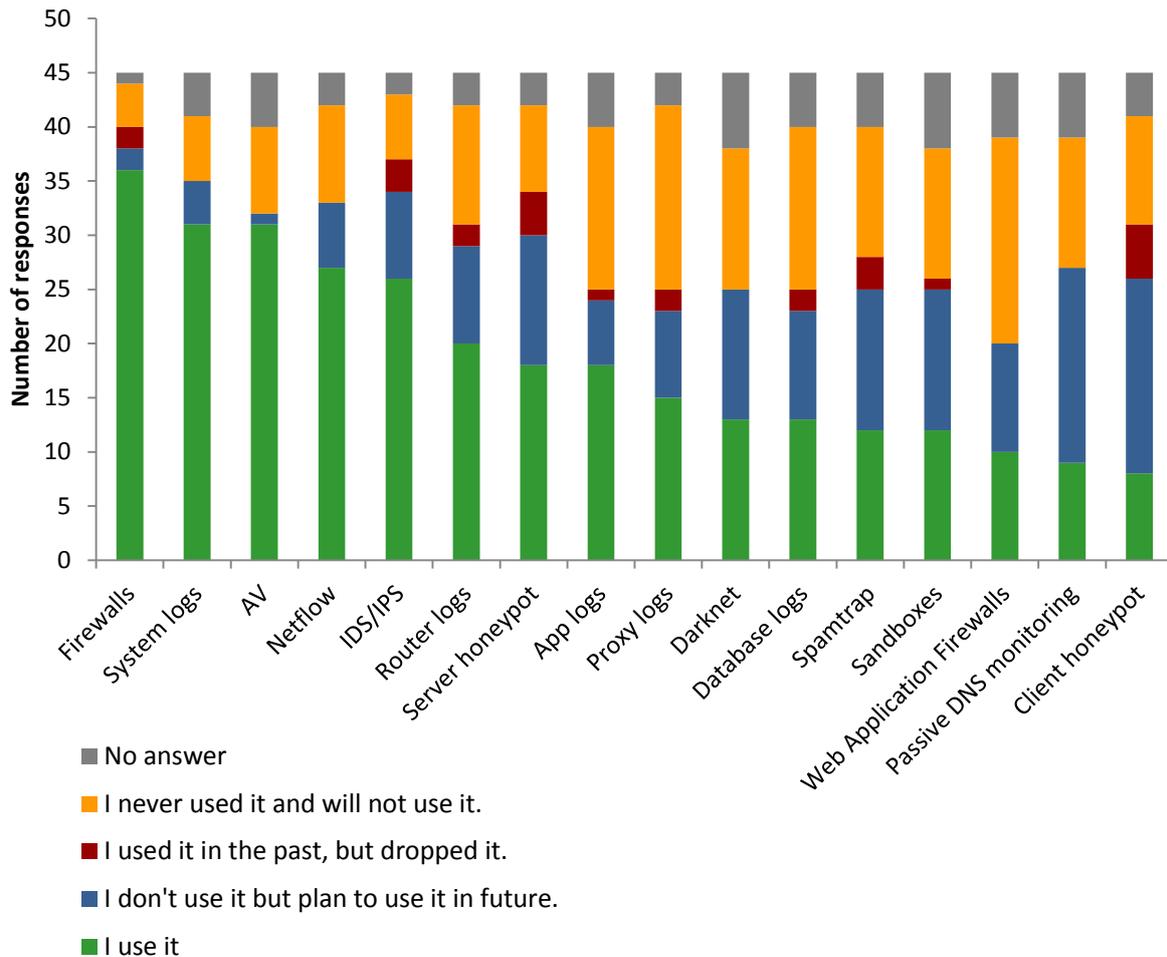


Figure 8: Categories of tools for gathering information from the network (see section 5.3 for more detailed descriptions of these tools)

Figure 8 above presents information about the use of tools for gathering information from the network (internal monitoring). Not surprisingly, firewalls, antiviruses and system logs are the most frequently used tools. The more sophisticated the tools are, the more workload and skills they require – therefore they are used less often.

One of the most interesting issues on this chart is position of the client honeypots, and an observation that many survey respondents pointed out both types of honeypots as a source of information they have used before but stopped. The question was therefore discussed during the workshop. One of the possible explanations given for this was that there is an additional workload involved in setting them up. Also, it was suggested that it is more interesting to see results from other honeypots than to deploy one’s own (preference to receive data about problems on one’s own network more than give data to others about their problems).

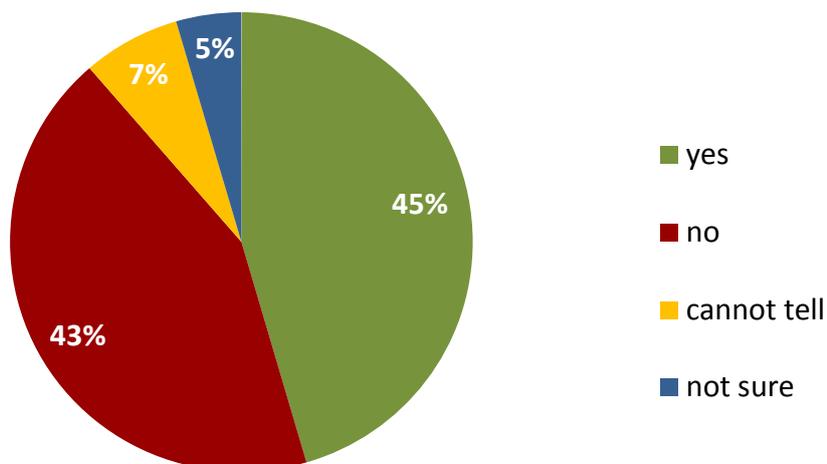


Figure 9: Do you collect information from your internal tools about incidents related to other constituencies?

Nearly half of the respondents claim that they do not collect information about other constituents in their internal tools. This is clearly an area of improvement for CERTs, as such information can be relatively easily collected with few resources invested.

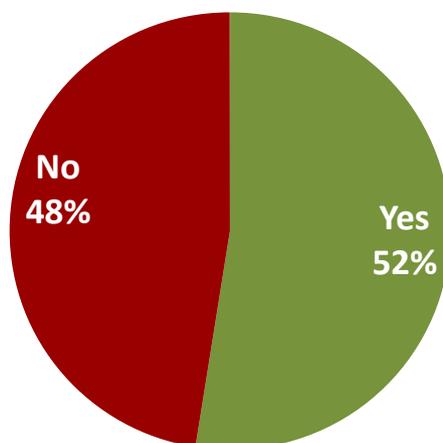


Figure 10: Do you share the data? (collected about other constituencies)

This question was only asked to those who admitted that they collect data about other constituencies. Interestingly, only half of them decided to share the data they collected. Almost half of the responding CERTs do not share data – either because they do not want to or are unable to share it for other reasons. From additional respondent information and workshop discussion it turned out that legal reasons can be a significant obstacle in sharing information. Experts pointed out

that in some countries privacy laws are very restrictive (for example, IP addresses are often classified as private data and are differently treated in different countries).

Furthermore, data that are shared require processing, storage and analysis. This is problematic for multiple respondents who lack the resources, both human and financial, to do so.

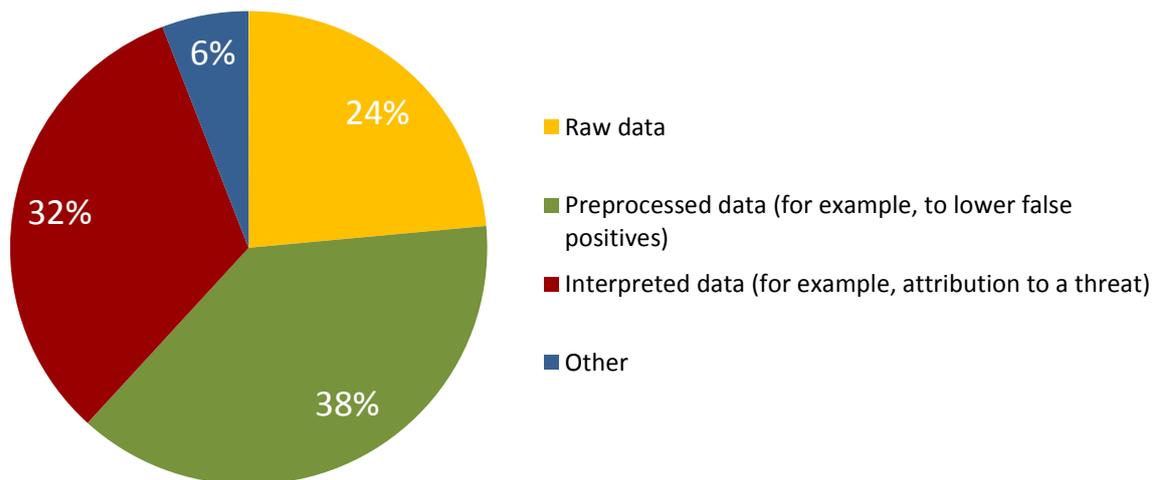


Figure 11: In what form do you share information with others?

Figure 11 presents the results of answers related to the way data are shared. More than one option could be selected.

A tendency to improve the data shared with others is clearly visible. Only 24% of answers reveal that respondents provide raw data. All the remaining answers (*preprocessed, interpreted and other*) show that CERTs make an effort to improve the data quality before passing it on further. This is done by lowering the number of false positives, internal correlation, interpretation, enriching the data, etc. Lowering the number of false positives has an influence on the perception of the quality of a source of information by receiving entities. It also allows more automation to be done on the receiving side. However, while sharing interpreted data is of course useful, using only such data reduces the chance of detection of more subtle or new malicious activities which may be otherwise detected if raw data were to be exchanged.

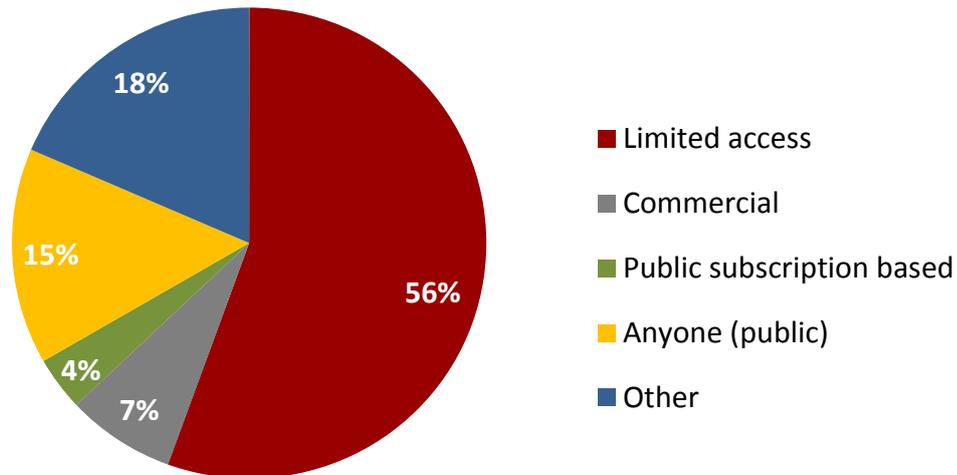


Figure 12: Under what conditions do you share information?

Answers show that in the CERT community, data are mostly shared under limited access. Limited access is necessary because these data usually contain sensitive or confidential information. This kind of information cannot be shared based only on a simple subscription model. The CERTs which share must have some assurance that the data they provide will be handled properly and not revealed to unauthorised parties. CERTs usually share with trusted partners, whose reliability is verified by the community.

4.5 Resources available

We asked the respondents about their ability to handle the current amount of information concerning detected security incidents.

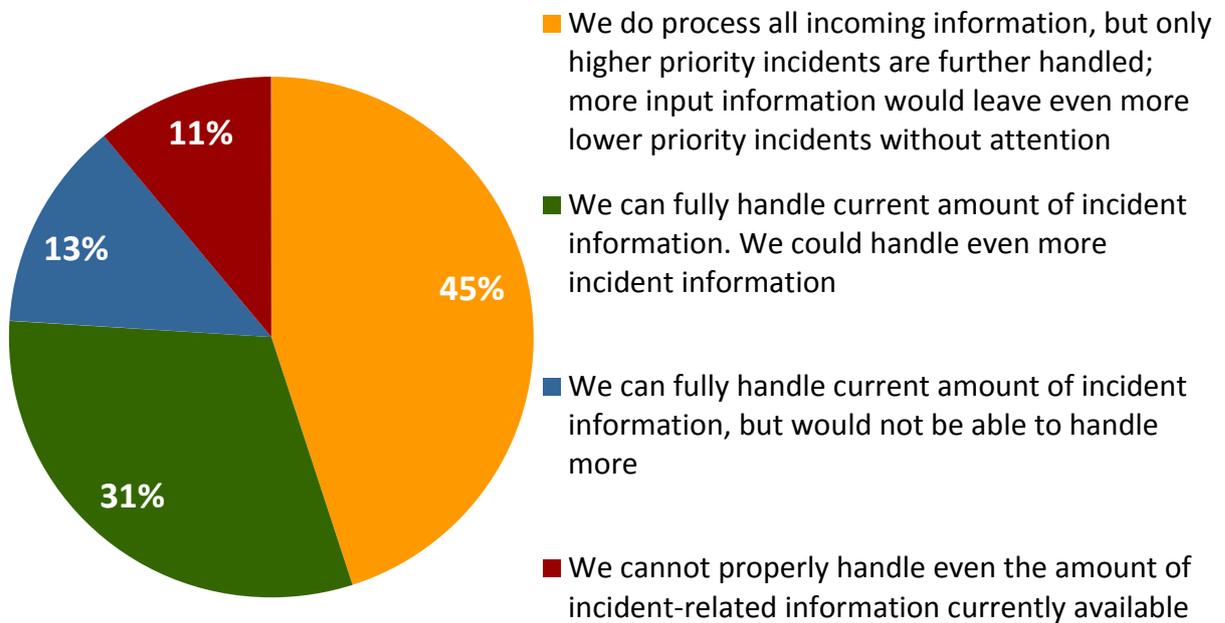


Figure 13: Resources available

The majority of respondents (89%) are capable of handling all the incoming information. Eleven percent of respondents claimed that they are not able to handle even the current amount of incident information. Among the majority the opinions on capability of processing differ.

- 45% claim that they process all incoming information, but only higher priority incidents are handled further. The explanations here were that there are insufficient human resources to process all information, hence in most cases the solution was to focus on most critical events.
- 31% claim that they can handle all information and would be able to handle even more. Additional comments on these answers give the impression that this option was chosen by flexible CERTs that are open to automatisisation, optimisation of the incident handling process and have enough resources to enable this.
- 13% claim that they cannot handle more information than they do now.

Interesting observations that emerged from additional analysis concerning follow-up of this question were:

- CERTs that claimed they are fully satisfied with the sources of information they already have (4%) often claim that they have enough resources to handle more. What is worth mentioning

is that these CERTs point to reactive reports or internal monitoring as a main source and do automatic correlation.

- Three CERTs answered that they would try new sources despite the fact that they cannot handle the current ones. This may be interpreted as an eagerness to take on new sources in the hope of better data quality.
- All the CERTs that cannot handle the current number of incoming incidents have problems with automatisisation. They either do not do any correlation at all or only partly automatically; moreover they point out problems with existing solutions.

In additional comments, respondents put emphasis on the variety of data formats of sources and thus problems with their proper processing: analysis and correlation. CERTs underscore that it always takes time and effort to adopt a new source. There were numerous comments stressing the need for tools for good automatic processing and unification of data formats.

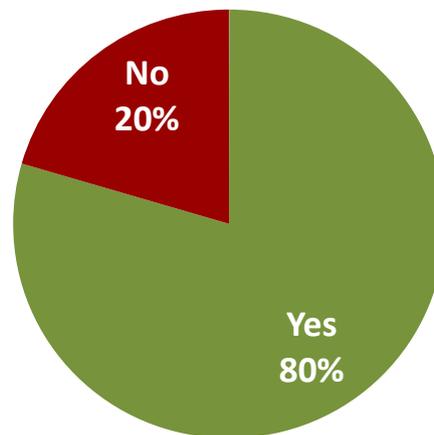


Figure 14: Do you correlate information from multiple sources in order to generate/confirm incidents?

The significant majority (80%) of CERTs correlate network security incident data. This influences the quality of data in a very positive way. Information about a particular incident becomes more complete and exhaustive.

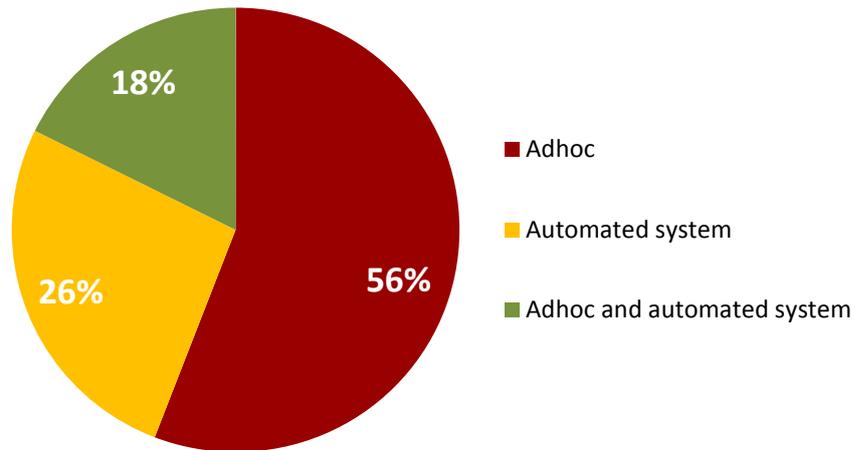


Figure 15: How do you correlate information from multiple sources?

Fifty-six percent of respondents claim that they correlate the information in an ad hoc way, which usually means a high additional requirement for human resources. 44% claim to involve automatization in the correlation process, which contributes significantly to a more effective and scalable incident handling process.

5 Inventory and description of identified services and tools

This chapter contains an inventory of identified services and tools (jointly called measures) for the proactive detection of network security incidents. The list is a result of desktop research by the study group, their direct experience with some of the services/tools, survey responses and discussion with experts. The tool section contains categories of tools with specific examples, rather than a full list of specific existing tools. Furthermore, the term ‘tool’ is understood to include protocols and mechanisms (such as netflow).

The study has also attempted to evaluate the measures listed in the inventory based on different criteria (discussed below). This is an extremely difficult process, somewhat like comparing apples and oranges. Nevertheless, the grades may be indicative of the usefulness of a particular measure. While definitions of each criterion provided are introduced, it must be remembered that the grades **given are essentially subjective in nature and to be viewed as suggestions only**. Note that services in particular can change over time. The grades given reflect the situation at the time of writing this report (September 2011).

5.1 Evaluation criteria

The services and tools described in the inventory are evaluated based on the following seven criteria: timeliness, accuracy, ease of use, coverage, resources, scalability and extensibility (which is applicable to tools only). The criteria set was chosen by the research team to provide some means of comparison between services and tools respectively. Grades used are: poor, fair, good and excellent. Please note that assigned ranks are a subjective evaluation of features and functionality of listed services and tools, based on expert judgement, survey responses and work group expert opinions.

5.1.1 Timeliness

The ‘timeliness’ criterion of service describes the length of time between the observation of the events and delivery of data by the service. The time of event is taken from the data feed as delivered by the service or tool and compared to the time of data retrieval. Based on this comparison a delay in notification about the incident is calculated which then directly influences the timeliness rating. There are four proposed classifications for the timeliness criteria:

Poor

Received data are more than seven days old from the actual date of the incident. The poor classification is assigned to services which provide mostly historical data, not fit for incident handling but useful for statistical analysis.

Fair

Received data are more than 24 hours and less than 7 days old. Services in this group provide data mostly on a daily basis, but report incidents from the previous day or before. The data source with

such a rank for timeliness is somewhat useful for incident handling, but cannot be an exclusive source of information.

Good

Received data are less than 24 hours old. Services assigned this classification deliver data with little delay from the time of actual incident occurrence and are potentially very useful for incident handling purposes.

Excellent

Data are received from service in near real time manner. Such information sources allow best responsiveness to incidents occurring on network and are potentially of high value to any CERT.

Not applicable

The 'N/A' rank is assigned only when there is no possibility to measure timeliness of a service and assign objective classification, or in case such information is not available or not important in context of the data that the service provides.

5.1.2 Accuracy of results

The 'Accuracy of results' criterion describes the quality of the service data feed. The more false positive results are delivered by the service, the poorer the quality of data and therefore the rank is lowered. The accuracy of delivered data is determined by expert judgement, in some cases previous experiences, survey responses and work group expert suggestions.

Poor

Data delivered by the service are of very poor quality and are almost unusable for effective incident handling process. Such services do not perform any data filtering, enrichment or correlation to filter-out false positive classifications.

Fair

Data delivered by the service are not meant to be used as a sole source of information about incidents but rather as an enrichment of existing ones.

Good

Data delivered by the service in most cases are ready to be used for incident handling purposes. The rate of false positive classifications is fairly low and the service can be used as one of main information sources for CERTs, though with minimal human supervision required to recheck received data occasionally.

Excellent

Data received from a service can be fully trusted. Services with this rank produce almost no false information and can be used as one of the main information sources for incident handling operations.

5.1.3 Ease of use

The 'Ease of use' criterion describes how easy is to access and use data delivered by the service. A low value for this criterion means it is harder to get access to data and parse it, mostly because some non-standard technologies are used. Higher rank means less work is needed to interpret results by a human or reformat the data to human-readable form.

Poor

Data received from service are either very hard to parse or very difficult to obtain, which renders the service hardly usable in most cases. Access to a service is restricted and/or requires sophisticated methods or tools to collect the data. Format of the fetched data is hard to parse and/or dependent on use of specific tools. The service lacks documentation, which renders the service very difficult to use.

Fair

It is hard to obtain the data or to parse them. Either access to the service is restricted or format of the fetched data is hard to parse. Service may require sophisticated methods or tools to either collect or interpret the data. The service lacks decent documentation.

Good

Service delivers data in a commonly used way and the format is easy to understand and interpret. The service does not require registration or this process is straightforward. It provides proper documentation.

Excellent

Service delivers data in a commonly used way and its format is easy to understand for a human reader and easy to parse by an automation mechanism. The service provides comprehensive documentation and support.

5.1.4 Coverage

Evaluation based on this criterion takes three aspects into consideration: constituency network range monitored by the tool or service (is it a part or the whole of constituency network?), incident detection/collection infrastructure (for instance, are all data collected from one sensor or is there a large network of sensors used – or perhaps the whole Internet is being crawled for malicious URLs rather than specific regions) and whether the service delivers different types of incidents (for instance, is it just spam, or specific botnet C&C servers, or malicious URLs or all of these). Essentially, coverage is about 'false negatives' – not in the accuracy sense, but in a broader sense, independent of incident type. Note that while the fact that a service or tool can detect/report different types of incidents is considered a plus, it does not diminish the rating for services that focus on just one type of incident (and do it well).

Poor

Service or tool delivers data concerning only small fragments of constituency network. The incident collection infrastructure is small.

Fair

Service or tool delivers data with significant gaps in regard to constituency network coverage. Incident collection infrastructure is medium.

Good

Service provides data feeds covering most of the constituency network address space. Incident collection infrastructure is fairly large. Data obtained from the service are considered highly usable for incident handling purposes and service can be used as a very good source of information about events occurring on the network.

Excellent

The whole address space of constituency network is covered by the data feeds. Incident collection infrastructure is large.

5.1.5 Resources required

The 'Resources required' criterion tries to describe the implementation impact from the financial, technical and human resource perspective. Lower rank means that the service is harder to implement and the cost of maintenance is higher.

Poor

Service requires vast resources for both successful implementation and usage. The cost of obtaining data from the service in relation to its usefulness is on the edge of viability.

Fair

Service requires significant resources for successful adoption. The expense of implementation is high but acceptable in relation to usefulness of data the service delivers.

Good

Service requires significant resources for successful adoption but usage is less demanding. Implementation is still considered costly but data retrieval and processing do not impose high expense.

Excellent

Service requires few resources and time for successful implementation and usage.

5.1.6 Scalability

The 'Scalability' criterion is applied only for tool evaluation and characterises the ability of the tool to handle growing data volumes and network growth. The tool is called scalable when it is capable of working properly with different volumes of data. The lower the value the less scalable the tool.

Poor

Tool is not scalable at all. It cannot cope with data/network growth and it might become unusable with a large amount of data.

Fair

The tool is not handling data/network growth well, but it can be configured/customised to do so.

Good

The tool is capable of managing data/network growth, but it has some minor limitations.

Excellent

The tool operates just as well with a small amount of data/network as with a large amount.

It can work properly within a small network as well as a large one.

5.1.7 Extensibility

The 'Extensibility' criterion is applied only to the evaluation of tools and defines the aptitude of the tool to extension of its functionality. This feature refers to the use of additional plugins or modules, and development of particular software in order to fit to one's needs.

Poor

The tool has a closed architecture and closed source. It does not support additional plugins. Extension is almost impossible.

Fair

The tool has a closed architecture, but its source is open and thus can be extended with significant effort and custom coding.

Good

The tool has a modular architecture, some modules are available. Writing new ones might require significant effort.

Excellent

The tool has a modular architecture and it can be easily extended. Vendor provides many plugins and/or writing custom ones is easy.

5.2 Services for the proactive detection of network security incidents³

5.2.1 DNS-BH Malware Domain Blocklist⁴

Description

Description of the service is based on information presented on the service website.

The DNS-BH Malware Domain Blocklist service provides information about malicious domain names responsible for propagating malware on the Internet. The project maintains lists of known malicious domains and provides DNS operators with zone files allowing fast and easy deployment on their networks. Service also provides information about classification source of each blacklisted domain. The data shared by the service can be used to create long-term filters allowing monitoring of the traffic and creating alerts when users try to access blacklisted content.

Evaluation of service

Timeliness

The DNS-BH Malware Domain Blocklist service updates information about malicious domain names every three days, and thus has limited monitoring functionality.

Accuracy of results

The service gathers information about malicious domains from sources such as Google Safe Browsing, mal0de.com database, PhishTank and many others. Because the service aggregates information, the quality of provided data depends strictly on information sources. The accuracy of results can be weakened because of a three-day delay between updates and may sometimes lead to false positive classification of a domain. The service provides information on dates when the threat was first observed and of last verification. Nevertheless, data receivers ought to verify the data which are considered outdated.

Ease of use

The service provides information about malicious domains in BIND,⁵ Microsoft Boot⁶ and TSV formats. Therefore it is very easy to deploy and maintain a DNS blackholing system. The service can also be used for monitoring a constituency DNS space checking for blacklisted domain names. The service provides classification reason for each domain allowing verification of data correctness.

Coverage

The service does not differentiate malicious domains, but provides information on classification reason which puts incident in proper context which can be used by data receiver to create different incident types. Because of the many information sources that the service integrates, its incident

³ Note that the ratings presented here were proposed by the authors of the study – members of the CERT Polska team

⁴ <http://www.malwaredomains.com>

⁵ Internet Systems Consortium, BIND: <http://www.isc.org/software/bind>

⁶ Microsoft Support, The Structure of a DNS Boot File: <http://support.microsoft.com/kb/194513>

collecting infrastructure can be considered large and coverage of the observed malicious domains is not limited to any specific region.

Resources required

Resources required to utilise the data delivered by the service depend on the use case. It can be a self-implemented monitoring solution or a DNS blackholing system based, for example, on BIND. Overall use of the service does not impose high requirements on resources, whether human or hardware.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Fair	Data feed is updated every 3 days, thus information provided could be older than 24 hours but should not be older than one week.
Accuracy of results	Good	The accuracy of data is assured by trusted external sources.
Ease of use	Excellent	Information is provided in two standardised formats. Additionally the service provides data in a third, easy-to-parse feed.
Coverage	Excellent	Provided data are not limited to any particular networks. Service provides different data feeds in different formats with contextual information allowing users to extract different types of incidents. The service utilises many incident information sources; thus its collecting infrastructure can be considered large and distributed worldwide.
Resources required	Excellent	It is very easy to implement a monitoring or blocking solution based on information provided by the service.

5.2.2 MalwareURL⁷

Description

Description of the service is based on information presented on the service website.

The service provides a database containing information about recently observed malicious activity on the Internet. It provides information for both non-commercial and commercial use. Registration is required to gain access to the service database. The database contains up-to-date information about newly observed malicious URLs, IP addresses the domain name resolves to and AS numbers to which the servers belong. The types of malicious activity covered by this service include phishing attacks, C&C servers, web pages containing exploits, scams, trojan infections and many others. The service provides data on its web page or in easy-to-parse CSV format and RSS feed. The service requires

⁷ <http://www.malwareurl.com>

registration to acquire access to data feeds. It is free to test for first 30 days and possibly free for non-commercial usage.

Evaluation of service

Timeliness

The service is updated daily with new information and retains old information in the database. The 24-hour period of updating is efficient enough to build an internal system of monitoring of constituency network and assure quick reaction time in case a new malicious activity is discovered.

Accuracy of results

MalwareURL uses proprietary software as well as widely recognisable services as sources of information and self-developed mechanisms which ensure high quality of gathered data. It also monitors changes in the behaviour of already observed malicious domains and delivers information about it. The short time between updates assures the information is always current.

Ease of use

Monitoring of constituency network will involve parsing provided CSV files or checking for updates on the RSS feed. The information can be filtered using AS numbers, thus reducing the dataset to only constituency-related entries. IP addresses and AS numbers enable the source of the malicious activity to be quickly determined so that proper actions can be taken to mitigate it. Provided data are straightforward and do not require additional investigation or analysis and can be used for alerting.

Coverage

The service does not differentiate detected malicious activity based on criteria such as location. The service delivers information as CSV and RSS feeds concerning various types of threats on the Internet such as scams, phishings, possible trojan infections, exploits and many other. The service seems to have a big incident collecting infrastructure. The assessment is based on the amount of data delivered each day.

Resources required

Resources required to implement a solution based on the information produced by the service are fairly low because of the format of supplied data. Information provided as RSS feeds and CSV files allows easy parsing and reprocessing of data.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Service updates data in 24-hour periods.
Accuracy of results	Good	The service uses known and trusted sources of information.
Ease of use	Excellent	Data are provided as ready to use and in easy-to-parse formats.
Coverage	Excellent	The service monitors all detected malicious services with no exception to any particular network. The range of observed incident types is broad. The service seems to possess a wide incident collecting infrastructure.
Resources required	Excellent	It is very easy to implement a solution employing data feed from the service for monitoring purposes.

5.2.3 Dshield⁸*Description*

Description of the service is based on information presented on the service website.

DShield is a free and open service that provides a platform for users of firewalls to share intrusion detection information. Any party can submit firewall logs to DShield database. The service provides summarised reports on observed attacks in the form of text files showing trends in attacks on particular ports, top 100 of attacking IP addresses, all observed IP addresses making attacks and a blocklist of IP address ranges. Because of heterogeneous sources of information, the data provide information only about an attack taking place without the context, such as attack vector being used. Interested parties can use the data to create rankings based on acquired information and to correlate trends observed on their own networks with the worldwide picture. The rankings give a more general view of attack, which is usually sufficient to determine its severity.

The blocklist provided by the service can be implemented as a filtering mechanism and used to produce alarms when traffic from blacklisted IP addresses is observed. Also lists of malicious IP addresses are a valuable source of information about malicious activity occurring from within the constituency network and targeting the outside world.

Evaluation of service

Timeliness

The time between updates of the data feeds delivered by the service differs depending on the type of feed. The trends for attacked port numbers are updated in real time. Feed for top 100 attacking IP addresses is updated every hour or so and list of all attacking IP addresses is updated in 24-hour

⁸ <http://www.dshield.org/>

periods. The blocklist of subnets is updated every hour and contains information from the last three days.

Accuracy of results

False positive classifications occasionally occur in data feeds. Because of that the DShield service is meant to be used as an additional source of information about already observed malicious traffic and cannot be regarded as a sole source of information about them.

Ease of use

The service provides data in easy-to-parse form of TSV files. The party interested in using provided data for monitoring purposes has to filter provided lists to extract information considering only the constituency set of IP addresses.

Coverage

The system utilises information from Firewalls and IDS distributed around the world (hence a large network collection infrastructure); therefore information about attacks considers a wide range of networks.

Resources required

The party interested in taking part in the DShield project is required to supply log files from its firewall and IDS. For this purpose a wide range of clients, tailored for many different brands of equipment, have been developed which ease the process considerably. Implementing a monitoring solution based on information gained from the service should not be a problem because the data provided are straightforward and easy to process.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Timeliness depends on the data feed used and varies from three days to real time.
Accuracy of results	Fair	Service uses firewall logs to build its data feed and because of this false classifications may occur.
Ease of use	Good	Data are provided in easy-to-use TSV format. Information about constituency needs to be extracted from the data feed.
Coverage	Excellent	System uses information from volunteers posting data from their firewalls distributed around the globe.
Resources required	Excellent	Data are easy to process and interpret.

5.2.4 Google Safe Browsing Alerts⁹

Description

Google Safe Browsing Alerts is a service that relies on Google Safe Browsing database of malicious websites. It allows network administrators to receive notifications with information of malicious content that is being hosted on their network.

Evaluation of service

Timeliness

The service updates its database constantly, but fresh information about a particular website seems to be supplied in an irregular fashion. The higher the website seems to be in the Google PageRank the more frequently it is scanned by the robots searching for malware. It is impossible to determine the exact time between the moment a compromised website starts serving malware and it being registered on the Safe Browsing list. It seems the time between successive scans varies from a day to a week. The database keeps a history of infections of a website showing scan results from the previous 90 days. It would seem the alert concerning a malicious website is generated the moment it is entered on the Safe Browsing list.

Accuracy of results

Although Google does not guarantee the accuracy of results and classification, the Google Safe Browsing is implemented in various software including the most popular web browsers. It is required to verify information received via Google Safe Browsing Alerts to be sure that no misclassification took place.

Ease of use

It is very easy to enrol for the service. The AS administrator is required to validate contact email address and ownership of the autonomous system to receive notifications. Alerts are sent automatically on the verified email account.

Coverage

The service alerts about malicious websites exploiting users as well as ones reported to commit phishing scams. Considering the fact that the Google search engine is one of the sources of information, the collection infrastructure is huge. Network administrator declares AS numbers for monitoring which belong to the network of constituency. The service sends only information concerning declared AS numbers.

Resources required

An email account is needed to receive notifications. The service requires information about networks which belong to organisation and it verifies ownership by sending verification email to one of AS contacts. Alerts will be received on separate email account provided during registration. Information received in the alerts usually needs to be verified. Verification whether a website is malicious is

⁹ <http://safebrowsingalerts.googlelabs.com>

usually not a trivial task and a highly skilled security expert is needed. The process can be speeded up by using existing software solutions aiding the verification process.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Data are updated in an irregular fashion. Some information can become false if update is delayed.
Accuracy of results	Fair	Information received in alerts can contain some false positive classifications mainly due to data aging.
Ease of use	Good	Alerts are sent as emails in easy-to-understand form.
Coverage	Excellent	Registering requires providing AS numbers for monitoring, therefore whole constituency network range can be covered by the service.
Resources required	Good	An email account is needed to receive notifications. The organisation is required to own the network in order to receive notifications.

5.2.5 HoneySpider Network¹⁰

Description

The HoneySpider Network service provides information about URLs classified by the system of client-side high- and low-interaction honeypots. The service performs periodical scans and generates alerts containing brief information about classification type (malicious, suspicious, benign), reason and context of scan. More elaborate results are available in the form of text reports or via web GUI interface.

The HoneySpider Network toolset was distributed among interested parties and allows them to create independent services providing help in incident handling for their organisations and constituency. It can also be used as a service that is provided by CERT Polska (you can get information by sending mail to info@cert.pl). The HoneySpider Network is also part of the WOMBAT project.

Evaluation of service

Timeliness

The service performs periodical scans on URLs. The time between scans can be defined by the operator of the service with granularity of one hour. The alert notifications can be sent either instantaneously or in batches – at a predefined period of time – depending on how the operator configures the service.

¹⁰ <http://www.honeyspider.net>

Accuracy of results

The service is based on the HoneySpider Network tool, consisting of client-side honeypots. It is a research project aiming at developing new methods of detecting malicious content served via the WWW network. Due to the project's research nature, the accuracy of provided data needs to be considered limited especially in terms of false positive classifications, but the spectrum of detected threats is much broader than in competing solutions, providing high true-positive classification rate.

Ease of use

The service provides information about the classification of scanned URLs as email or RSS alerts, each containing classification type, reason and context of scan process. The data are in human-readable format, but still easy to parse by automated systems. More elaborate information about classification of an URL is available via API called WAPI or as text report but only in human-readable form and requires expert knowledge to fully understand the data.

Coverage

The system scans a variety of URLs. Most of them come from spamtrap systems and lists prepared by experts mainly for monitoring of governmental and critical websites. It is possible to supply a custom list of URLs and receive alerts only for alarms generated in the context of the list processing.

Resources required

Receiving alerts requires only an email account. Taking full advantage of information provided by the service requires an expert able to interpret received results and advise the Incident Response Team in taking proper actions to mitigate the threat.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Service can perform scans with granularity of one hour.
Accuracy of results	Fair	Service is based on research project, therefore some false positive classifications may occur but data delivered for detected malicious websites are ready to be used as incident report base.
Ease of use	Good	Notifications about detected malicious URLs are provided as email alerts in easy-to-understand form, but may nevertheless require additional interpretation.
Coverage	Fair	Service monitors URLs provided as a list and does not crawl websites in depth. The coverage depends only on extensiveness of provided lists. If using the service provided by CERT Polska, URLs checked may be focused on Poland or from Polish spamtraps.
Resources required	Excellent	Receiving alerts requires an email account or an RSS reader.

5.2.6 AusCERT¹¹

Description

Description of the service is based on information presented on the service website.

The AusCERT organisation provides services for the public, its members and the higher education sector. It is a full member of the international Forum for Incident Response and Security Teams (FIRST), and Asia Pacific Computer Emergency Response Team (APCERT), which gives it access to accurate, timely and reliable information about emerging computer network threats and vulnerabilities on a global basis. It offers a wide range of services. For the purpose of this study, three services are considered: Malicious URL feed, Remote Monitoring and Incident Management Services for registered members. The Incident Management Services are both proactive and reactive and can be used to obtain assistance on responding to ongoing computer attack. As a globally recognised party AusCERT often coordinates incident handling between its members or delivers information about incidents and helps with investigation. Becoming a member within the organisation requires registration and providing information about autonomous system numbers the organisation owns. The malicious URL feed delivers information on phishing, malware, logging sites used to capture information from computers compromised with malware or fraudulent websites and mule recruitment websites. It can be used to build blacklists and filtering mechanisms in order to protect members of the constituency network. Remote Monitoring is a service that provides trusted mechanisms to inform about interruptions in operations of members services as well as possible intrusions (e.g. defacements).

The AusCERT provides information about incidents mainly to its members, but also forwards it to national CERTs or other recognised institutions. To obtain full cooperation on incident handling and alerting it is necessary to become a member.

Evaluation of services

Timeliness

The services are operational full time. The malicious URL feed is updated every 24 hours.

Accuracy of results

The services are globally recognised as delivering accurate and reliable information about incidents and emerging threats.

Ease of use

The feed is available in two types: a list of malicious websites from the previous 24 hours and a combined list of malicious websites from the last 7 days.

Remote Monitoring sends notifications as emails or SMS mobile messages about changes in the monitored services such as intruder activity or network interruption.

¹¹ <http://www.auscert.org.au>

Information concerning proactive detection of incidents from the Incident Management Services is usually provided via email or telephone contact.

Coverage

The services rely on information gathered by the AusCERT coordination centre, which monitors Internet threats and vulnerabilities from numerous sources throughout the year. There is no publicly available information on distribution of monitoring sensors.

The Remote Monitoring service coverage is limited to the list of services defined by the member.

Resources required

Receiving the SMS alerts requires an operational mobile phone. Also a delegated person is needed to interpret and handle the alerts.

Lists of malicious URLs are available for download at the AusCERT website. An automated system can be employed to download lists daily. The information can be easily reprocessed into blacklists preventing users from visiting malicious websites or to detect ones placed in the network of organisation.

The Remote Monitoring service in its basic form only sends email alerts about possible interruptions in monitored services, so a valid email account is needed to receive them and a delegated person to handle the alerts.

The AusCERT services deliver information to members mostly via cryptographically secured email communication. An interested party needs to become a member of the community and set up proper communication channels.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	The Malicious URL data feed is updated every 24 hours.
Accuracy of results	Good	Accurate and reliable information about incidents and emerging threats.
Ease of use	Good	Registered users can download malicious URL data feed from AusCERT website. Information on incidents is also provided via emails.
Coverage	Good	AusCERT monitors Internet traffic and services but registered users can provide AS numbers to receive information about incidents concerning specific networks.
Resources required	Excellent	Services provide information through email and SMS notifications.

5.2.7 Cert.br Distributed Honeypot Project¹²

Description

Description of the service is based on information presented on the service website.

The CERT.br has established a service which delivers information about network traffic observed by a system of honeypots distributed across Brazil. Information presented by the service is divided in categories like country of origin of traffic, protocol traffic and operating system. Also statistics about the most popular TCP and UDP ports are shared. The data can be useful for checking and correlating observed traffic in case of global-scale worm infections or attacks and monitoring of known malware sources.

It is also possible to receive more sophisticated data from the service as emails with information concerning constituency network. It requires establishing a contact with the organisation and agreeing on terms of data sharing. The data contain information about IP addresses suspected of malicious activity together with a number of events generated by the IP in the system of honeypots. Beside just the list of IP addresses, the data feed is extended with detailed information about flows the suspicious IP addresses have generated, such as date of event, source and destination ports, TCP flags and fingerprinted operating system of the attacker. The data feed can be used to monitor the constituency network for malicious behaviour which targets external networks.

Evaluation of service

Timeliness

The service delivers data on a daily basis presenting information from a 24-hour period preceding the current date.

Accuracy of results

The system uses a network of low-interaction honeypots (Honeyd) as means to gather incoming traffic. The Honeyd honeypot is software recognised by many security experts and often used as a reference.

Ease of use

The service is accessible via a website and delivers information as graphs and tables. Presenting the data as static graphs and tables hinders the automatic processing.

The data feed delivered by emails is fairly easy to parse and interpret. Additional information can be used by the operator to determine the severity of the attack and to filter out false positive alarms.

Coverage

The service delivers information about network traffic directed to honeypots from the Internet. On its web page it provides information about top 10 AS numbers from which such traffic originated. Such information can be useful to network administrators as a starting point for investigation of

¹² <http://honeytarg.cert.br/honeypots/>

possible infection on their network. The incident collecting infrastructure consists of over 20 sensors distributed in Brazil.

Information in the data feed concerns only IP addresses which belong to declared networks. The system does not deliver information about the threat type, just plain data about flows.

Resources required

Using the service requires only an Internet connection and a delegated operator to interpret presented information. Receiving the data feed requires an email account. The emails are encrypted with PGP keys, but this should not be much of an obstacle in automated processing of the information.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Data feed is updated every 24 hours.
Accuracy of results	Good	Service uses Honeyd to gather information on attacks.
Ease of use	Fair	The data are delivered through a web interface as graphs and tables which makes it inconvenient for automatic processing. The data feed concerning constituency networks is delivered as encrypted emails with information on flows which are meant for human interpretation.
Coverage	Good	Service provides information about attacks directed to a network of Brazilian honeypots, therefore data considering specific constituency is limited.
Resources required	Good	Service can be used for manual correlation of data with other sources. Automatic processing of data feed delivered by emails would require specialised solution.

5.2.8 FIRE (Finding Rogue nEtworks)¹³

Description

Description of the service is based on information presented on the service website.

FIRE is a service to identify rogue networks that persistently host malicious behaviour such as botnet command and controls servers, drive-by-download websites and phishing scams. The system tries to locate such networks and identify service providers responsible for keeping them. The service can be queried for specific information about an autonomous system. Details given are IP addresses of malicious servers, country where the server is located and type of malicious behaviour. Service keeps history of detected malicious networks allowing tracking changes in distribution of malware over

¹³ <http://www.maliciousnetworks.org>

time. Discovered information is presented to the public on the service website. Service is a part of the WOMBAT project.

Evaluation of service

Timeliness

The system provides new information every 24 hours.

Accuracy of results

The service gathers data from Anubis, Wepawet, PhishTank and HoneySpider Network, which are known and often referenced by security experts. The accuracy of data depends solely on those services as well as on algorithms implemented by the FIRE authors. However, survey feedback suggested a mix of good and poor results.

Ease of use

The service provides information on its website. Organisations can monitor their autonomous system numbers by querying the online database with their AS numbers which returns list of detected malicious IP addresses together with type of malicious activity. Unfortunately there are no other public methods of acquiring information from the service.

Coverage

The service coverage depends on monitoring and detection capabilities of its sources of information. Most of them provide public interfaces to submit data for analysis, thus greatly improving detection rate and coverage and creating a large incident collecting infrastructure. Service provides information about autonomous systems distributed worldwide. Information the service shares includes AS numbers of malicious networks, IP addresses of servers responsible for threats and their type.

Resources required

The service does not require any resources to gain access to information aside from Internet access and a modern web browser. Manual usage of service is easy and straightforward. Implementing an automated solution for monitoring organisations' AS numbers will require developing a specialised solution for data retrieval.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Service updates data every 24 hours.
Accuracy of results	Good	Data are gathered from various sources and processed with advanced algorithms to assure minimal amount of false positive classifications.
Ease of use	Fair	Service provides access to data through web interface which hinders automatic processing. It is not possible to post AS numbers for monitoring.
Coverage	Good	Service depends on various sources of information such as Anubis, Wepawet and HoneySpider, thus it has a distributed and fairly big incident collecting infrastructure. Service delivers information on malicious AS numbers, IP addresses and types of threat.
Resources required	Good	Manual usage of service is easy but automatic processing requires building a customised solution.

5.2.9 Team Cymru – TC Console¹⁴

Description

Description of the service is based on information presented on the service website.

Team Cymru – TC Console is a web-based user interface visualising malicious activity inside the organisation's network. The data are collected from the community of contributors to the project. The service aggregates the data and helps in analysis of observed traffic. It is also possible to fetch information from the service as TSV files which can be used for automated processing of incidents. The service provides details on malicious activity like open resolvers, proxies, machines distributing malware or phishing pages, brute force attacks, spam, compromised routers, and machines that are infected with malware or connecting to any known C&C server.

The service allows its users to share information about incidents. A user can submit information about an IP address and description of malicious behaviour that was observed inside the submitter's network. A moment later, the information is available as a incident report for the network administrator of the submitted IP.

Organisation has to register to join the community to be able to take information from the service and share incident reports. Registration involves declaring the AS numbers the party owns and is interested in receiving information about.

¹⁴ <https://www.tcconsole.com>

Evaluation of service

Timeliness

The service collects and presents data in almost real time. Depending on the information feed, the data are displayed from last 24 hours or 30 days. Older data are stored as history and kept for 90 days.

Accuracy of results

Team Cymru is a recognised organisation providing valuable services and information for many security specialists around the world. The accuracy of data is increased by feedback about incidents from the project partners.

Ease of use

The TC Console service presents data in an easy-to-use web interface and data feeds as TSV files. No software installation is needed.

Coverage

The service displays data gathered from sources distributed around the world. Presented data concern only the network the user owns and malicious activity that originates from it either discovered by Team Cymru sensors or reported by the community.

Resources required

Organisation willing to use the service needs just Internet access and a modern web browser. Automated incident-generation based on information from the TC Console service is possible with data delivered with TSV files.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Data are presented almost in real time. Depending on the datafeed it is possible to access information concerning last 24 hours or 30 days.
Accuracy of results	Good	Service gathers data from various sources. Accuracy can be increased if organisation is willing to share data gathered by internal monitoring.
Ease of use	Good	Service is very easy to comprehend and use. It provides a web interface and delivers information through convenient data feeds (TSV files).
Coverage	Excellent	Data come from distributed sources and provide broad overview of situation on the network. The coverage can be extended if organisation is willing to share data from its darknets and other sources.
Resources required	Excellent	Internet access and a modern browser are all that is needed to take full advantage of the service.

5.2.10 EXPOSURE¹⁵

Description

Description of the service is based on information presented on the service website.

EXPOSURE is a service that identifies domain names that are involved in malicious activity by performing large-scale passive DNS analysis. The service keeps a list of recognised malicious domain names and shares them as a blacklist. The mechanism of the domain removal from the blacklist is not clear from the description of the service. The service also stores historical data on detected malicious domains and allows tracking changes as they happen over time. The service is a part of the WOMBAT project.

Evaluation of service

Timeliness

The service publishes new data every 24 hours.

Accuracy of results

System delivers information about malicious domain names which are confirmed to be dangerous by services like Anubis, Wepawet, PhishTank and others. These sources of information are recognised to provide useful and accurate data concerning malicious activity on the Internet.

Ease of use

The service delivers a web interface to browse information and a search engine to query the service about a particular domain name. It also provides a list of blacklisted domains in the form of a text file. Results of a query for a particular domain contain information about history of observed DNS traffic considering the domain, list of IP addresses the domain resolved to, country code and autonomous system number. The data are presented in easy-to-interpret form of graphs and tables.

Coverage

The service monitors DNS traffic and receives additional data from Anubis, PhishTank, Wepawet and other systems; thus its data collecting infrastructure can be considered large. The services provide global-scale information, therefore EXPOSURE data feed has the same global scale of coverage.

Resources required

Accessing the service requires only an Internet connection and a modern web browser. It is also possible to automate the process of downloading the blacklist of domain names and using it for monitoring constituency network or creating a DNS blackholing service. In both cases a delegated person would be required to build and maintain the system utilising the service.

¹⁵ <http://exposure.iseclab.org>

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Service updates data every 24 hours.
Accuracy of results	Good	Data gathered by the service are verified with known and trusted tools and correlated with advanced algorithms but some false positive classifications are possible. Data-aging principles are not clear based on the description of the service.
Ease of use	Excellent	Manual usage via web interface is straightforward and easy. Service also provides list of domains for blocking or checking for entries from constituency as an easy-to-parse text file.
Coverage	Good	The service is not limited to monitoring any particular network range and its data collecting infrastructure is big. It delivers data through web-interface and text file of detected malicious domains. It does not provide reason for classification.
Resources required	Excellent	A delegated person is required to interpret data delivered by the service. The list of blacklisted domain names needs to be parsed and prepared for proper software but it is not difficult.

5.2.11 Zeus/SpyEye Tracker¹⁶

Description

Description of the service is based on information presented on the service website.

The ZeuS Tracker is a project by abuse.ch. The service tracks and monitors malicious ZeuS¹⁷ C&C servers, ZeuS configs, ZeuS binaries, ZeuS dropzones and fake URLs. It provides blocklists in different formats (e.g. for Squid Web-Proxy or iptables) preventing infected clients from accessing the C&C servers.

The SpyEye Tracker is similar to the ZeuS tracker, but focused on the SpyEye malware. It is also another project by abuse.ch. It tracks and monitors malicious SpyEye C&Cs and provides blocklists in different formats. Main lists categories are: SpyEye BinaryURLs, SpyEye ConfigURLs, SpyEye Dropzones. Additionally, SpyEye Tracker should help ISPs, CERTs and Law Enforcement to track malicious SpyEye C&C servers which are their responsibility.

¹⁶ <https://spyeyetracker.abuse.ch> , <https://zeustracker.abuse.ch>

¹⁷ SpyEye Bot versus Zeus Bot: <http://www.symantec.com/connect/blogs/spyeye-bot-versus-zeus-bot>

Lists contain only active IP addresses or URLs. Additionally separate lists of removed addresses are provided (hosts which were removed from the tracker). An entry from the list may be removed if organisation contacts the list administrator.

Both services provide a feature which allows to filter C&Cs for specified nameservers, status, AS number, IP addresses, etc. Additionally ZeuS Tracker provides DNS Service, which is similar to normal DNS blackhole lists. It is an advanced lookup to check an IP address or a domain name against the ZeuS Tracker IP or Domain Blocklist.

Evaluation of service

Timeliness

The database is probably updated in real time, but not regularly. It depends on present ongoing analysis or results of bot snooping.

Accuracy of results

The information is gathered from analysis and spying on real malware samples. This makes the quality of the information quite good. However there is no guarantee there will be no false positives.

Ease of use

The services provide information in four main formats: lists in text format, HTML format, application/system configuration files and RSS feed. Lists are sorted in many ways:

ZeuS Tracker:

- ZeuS domain blocklist
- ZeuS IP blocklist

application/system config files:

- ZeuS combined (IP addresses + domain names) blocklist for Squid¹⁸
- ZeuS IP blocklist for iptables
- ZeuS domain blocklist for Windows (Hosts-File)
- ZeuS combined blocklist for unix (hosts.deny)

RSS feeds:

- ZeuS Tracker RSS Feed (notification about new ZeuS hosts)
- ZeuS Tracker removal list RSS Feed (notification when a ZeuS host will be removed from the ZeuS Tracker)
- ZeuS Tracker config URL RSS Feed (notification about new ZeuS config URLs)
- ZeuS Tracker binary URL RSS Feed (notification about new ZeuS binary URLs)

¹⁸ Squid: Optimising Web Delivery: <http://www.squid-cache.org/>

- ZeuS Tracker dropzone URL RSS Feed (notification about new ZeuS dropzone URLs)
- ZeuS Tracker ASN RSS Feed (notification about new ZeuS hosts in the specified AS)
- ZeuS Tracker country RSS Feed (notification about new ZeuS hosts in the specified country)

Additionally ZeuS Tracker provides DNS Service, which is an advanced lookup to check an IP address or a domain name against the ZeuS Tracker IP or Domain Blocklist.

SpyEye Tracker:

- SpyEye domain blocklist
- SpyEye IP blocklist

application/system config files:

- SpyEye combined (IP addresses + domain names) blocklist for Squid
- SpyEye IP blocklist for iptables
- SpyEye domain blocklist for Windows (Hosts-File)
- SpyEye combined blocklist for unix (hosts.deny)

RSS feeds:

- SpyEye Tracker C&C RSS Feed (it contains the latest forty SpyEye C&Cs)
- SpyEye Tracker ConfigURL RSS Feed (it contains the latest forty SpyEye ConfigURLs)
- SpyEye Tracker BinaryURL RSS Feed (it contains the latest forty SpyEye BinayURLs)
- SpyEye Tracker DropURL RSS Feed (it contains the latest forty SpyEye DropURLs)
- SpyEye Tracker ASN RSS Feed (it notifies about new SpyEye C&Cs in the specified AS)
- SpyEye Tracker Country RSS Feed (it notifies about new SpyEye C&Cs in the specified country)

Coverage

The service provides information that is not focused on a specific world region (country or network). However services are focused on tracking only two malware families: ZeuS and SpyEye. About one up to tens of new incidents are generated daily. However, no detailed information about the collection infrastructure is available.

Resources required

Few resources, whether human or hardware, are required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	The database is probably updated in real time, but not regularly. It depends on ongoing analysis or snooping of infected bots.
Accuracy of results	Excellent	The information is gathered from analysis and snooping on real malware samples. This makes the quality of the information quite good. However there is no guarantee there will be no false positives.
Ease of use	Excellent	Many formats of data: lists in text format, HTML format, application/system configuration file and RSS feeds. Lists are sorted in many ways. Additionally service provides DNS Service, which is an advanced lookup.
Coverage	Fair/Good	From one up to tens of new incidents daily, only two types of botnet. The amount of collect samples is satisfactory for such a service.
Resources required	Excellent	Few resources, whether human or hardware, are required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

5.2.12 AMaDa¹⁹**Description**

Description of the service is based on information presented on the service website.

The AMaDa (abuse.ch Malware Database) is a service provided by abuse.ch. This service provides lists of C&C servers grabbed mainly from sandboxes. Monitored malware and botnets are of various types. IP or domain name of the C&C server lists are available. Additionally there is an option to search and monitor (i.e. by RSS feed) IP addresses, domain names or AS numbers.

Evaluation of service**Timeliness**

The database is probably updated in real time, but not regularly. It depends on ongoing analysis and snooping of bots.

¹⁹ <http://amada.abuse.ch>

Accuracy of results

Information is gathered from analysis and tracing of real malware samples. This makes the quality of the information quite high. However it cannot be guaranteed that there will be no false positives. The list can be sorted by add date (old/non-active data are not removed).

Ease of use

The service provides information about C&C servers in two different ways: in text files (lists), one file for domain names and another for IP addresses. The service does not differentiate both lists based on location, AS number, country code or TLD (Top-level domain). The CERT team interested in using lists needs to filter the data and parse only relevant information.

The second way is monitoring a particular IP address or AS number or domain name by RSS feed or web page (HTML format). All this data can be used to monitor CERT's constituency networks for association with particular botnet infrastructure. Another option is to create filters from the data to detect infected computers when they 'phone home' (make connection to the C&C servers).

Coverage

Monitored malware and botnets (mainly from sandboxes) are of various types. About 1–5 new incidents are generated daily (focus on C&C servers). No detailed information about collection infrastructure is available; however the collection infrastructure need to achieve such a result can be rated as reasonable.

Resources required

Few resources, whether human or hardware, are required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	The database is updated in real time, but not regularly. It depends on a present ongoing analysis or infected bots spying.
Accuracy of results	Good	The information is gathered from analysis and spying the real malware samples. This makes the quality of the information quite good. However there is no guarantee that there would be no false positives.
Ease of use	Excellent	Multiple ways: in text files (different file for domain names and another for IP addresses), and by RSS feed or web page (HTML).
Coverage	Fair	About 1–5 new incidents are generated daily – in terms of C&C servers. This implies that the collection infrastructure is reasonable.
Resources required	Excellent	Few resources required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

5.2.13 Malware Domain List²⁰

Description

Description of the service is based on information presented on the service website.

Malware Domain List is a service that provides a list of URLs that are dangerous (involved in infection process, botnet management or drop zone, malware hosting, etc.). Additionally a list of IP addresses of web servers is provided. It is a non-commercial community project and can be used for free by anyone. The service provides both a full list and updates lists.

Evaluation of service

Timeliness

The database is updated in real time, but not regularly. It depends on a present ongoing analysis or bots snooping. The service provides updates lists which are useful if a CERT wants to gain only the most up-to-date information. The full lists are not cleaned and contain both obsolete and up-to-date data.

Accuracy of results

The information is gathered from analysis and spying on the real malware samples. This makes the quality of the information quite high. However there is no guarantee that there would be no false positives.

Ease of use

The service provides information in two main ways: lists in CSV format or simple text, and second: RSS feed. Additionally lists are sorted in many ways:

- complete database (all fields) – full or only updates
- only URLs – full list or only updates
- ZeuS URLs – full list or only updates
- sites which are offline or have been cleaned – full list
- list of active IP addresses

Complete database lists have additional fields, which are useful in the parsing/filtering process, like IP address of the server hosting site, AS number, country code, registrant or reverse DNS lookup.

Coverage

The service provides information of a wide scope, not dedicated to a specific world region (country or network). Listed malicious URLs and IP addresses come from various sources (not only from one

²⁰ <http://www.malwaredomainlist.com/>

type of botnet or malware family). Up to several hundred new incidents are generated daily. However, no detailed information is available about the collection infrastructure.

Resources required

Few resources, whether human or hardware, are required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	The database is updated in real time, but not regularly. It depends on a present ongoing analysis or infected bots spying.
Accuracy of results	Good	The information is gathered from analysis and spying the real malware samples. This makes the quality of the information quite good. However there is no guarantee that there would be no false positives.
Ease of use	Excellent	Many formats of data delivery: lists in CSV format or simple text, and RSS feed. Additionally lists are sorted in many ways: complete database (all fields) – full or only updates; only URLs – full list or only updates; ZeuS URLs – full list or only updates; sites which are offline or have been cleaned – full list; list of active IP addresses. Complete database have additional fields, which are useful in parsing/filtering process, like IP address of the server hosting site, AS number, country code, registrant, or reverse DNS lookup.
Coverage	Good	Several up to hundreds of new incidents daily. This can be considered a pretty good result given the type of incident being detected. The service very actively uses submissions from the public, which gives a wide coverage.
Resources required	Excellent	Few resources, whether human or hardware, are required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

5.2.14 The Spamhaus Project (Spamhaus DNSBL Datafeed)²¹

Description

Description of the service is based on information presented on the service website.

²¹ <http://www.spamhaus.org/>, <http://www.spamhaustech.com/datafeed/>

The Spamhaus Project tracks the Internet's spam operations and provides a number of real time spam-blocking databases ('DNSBLs'), including the Spamhaus Block List (SBL), the Exploits Block List (XBL), the Policy Block List (PBL), the Domain Block List (DBL) and the Don't Route Or Peer List (DROP); or one big ZEN list that merges all lists in one. All lists are dedicated to mail servers to filter incoming emails.

Spamhaus also offers a powerful and professional service called Spamhaus DNSBL Datafeed. It is dedicated for users with professional DNSBL query requirements, such as corporate networks and ISPs. It offers two subservices: Datafeed Query Service, which provides real time access to a private network of Spamhaus DNSBL servers and uses traditional DNS; and Datafeed Rsync Service which provides rapid data synchronisations between Spamhaus DNSBL servers and local servers on client networks (recommended for high-volume users). Both services provide access to all Spamhaus DNSBLs. De facto only Datafeed services could be used as effective and useful in proactive incident detection.

The Spamhaus Project is an international non-profit organisation.

NOTE!

The Spamhaus Project has DNSBL Usage Terms – the free use is limited. Every organisation should check and accept these terms. To use the Spamhaus Datafeed Service a yearly fee needs to be paid.

Evaluation of service

Timeliness

The database is maintained in real time. Using Datafeed Query Service we have a guarantee that answers are up to date. Using Datafeed Rsync Service client has to regularly synchronise data with Spamhaus's servers. However the delay should not be an issue (recommended synchronisation time is every 5 minutes).

The SBL DNS zone is rebuilt and reloaded every 30 minutes, 24 hours a day. The XBL, PBL DNS zone is rebuilt and reloaded every 15 minutes, again throughout the day. The DBL DNS zone is rebuilt and reloaded every 60-seconds, 24/7. DROP is updated not more than once an hour (probably once per day).

Accuracy of results

All Spamhaus's services are maintained by a dedicated international team of investigators and forensics specialists. This makes Spamhaus reliable. However false positives are possible due to its aggressive rules in blacklisting spam bots.

Ease of use

The Spamhaus Datafeed Query Service is easy to use and set up – it uses traditional DNS queries. However the setup process requires some technical knowledge. Rsync Service is more complex to set up, but similarly easy to use. Setup process requires some advanced tools and more advanced technical knowledge.

Coverage

Only hosts involved in sending spam are listed (one type of security incident). Spamhaus is an established initiative and as such has a wide network of distributed sources for gathering of data.

Resources required

Technical knowledge is required to set up this service. Also hardware with specific software could be required (in case of Datafeed Rsync Service: DNS server).

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	The database is maintained in real time or near.
Accuracy of results	Good	All Spamhaus's services are maintained by a dedicated international team of investigators and forensics specialists. However, false positives are possible.
Ease of use	Good	Spamhaus Datafeed Service uses traditional DNS queries. However the setup process requires some technical knowledge. Rsync Service is more complex to set up, but easy to use.
Coverage	Excellent	Hosts involved in sending spam are listed (one type of security incident).
Resources required	Good	Technical knowledge is required to set up this service. Also hardware with specific software could be required (in case of Datafeed Rsync Service: DNS server). NOTE: The Spamhaus Datafeed Service is not free (chargeable)

5.2.15 Shadowserver Foundation²²

Description

Description of the service is based on information presented on the service website.

The Shadowserver Foundation is a volunteer watchdog group of security professionals that capture, gather, analyse, monitor, track and report on malware, botnet activity and electronic fraud. It also helps in coordinating incident response. The Shadowserver provides service that could be used in proactive incident detection.

The ASN & Netblock Alerting & Reporting Service allows organisations to receive customised reports about IP addresses from specified netblocks and ASNs that are involved in malicious activity. Service

²² <http://www.shadowserver.org>

is designed for ISPs, enterprises, hosting providers and other organisations that directly own or control network space (although national CERTs may apply for wider access). This service provides many types of reports.

The service is free of charge.

Evaluation of service

Timeliness

Shadowserver runs the reports starting every morning for the previous 24 hours (UTC time-based).

Accuracy of results

Service is maintained by an international team of investigators and forensics specialists.

This makes Shadowserver reliable. However, there is guarantee there would be no false positives. Daily summary reports contain only up-to-date data (no outdated information).

Ease of use

Shadowserver filters receive data in an analysis engine to classify the attacks. Reports are split into the following categories:

- Botnet URL Report (any URL that was seen in a botnet channel)
- Compromised Host Report
- Click-Fraud Report
- Command and Control Report
- DDoS Report
- Drone Report (Any IP that was seen joining a known C&C server)
- Honeypot URL Report (source URLs of where malware was downloaded from by the Honeypot systems)
- IRC Port Summary Report (summary of the ports used by C&C)
- Proxy Report (Drones are used frequently as proxies or jump points)
- Scan Report
- Sandbox URL Report (URLs that were accessed by malware)
- Sandbox Connection Report (summary of all the network traffic that the sandbox has seen for the specific interval)
- Sandbox IRC Report (list of all the new IRC C&C systems that were found after analysing malware)
- Sandbox SMTP Report (a list of email addresses that was used by malware during a sandbox run)

- Sinkhole HTTP Drone Report (IP addresses that joined the sinkhole server that did not join via a referral URL)
- Sinkhole HTTP Referrer Report (list of referral URLs that pushed systems to the sinkhole server)
- Spam-URL Report (list of the URLs and relays for Spam that was received)

Reports are available in formats:

- CSV
- HTML
- XML
- Text
- URL to download

Coverage

The Shadowserver Foundation filters data received from its worldwide sensor and monitoring networks. Data inserted into database are not limited to the specific region or network. Only organisations that directly own or control network space are allowed to receive reports; however, it is possible for national/government CERTs to obtain data relating to their constituency.

Resources required

Few resources, whether human or hardware, are required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information. Applying organisation has to be owner of the networks that it wants to monitor. National and government CERTs are allowed to apply for data related to all networks in the country.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Shadowserver runs the reports starting every morning for the previous 24 hours (UTC time-based).
Accuracy of results	Good	The service is maintained by an international team of investigators and forensics specialists. This makes Shadowserver reliable. However no one guarantees there would be no false positives.
Ease of use	Excellent	Reports are available in formats: CSV, HTML, XML, Text. Reports are split into many categories.
Coverage	Good/Excellent	Shadowserver uses worldwide sensor and monitoring networks. National/government CERTs may apply for coverage of their national networks.
Resources required	Excellent	Few resources required: a crawler to download lists or feeds on a regular basis, and a filter to parse relevant information.

5.2.16 SGNET²³ / Leurre.com HoneyNet Project²⁴

Description

SGNET / Leurre.com HoneyNet Project is a distributed network of honeypots. It is a part of the WOMBAT²⁵ project. Note that SGNET is a replacement for the original Leurre.com platform, which was based on low-interaction honeypots only; thus SGNET is less susceptible to false positives than Leurre.com was. This description is focused on SGNET itself. The project is based on protocol learning and high-interaction honeypots (Argos²⁶). Platforms are located in a number of different places all over the world. Honeypots capture all traffic to and from these platforms, store it in a database, add some contextual information (i.e. geographical information, OS fingerprinting, etc). As mentioned in WOMBAT deliverable D13 the memory tainting information generated by Argos combined with simple heuristics allows SGNET honeypots to identify shellcodes. SGNET takes advantage of part of the Nepenthes modules to understand the intended behaviour of the observed shellcodes and emulate network actions associated to it. It features a protocol learning technique based on bioinformatics. The proposed learning technique allows the creation of low-cost protocol responders, which serve as low-interaction honeypots. To become a partner, an institution needs to agree to host one of the platforms in their network. They must also sign a Non-Disclosure Agreement where they agree not to reveal the names of the partners or information about attack sources. Partners have access to the whole database, a graphical interface and a scriptable WAPI²⁷ interface, developed as part of the WOMBAT project.

Evaluation of service

Timeliness

Information in GUI is updated in real time, but data (network traffic trace) from the platforms are sent to the central server periodically on a daily basis.

Accuracy of results

The project is based on protocol learning and high-interaction honeypots. It leverages elements of the nepenthes honeypot, with protocol learning and a high-interaction honeypot – Argos. The memory tainting technique used by Argos ensures a low level of false positives.

Ease of use

SGnet has a WAPI enabled (scriptable) and a graphical user interface, accessible through a browser.

Coverage

A network of honeypots is distributed around the world. All network traffic is captured and analysed, no filter is applied. Only attacks that propagate in an active manner can be detected.

²³ http://wombat-project.eu/WP3/FP7-ICT-216026-Wombat_WP3_D13_V01-Sensor-deployment.pdf

²⁴ <http://www.leurrecom.org/>

²⁵ <http://www.wombat-project.eu>

²⁶ <http://www.few.vu.nl/argos/>

²⁷ <http://wombat-api.sourceforge.net/>

Resources required

Basic technical knowledge is required to set up honeypot platform. Technical knowledge is required to operate and interpret information. To set up a platform a server with at least Pentium II, 500 MHz, with 1 GB Hard Disk and 128 MB of memory and four public IP addresses are required.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Information in GUI is updated in real time, but data (network traffic trace) from the platforms are sent to the central server periodically on a daily basis.
Accuracy of results	Excellent	Risk of false positives is low – attacks are detected through a memory tainting technique employed in Argos.
Ease of use	Good	SGNet has a WAPI enabled interface as well as a graphical interface. The WAPI interface can be used for automation.
Coverage	Fair	Distribution is worldwide, but not on a very wide scale in comparison to other solutions. Only scanning threats are detected.
Resources required	Good	Basic technical knowledge is required to set up honeypot platform. Technical knowledge is required to operate and interpret information.

5.2.17 ARAKIS²⁸

Description

ARAKIS is an early warning system operated by NASK / CERT Polska. ARAKIS aggregates and correlates data from various sources, including honeypots, darknets, firewalls and antivirus systems in order to detect new threats. The project focuses on detection and characterisation of automated threats with a focus primarily, though not only, on exploits used in the wild, not malware. The system detects threats that propagate actively through scanning. ARAKIS uses four types of sources: a distributed network of low-interaction honeypots (main source of the data), firewalls, antivirus systems and darknets. Sensors are distributed in the Polish network. Partners get information about suspicious or malicious connections detected from their constituencies' IP addresses. Partners do not receive a sensor to plug in to their network. To gain access, CERTs can apply by email to arakis@nask.pl.

Evaluation of service

²⁸ <http://arakis.pl/en/index.html>

Timeliness

Information is distributed on a daily basis (snapshot from last 24 hours). Daily snapshots are available all the time (distribution via HTTPS file(s) download based on the certificates authentication).

Accuracy of results

The project is based mainly on low-interaction honeypots – a number of false positives is possible. However, for the data feeds a series of algorithms are applied to remove potential false positives and only forward more reliable data. Low-interaction honeypots have limited interaction with attacker, but can emulate more services and are more difficult to compromise. Additional sources are darknet and firewall. They are used primarily as anomaly detectors and signal increases and decreases of port activity. Antivirus systems are less useful, but give information about known threats on the network. Every flow (its payload) observed to the honeypots is captured and analysed using various algorithms to find correlation between them. After that information from many sources and distributed sensors is aggregated together. The network of the sensors is logically and geographically (in Poland) distributed. Filtering of data at different levels makes the data quite reliable.

All data have timestamps (data aging is present). Daily summarisation reports contain only up-to-date data (no outdated information).

Ease of use

Partners have access to download via HTTPS (authentication with certificates) files with daily snapshot of:

- Details of all connections that performed an suspicious or malicious scanning (source IP from defined subnet)
- Details of all connections that triggered any Snort IDS rule (source IP from defined subnet)

Coverage

Sensors are located only in Polish networks. Only attacks that propagate in active manner (mostly scanning threats, including bots, worms and mass exploitation tools) can be detected.

Resources required

Basic technical knowledge is required to interpret/understand information. Few resources, whether human or hardware, are required: a crawler to download lists on a regular basis, and a filter to parse relevant information.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Information is distributed on a daily basis (snapshot from last 24 hours). Daily snapshots are available all the time (distribution via HTTPS file(s) download with certificate-based on the authentication).
Accuracy of results	Good	The project is based mainly on low-interaction honeypots (a number of false positives is possible). Nevertheless additional algorithms are applied to remove false positives.
Ease of use	Excellent	One format: CSV (easy to parse). Reports are split into two categories.
Coverage	Good	Sensors are located only in Polish networks. Only active propagation can be detected.
Resources required	Excellent	Basic technical knowledge is required to interpret/understand information. Few resources, whether human or hardware, are required: a crawler to download lists on a regular basis, and a filter to parse relevant information.

5.2.18 Malc0de database²⁹

Description

Description of the service is based on information presented on the service website.

Malc0de database delivers information about URLs which serve malware, i.e. malicious executables. There is an IP address and AS number associated with every URL as well as MD5³⁰ hash of binary with hyperlink to report from ThreatExpert service.

Evaluation of service

Timeliness

URLs with malicious binaries are updated several times a day. It may happen that websites are removed by hosting providers; thus binaries are not available for download. Data aging policies are unclear.

²⁹ <http://malc0de.com/database/>

³⁰ <http://en.wikipedia.org/wiki/MD5>

Accuracy of results

Samples in malc0de database are published based on ThreatExpert classification. Since domain names/IP addresses of servers distributing malware change rather often it happens that URL provided by database could lead to nowhere – but malware hashes remain present.

Ease of use

Because service shares information on the web page access to it is very easy. Data are displayed in HTML table. It is also accessible as an RSS feed. Malc0de database can also be searched based on following criteria: MD5 hash, IP address, AS number, AS name or country code.

Coverage

Data presented on the website present URLs along with malware that was delivered through them. Information about data collecting infrastructure is undisclosed.

Resources required

In order to browse database an internet browser is required. Data can be parsed in automatic manner using an RSS feed.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Updated several times a day.
Accuracy of results	Good	Based on ThreatExpert classification.
Ease of use	Excellent	Provides an RSS feed, which allows automation.
Coverage	N/A	Collecting infrastructure is unknown.
Resources required	Excellent	Data can be viewed in regular browser or RSS reader.

5.2.19 ParetoLogic URL Clearing House / malwareblacklist.com³¹

As of 19 April 2011 The ParetoLogic URL Clearing House changed its name and moved to a new domain.

Description

Description of the service is based on information presented on the service website.

Service provides information on malicious URLs. These are collected from client honeypots and include viruses, trojans, exploits, etc. Every honeypot is sent browsing a list of sites, which is updated daily. Most sites are known to be bad; some are harvested from spam messages or malware payloads. Regular sites are crawled as well, but with lower priority.

³¹ <http://malwareblacklist.com/>

Every URL is associated with an IP address and description. Executables or regular files are stored on MalwareBlacklist's server and can be downloaded from there. There are also phishing reports available if a URL is classified as one.

List is open for everybody, but there is a free registration possible, which allows users to query the database, submit an URL and download malicious samples.

One can become ParetoLogic's partner and gain access to FTP server as well as custom APIs to query their database.

Evaluation of service

Timeliness

List is updated several times a day. It may happen that websites are removed by hosting providers; thus binaries are not available for download. Every URL includes information when it was submitted to the database.

Accuracy of results

Websites are classified by client honeypots and malicious classifications are confirmed with sandbox technology.

Ease of use

In order to search through the database or download samples free registration is required. To browse through the database an internet browser with JavaScript support is required since list of URLs is populated using asynchronous XML HTTP requests (AJAX).

Automation is possible using provided API. ParetoLogic provides the entire URL list as hourly feeds directly from their FTP server. These features are available for registered users.

Coverage

Coverage is unclear.

Resources required

To take advantage of all possibilities of the service, a registration and acceptance of terms and conditions is required.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Updated several times a day.
Accuracy of results	Good	Based on client honeypots, verified through sandboxes.
Ease of use	Good	In order to access the data, browser with JavaScript support is required. Automation using FTP server or API is possible for registered partners.
Coverage	N/A	Coverage is unclear.
Resources required	Good	Data can be viewed in a regular browser or gathered in automated manner using API.

5.2.20 SpamCop³²*Description*

Description of the service is based on information presented on the service website.

SpamCop is a service for reporting spam. It determines the origin of unwanted email and reports it to the relevant Internet service providers via email.

SpamCop also provides a Blocking List (SCBL). It lists IP addresses which have transmitted reported email to SpamCop users. One can query SpamCop database in order to check if your IP address is listed. There is also a possibility to browse for whole subnets,³³ as well as to implement local mirror³⁴ (using rsync³⁵ and SSH) of the blacklist. The mirroring service is free if access to the mirror is given to public, or for the annual fee, if the mirror is private.

Evaluation of service

Timeliness

Data in the blacklist is near real time. A reported address stays on the SCBL for only 24 hours, i.e. it will be removed from the list if there are no reports against it within 24 hours.

Accuracy of results

SpamCop Blocking List benefits from a number of report sources, including automated reports and SpamCop user submissions. The SpamCop team manually balances the threshold in an effort to make the list as accurate as possible.

³² <http://www.spamcop.net/>

³³ <http://www.spamcop.net/w3m?action=map>

³⁴ <http://www.spamcop.net/fom-serve/cache/340.html>

³⁵ <http://en.wikipedia.org/wiki/Rsync>

Ease of use

Since SpamCop uses DNSBL format its implementation is very simple. They distribute sample configuration for most popular Message Transfer Agent, like Postfix³⁶ or EXIM.³⁷ There's a possibility to create an automated solution to query the service for classification on a specific IP address.

Coverage

Service provides information regardless of its origin. It is open to user submissions from anywhere, but only catches spamming hosts.

Resources required

Basic knowledge of MTA configuration is required. It is also possible to use the list manually using tools such as nslookup. To browse the information on subnets, an Internet browser is necessary. To set up and use a local mirror, additional resources are required (hardware, bandwidth, knowledge; in case of private mirror – also annual fee). In order to receive abuse information by email one has to be a network owner.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Data in the blacklist is near real time.
Accuracy of results	Good	The SpamCop team keeps the list as accurate as possible. Low rate of false positives.
Ease of use	Good	Data are sent as regular email messages.
Coverage	Excellent	Spamcop has a very wide collection infrastructure – amongst others, manual user reports.
Resources required	Good	An email client software is required or MTA (in order to use blocking list). Additional resources needed to set up a local mirror.

5.2.21 Arbor ATLAS³⁸

Description

Description of the service is based on information presented on the service website.

ATLAS stands for Active Threat Level Analysis System and it is a globally scoped threat analysis network. Arbor collectively analyses the data traversing different 'darknets' to develop a truly globally scoped view of malicious traffic traversing the backbone networks that form the Internet's core.

³⁶ <http://www.postfix.org/>

³⁷ <http://www.exim.org/>

³⁸ <http://atlas.arbor.net/>

The ATLAS portal is a public resource that delivers a subset of the intelligence derived from the ATLAS sensor network on host/port scanning activity, zero-day exploits and worm propagation.

In addition to the global information ATLAS provides data on specific IP addresses, TCP/UDP ports, AS numbers or countries.

In order to benefit from all ATLAS's services, free registration is required.

Evaluation of service

Timeliness

Charts updated in real time and feed is updated daily. Feed contains data from last 24 hours, single entries include timestamps.

Accuracy of results

Delivered data are attributed to real threats (e.g. scans) discovered by ATLAS. Besides that, Arbor Networks has proven itself as a provider of high-quality data. Therefore information from this service can be trusted and used as one of the main sources of external data.

Ease of use

Arbor ATLAS shares data on their web page. They provide charts as well as plain data in XML or CSV formats, thus access to information is easy and straightforward. Data feeds can be built for users in the subscriber reputation feed (IODEF or CSV formats).

Coverage

Service includes globally collected data about:

- Honeypot-captured payloads
- IDS logs
- Scan logs
- Internet DoS statistics
- News & vulnerability reports
- Captured malware samples
- Phishing infrastructure data
- Botnet command & control data

Resources required

Information can be studied online on Arbor's web page. Another, more efficient, way is to download data in XML or CSV format and implement tool to gather information.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Charts updated in real time and feed is updated daily. However, sometimes information seems to be delayed.
Accuracy of results	Good	Delivered data are attributed to real threats.
Ease of use	Excellent	Service provides charts as well as plain data in XML/IODEF or CSV format.
Coverage	Excellent	One can provide a list of networks to monitor. Also, as a source of data, a darknet is used as well as honeypot and data from Arbor devices.
Resources required	Excellent	Internet browser or XML/CSV parser is required.

5.2.22 Composite Blocking List³⁹

Description

Description of the service is based on information presented on the service website.

Composite Blocking List is a DNS-based blackhole list of suspected email spam senders.

The CBL takes its source data from very large spamtraps/mail infrastructures, and only lists IP addresses exhibiting characteristics which are specific to open proxies of various sorts (HTTP, socks, AnalogX, wingate, etc.) and dedicated Spam bots which have been used to send spam, worms/viruses that do their own direct mail transmission, or some types of trojan horse or 'stealth' spamware, dictionary mail harvesters, etc.

The CBL provides lookup utility, which enables users to find out whether a particular IP address is listed in database.

As stated on CBL's website, they 'prefer users to use the SpamHaus DNSBL system to get access to the CBL, instead of the CBL directly. This has a number of benefits including more DNS servers answering queries (hence less chance of overload/delay on queries) as well as being able to query all of their DNSBLs in one query. The CBL is wholly included in (and in fact is the largest part of) the Spamhaus XBL subzone.'

Evaluation of service

Timeliness

Spam emails are collected in real time from spamtraps. Data aging policies are unclear.

³⁹ <http://cbl.abuseat.org/>

Accuracy of results

Service is widely used and has a good reputation among experts.

Ease of use

The CBL is a standard IP-based blocking list. The user has to configure their MTA to use it. Rsync access is also provided for registered users.

Coverage

Service provides information regardless of its origin. Spamming hosts are covered.

Resources required

Basic knowledge of MTA configuration is required.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Spam message are collected in real time.
Accuracy of results	Excellent	Low rate of false positives.
Ease of use	Fair/Good	To block spam basic knowledge of MTA configuration is required. For detecting problematic hosts rsync and filtering setup is required.
Coverage	Excellent	Information on spamming hosts is provided regardless of their origin.
Resources required	Good	MTA, rsync, technical knowledge to set up filtering for specific ASN or netblock.

5.2.23 Team Cymru's CSIRT Assistance Program⁴⁰

Description

Description of the service is based on information presented on the service website.

Team Cymru provides daily lists of compromised or abused devices for the ASNs and/or netblocks within a CERT's jurisdiction. This includes such information as bot-infected hosts, command and control systems, open resolvers, malware URLs, phishing URLs, and brute force attacks.

The service is free of charge and any verifiable regional or national CERT is welcome to join the program. In order to participate a short memorandum of understanding is required. This service is addressed mainly at the newly established CERTs.

⁴⁰ <http://www.team-cymru.org/Services/CAP/>

Evaluation of service

Timeliness

New data feed is provided once a day. It is also possible to access archived data from the previous three days.

Accuracy of results

Team Cymru has proven itself as a provider of reliable data over a long period of time, so it is likely that their data feed will provide good accuracy of results.

Ease of use

Common methods, such as email message or access through HTTP, are used, thus getting data feed is straightforward and can be easily automated.

Coverage

Delivered information include bot-infected hosts, command and control systems, open resolvers, malware URLs, phishing URLs, and brute force attacks. Collecting infrastructure is undisclosed, but the amount of information implies a large infrastructure.

Resources required

An email client is required in order to receive alerts and HTTP client (e.g. web browser or command line one) to download provided data.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Updated daily.
Accuracy of results	Excellent	Delivered data are attributed to real threats.
Ease of use	Excellent	Service sends reports via email.
Coverage	Excellent	Delivers information about multiple types of incidents concerning networks owned by the organisation. Large infrastructure is needed for the collection of such information.
Resources required	Good	Alerts can be read in standard MUA, browser.

5.2.24 CERT.BR Spampots⁴¹

Description

Description of the service is based on information presented on the service website.

⁴¹ <http://honeytarg.cert.br/spampots>

The Spampots project is a part of the honeyTARG HoneyNet Project. It uses an infrastructure of low-interaction honeypots to gather and analyse spam traffic data. Currently the infrastructure consists of sensors distributed in nine countries in South America, Europe, Asia and Australasia. The data are collected periodically, analysed and distributed among trusted parties. Members of the project have exclusive access to the web interface, which presents various statistical information about volumes of spam traffic.

Evaluation of service

Timeliness

The data are constantly collected by honeypots, processed and displayed on the members' portal. The portal presents data in 15-minute periods, thus allowing near real time observation of spam traffic.

Accuracy of results

Because of the closed membership accuracy of the service data feeds was not determined.

Ease of use

The service presents statistical data in easy-to-understand graphs and tables, thus allowing quick interpretation of gathered results. The members' portal is designed for manual analysis of information.

Coverage

The service sensors are distributed in nine countries around the globe, giving information on spam campaigns targeting their geographical regions. The project is looking for ways of extending its coverage by inviting organisations to join and share resources. In return the service provides information on received spam, URLs found in emails, IP addresses which abused sensors, etc.

Resources required

Organisation willing to receive data needs to become a member of the project and is obliged to set up a sensor-gathering spam traffic. Access to data requires internet browser.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Service delivers data with almost no delay.
Accuracy of results	N/A	
Ease of use	Good	Data are accessible through web browser and presented in easy-to-interpret way. Automation possibilities are unclear.
Coverage	Fair	The project has a distributed infrastructure of monitoring nodes but is still limited to a small number of countries.
Resources required	Fair	Getting access to data requires becoming a member of the project and deploying a sensor that gathers data.

5.2.25 Project HoneyPot⁴²

Description

Description of the service is based on information presented on the service website.

Project HoneyPot is a free service. It is a distributed network of decoy web pages for identifying spammers. The service allows defining a custom-tagged email addresses which if harvested by spambots will receive unwanted email messages. The project gathers and processes the data and shares findings through its website. The project was created by Inspam Technologies Inc.

Non-members have only limited access to information produced by the service – the data feeds are shortened to contain only the top 25 abusing IP addresses. Becoming a member requires registration. Active members – organisations which installed honeypot software or donated an MX record – are given access to additional information.

Evaluation of service

Timeliness

The service sends an email every 24 hours containing a summary report on IP addresses detected as spam servers. Additionally there is a possibility to define an RSS channel for various data feeds.

Accuracy of results

The service produces ratings for detected IP addresses sending spam messages. The ratings are based on the number of emails sent to the honeypot, performed dictionary attacks, harvested addresses, posted spam comments to web forms, hosted bad web pages (phishing sites, etc.), and broken no-follow or certain robot.txt rules. The service has a disclaimer stating that lists may contain false positive classifications due to the potential theft of IP address or taking control over server.

Ease of use

The service uses mainly web interface to deliver data. Data feeds are also delivered via RSS and email. Interpretation of presented information is easy and does not require extensive technical knowledge. RSS or email feeds can be used to automate process of discovery of spam servers.

Coverage

The service monitors spam traffic with almost 68 million spam traps scattered over the world. Participants that will install honeypots in their network or donate an MX record can gain access to additional data feeds. Currently the service provides a variety of statistical information and a number of data feeds concerning IP addresses separated in groups depending on type of malicious behaviour (email harvesters, comment spammers, dictionary attackers, spam servers).

⁴² <http://www.projecthoneypot.org> – note that this project is not related to the HoneyNet Project (<http://www.honeynet.org/>)

Resources required

Becoming a member requires a simple registration process. It is possible to use the public interface of the service, but it has limited data feeds. Becoming an active member requires installing a honeypot sensor or donating an MX record.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Service updates its data feeds every 24 hours. Data are sent via email with the same delay.
Accuracy of results	Good	The data feeds may contain some false positive classifications.
Ease of use	Excellent	Web interface of the service is easy to use and data feeds can be used to automate discovery of new malicious spam servers.
Coverage	Excellent	Worldwide coverage of 68 million spam traps assures great discovery capabilities.
Resources required	Good/Excellent	Basic use of service requires just registration. Organisations willing to share data are obliged to install spam trap sensor or donate an MX record.

5.2.26 Malware Threat Center⁴³

Description

Description of the service is based on information presented on the service website.

Malware Threat Center is a service delivering information on various threats on the Internet. The service is free of charge and publicly open. The service provides data without any warranty, meaning that false positive classifications are possible. It is built using data produced from the Cyber-TA Honeynet Project hosted at SRI International.

Evaluation of service

Timeliness

The service delivers data feeds every 24 hours. Data delivered by the service contain information on first and last time an incident was noticed.

Accuracy of results

The service provides data without any warranty on accuracy. Data are meant to be used mainly for research purposes.

⁴³ <http://www.mtc.sri.com>

Ease of use

Service provides datasets in two formats: as a filter list in text file and as a web page with more detailed information on the observed threat. Both formats are easy to understand and the text file filter list can be put to use as an input data for some kind of monitoring mechanism.

Coverage

The service provides five data feeds: botnet command & control servers, malware attack source and filters, most aggressively spreading malware binaries, most effective malware-related Snort signatures and malware-related DNS names. It is possible to download each feed as a text file – a watch list from the previous 10 or 30 days. There is no precise information on collecting infrastructure aside from mentioning Cyber-TA HoneyNet Project as a source of information.

Resources required

Manual use of the service requires just an Internet connection and a browser. It is possible to import data from the text files and web page as well but it will require developing a specialised solution to parse it.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Service updates its data feeds every 24 hours.
Accuracy of results	Fair	Service does not provide any warranty on published data. False classifications may be present in the data feeds.
Ease of use	Excellent	Manual use of the service is very easy. Automated data processing will require a specialised solution.
Coverage	Fair	It is not precisely defined how service gathers data apart from the fact that it is built on data received from the Cyber-TA HoneyNet Project. The service delivers five data feeds which cover various incident types.
Resources required	Good	The service is free and open. The data can be browsed with a web browser. Importing it in an automated manner will require developing specialised software.

5.2.27 Smart Network Data Services⁴⁴

Description

Description of the service is based on information presented on the service website.

Smart Network Data Services (SNDS) is a Windows Live Hotmail initiative designed to allow everyone who owns IP space to contribute to the fight against spam, malware and viruses.

⁴⁴ <https://postmaster.live.com/snds/>

In order to get access to SNDS data one has to sign up using Microsoft® Windows Live™ ID credentials.

Then, it is necessary to request access to an IP address range or ASN that one owns. SNDS sends email messages to the address associated with the given ASN (based on *whois* records).

Evaluation of service

Timeliness

Every day data from the previous day are aggregated at midnight PST (Pacific Standard Time). Published data are available for 90 days.

Accuracy of results

Accuracy of results is unclear due to lack of actual experience.

Ease of use

One way is to access data is by logging into website and browsing it.

Besides that, data are distributed as CSV file, which can be accessed via HTTPS protocol and is easy to parse in automated manner.

Coverage

Hotmail offers a large network incident-collection infrastructure which detects not only spam but also virus-infected emails, malware hosting websites and open-proxy status information. Service delivers information for the defined AS numbers or IP address ranges.

Resources required

A browser is needed to access the data. Ownership of IP address range or AS has to be proved by the applying organisation.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Good	Updated daily.
Accuracy of results	Good	Delivered data are attributed to real threats.
Ease of use	Excellent	Online browsing or downloading CSV file via HTTPS.
Coverage	Excellent	Hotmail and Microsoft have a huge network incident-collection infrastructure.
Resources required	Good	Browser, proving ownership of IP address range and/or ASN.

5.2.28 Malware Patrol⁴⁵

Description

Description of the service is based on information presented on the service website.

Malware Patrol is a free service for verifying URLs for presence of malware. The service is provided for non-commercial use and is a result of community effort. Contributors of the service are CERTs, security groups, universities, security professionals and regular users from around the world. To become a contributor an organisation can set up a spamtrap or forward suspicious emails to the service for analysis. The service delivers blacklists containing domains spreading malware. Unsanitised lists of URLs are not public but available for organisations like CERTs. A CERT must apply to get access to such lists.

Evaluation of service

Timeliness

The service rechecks URLs in its database at least once a day. New URLs are checked no later than one hour after submission. It is not clear how fast domains that were cleared of malware are removed from the block lists.

Accuracy of results

The information on processing and malware detection mechanisms is undisclosed.

Ease of use

The service web interface is straightforward, easy to navigate and understand. The block lists are of many different types and provided in ready-to-use form.

Coverage

The service is based on contributors' input. The information on number of contributors and their geographical spread remains undisclosed.

Resources required

The service provides information in ready-to-use form of block lists. There is no need for reprocessing the data.

⁴⁵ <http://www.malwarepatrol.net>

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Updated feed is provided every hour.
Accuracy of results	N/A	Methods of analysis are undisclosed.
Ease of use	Excellent	Web interface is simple and provided block lists are in ready-to-use form.
Coverage	N/A	Aside of providing feeds in many different formats, the distribution of spamtraps or contributors is undisclosed.
Resources required	Excellent	There are no special requirements for using the service.

5.2.29 Zone-H⁴⁶*Description*

Description of the service is based on information presented on the service website.

Zone-H is a service that provides information about defacements on websites. All information is collected from public sources or a result of direct notification by Zone-H's community (everyone can anonymously notify about defaced sites). Registration is optional and free for everyone.

Zone-H's data are under Creative Commons: Attribution-Noncommercial-No Derivative (CC BY-NC-ND) licence. This translates to: **for non-commercial usage**.

Evaluation of service

Timeliness

All Zone-H's lists are updated in almost real time (no more than 24 hours gaps are expected).

Accuracy of results

All submitted sites are verified (probably by experts). There is an 'Onhold' list, where submitted incidents wait to be checked ('Onhold' list is publicly available). The assumption is that false positives should be eliminated when verified during the 'Onhold list' stage. Time (date) of a website notification is present.

Ease of use

Zone-H provides three lists: one is a list with non-confirmed defacements, second is a special defacements list (for example important websites, like governmental, known companies, military, etc.), third is a 'normal' sites list.

⁴⁶ <https://www.zone-h.org/>

All lists except 'Onhold' are available via RSS feed (last 20 defacements published, updated every 5 minutes). Additionally one can receive all the special defacements per mail everyday by subscribing to the mailing list (registration is free).

One can search the service for domain names (a wildcard query is possible, for example: '.gov.pl' finds all domains *.gov.pl). Searching **requires filling a CAPTCHA** (a human is needed)!

Coverage

There is no possibility to define websites that are checked periodically. Set of websites depends only on the community that submits information to the service (no active search is performed).

Resources required

Getting lists from website or RSS feed is not challenging (requires a crawler and writing a simple parser) and can be done automatically. Searching service through built-in search engine requires a human (for filling CAPTCHA).

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	All Zone-H's lists are updated in almost real time.
Accuracy of results	Excellent	All submitted sites are verified by hand.
Ease of use	Good	One can parse information from website or RSS feeds. Searching mechanism allows 'wildcard query', but it is secured by CAPTCHA, which prevents automation of the process.
Coverage	Good	The dataset depends only on activity of the community. However, defacers tend to brag, meaning that coverage may be quite good.
Resources required	Fair– Excellent	Few resources, whether human and hardware, are required in case of parsing information from website or RSS feed. Searching is limited by CAPTCHA.

5.2.30 Cisco IronPort SenderBase Security Network⁴⁷

Description

Description of the service is based on information presented on the service website.

SenderBase is an email and web traffic monitoring network. It examines more than 90 different parameters concerning email traffic and 20 different parameters concerning web traffic. Parameters tracked include global sending volume, complaint levels, spamtrap accounts, whether a sender's DNS

⁴⁷ <http://www.senderbase.org/>

resolves properly and accepts return mail, country of origin, blacklist information, probability that URLs are appearing as part of a spam or virus attack, open proxy status, use of hijacked IP space, valid and invalid recipients, and other parameters.

Evaluation of service

Timeliness

According to the information presented on the service website data are published in real time (immediately after detection).

Accuracy of results

According to the information presented on the service website, information contributed to SenderBase comes from over 100,000 organisations (including the largest networks in the world). It now collects data on more than 25% of the world's email traffic. Cisco is a recognised corporation and has a good reputation among information security experts. All this makes results reliable. Data aging is present (reputation scores are changing in real time).

Ease of use

One can query SenderBase about the 'reputation' for:

- IP address
- network range in CIDR
- domain name
- network owners that control the servers (names)

The searching is available only via web page and is free for everyone. No additional channels are available (RSS or email).

Every item on the list has a 'reputation' score ranging from -10 (for the worst) to +10 (for the very best). The score is grouped into Good (little or no threat activity), Neutral (IP or domain is within acceptable parameters; however, email or web traffic may still be filtered or blocked) and Poor (problematic level of threat activity has been observed, email or web traffic is likely to be filtered or blocked) for simplicity reasons.

The list could be exported to a text file (plain text or Sendmail/Postfix blocklist) with different parameters (level of details).

Coverage

According to the information on the website, over 100,000 organisations from round the world submit data to the global database. It makes the coverage globally distributed.

Resources required

Querying the website is simple for a human.

Automatic querying and parsing the list requires the development of a tool (it should not be complicated).

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Database is updated in real time.
Accuracy of results	Good/Excellent	Results are reliable.
Ease of use	Excellent	Searching mechanism allows IP range in CIDR, domain name or network owner. Data can be exported to text files (in CSV formats also).
Coverage	Excellent	The coverage is globally distributed.
Resources required	Good	Few resources, whether human or hardware, are required in case of parsing information from website. However, no API is provided for automatic querying. This requires the development of a tool.

5.3 Tools/mechanisms for the proactive detection of network security incidents

5.3.1 Client honeypots

Description

Client honeypots are security tools which actively search for malicious servers. Their goal is to interact with the server and determine whether an attack has occurred. Main focus of client honeypots are usually web browsers but it is possible to create a client honeypot based on any client software, e.g. FTP clients, SSH, email, etc. The honeypots were divided in three major families: high-interaction, low-interaction and hybrid.

High-interaction honeypots are fully functional systems usually deployed as virtual machines. The system that runs on the virtual machine is equipped with a real client software and usually with some monitoring capabilities. The machine is monitored for changes when client software interacts with a server and an assessment of malicious activity is made during that process. High-interaction client honeypots are very effective in detecting unknown attacks but at the cost of performance. The performance loss comes mostly from monitoring of the whole operating system, logging of changes occurring while client software operates and determining malicious behaviour. High-interaction honeypots are prone to evasion techniques such as delays in invocation or running when some conditions are met (logic bombs). Some malicious software is also able to detect virtualisation and prevent itself from committing harmful operations. Accurate detection of attacks requires vulnerable software to be present – the right patchlevels and plugins. Despite the downsides, high-interaction client honeypots are one of the most popular and widely used systems of detecting malicious behaviour targeting home users.

Low-interaction client honeypots represent a different approach from the high-interaction ones described above. These honeypots are lightweight applications that simulate behaviour only of the client software instead of the whole operating system. The honeypot interacts with a server and responses are directly examined to discover malicious behaviour. They are relatively easy to deploy and can be extremely efficient. The downside of low-interaction honeypots is that they only simulate real software and because of this they are different from it. Usually low-interaction honeypots can detect only known malicious behaviour, which is described by signatures or heuristic algorithms. Unknown exploits and attacks are usually unnoticed. Exploits also try to evade detection by such honeypots usually taking advantage of their simplicity or incomplete support of all functionality of the real client software.

Hybrid client honeypots are the third type. The software tries to combine the advantages of both types of previously described honeypots. The approach applied to building such honeypot can depend very much on the tools used. Combining the results of low-interaction honeypot and high-interaction gives a unique insight into the exploitation and infection process. Information gathered from hybrid systems is more comprehensive and allows users to better describe the vulnerability and possibly find a method of mitigating the threat.

During the desktop research the following tools were chosen to represent the client honeypot category:

- The HoneySpider Network – developed by CERT Polska/NASK, Govcert.nl and SURFnet is a free client honeypot solution that is available for trusted CERTs. It also features an improved high-interaction Capture-HPC client honeypot component that is much more stable than the original solution.⁴⁸ To gain access to HoneySpider software, send an email to contact@honeyspider.net or contact CERT Polska directly;
- phoneyC,⁴⁹ a low-interaction client honeypot that includes a crawler and emulates vulnerabilities;
- shelia,⁵⁰ a high-interaction client honeypot for detection of not just malicious websites but various attachments as well.
- Argos⁵¹ – an emulator for zero day attacks – has also recently received support for client side honeypot use.

Note that very good online services for the analysis of malicious websites are available, such as Wepawet⁵².

The modified version of Capture-HPC will be described as an example of approach to detection of malicious behaviour directed on web browser users. This version of Capture-HPC is taken from the

⁴⁸ <https://projects.honeynet.org/capture-hpc>

⁴⁹ <http://code.google.com/p/phoneyc/>

⁵⁰ <http://www.cs.vu.nl/~herbertb/misc/shelia/>

⁵¹ <http://www.few.vu.nl/argos/>

⁵² <http://wepawet.iseclab.org/>

HoneySpider Network hybrid client honeypot. The original tool is licensed under GPLv2 license. The HSN (modified) version of Capture-HPC is very soon to be released to public.

Evaluation of Capture-HPC

Timeliness

The timeliness of the tool solely depends on configuration. One of the major parameters is visitation time describing how long a browser should stay at a scanned web page. Proposed value is 30 seconds which together with time needed to start the virtual machine and receive results sums up to about 1 minute. Therefore every minute or so a classification of one submitted URL is received.

Accuracy of results

The honeypot accuracy of detection relies on contents of its whitelists used to filter out non-malicious activity on the operating system. The lists need to be prepared by a security expert and tailored to the specific configuration of the virtual machine. A basic set of lists is shipped together with the tool but is hardly sufficient for usage. The false positive rate produced by the tool can be very high at beginning and improves with updates to the whitelists.

Ease of use

Deployment of the tool requires intermediate knowledge of host and virtual machines' (guest) operating systems. During the research phase the tool was installed on Linux Debian and guest operating system was Windows.

When configured properly the tool is fairly easy to use. The tool takes a list of URLs as a text file. All operations on virtual machine used as honeypot are performed automatically and results are available as text log files and in case of infection a zip file containing modified files from the operating system. The log files produced during the infection are easy to parse and can be used to produce report about infection process. Log file with classification contains information about visited URL and classification result.

Coverage

The tool produces information about maliciousness of the web page it visited. In case of infection the tool produces log files containing information about changes in the operating system of the guest virtual machine. The changes consider: file system modifications (creating/modifying/deleting files and folders), processes (invoked and terminated processes), Windows registry modifications (added/modified/deleted keys). Together with log file, a zip file is created archiving created (downloaded), deleted and modified files. Both log files present changes as they occurred when infection took place. It is a valuable source of information about the infection process. Binary files caught during visit on a malicious web page can be analysed offline to produce additional information regarding the type and family of malware responsible for infection.

The tool can be used to monitor selected high-priority web services on the constituency network and beyond it. The coverage is dependent on how the URLs are fed into the tool. Having it fed constantly with URLs found in spam or constituency proxy logs will greatly extend coverage of observed attacks.

Required resources

To implement a solution based on this tool a high-performance server is needed. Of course, exactly how many servers are necessary depends on the number of URLs that need to be processed in a specific time slot. Research determined that a sufficient configuration of a server would be an 8-core processor, with a minimum 8 GB of RAM and no less than 32 GB of disk space, which allows processing about 700 to 800 URLs per hour. The more disk space the more log files can be kept as history. Internet link is needed with at least 5 Mbit/s connection speed. Such configuration allowed running 12 virtual machines at once with good performance and stability. The tool requires maintenance and produced results have to be interpreted by a security expert. Therefore a delegated person is needed to operate the tool.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Software provides data usually as soon as classification process is finished, with a classification of malicious, suspicious or benign.
Accuracy of results	Fair to excellent	Accuracy of results varies and depends on the implemented solution and approach; e.g. modified Capture-HPC from HoneySpider Network project with properly configured exclusion lists can provide accuracy of results evaluated as good.
Ease of use	Fair/Good	Ease of use very much depends on solution but usually it requires technical experience and knowledge. Results produced by tool are usually provided in common format, allowing its relatively easy interpretation and parsing.
Coverage	Fair/Good	Coverage depends on implementation. For most CERTs coverage will probably be Fair to Good.
Resources required	Good	Implementing a solution usually implies significant resources but later usage is not so demanding.
Scalability	Excellent	Solutions of this kind usually scale very well allowing tool throughput to be extended.
Extensibility	Fair	Tools usually have strictly defined architecture but some modifications are possible.

5.3.2 Server honeypot

Description

Server honeypot (often referred to simply as 'honeypot') is a trap that is set to detect and track attempts at unauthorised use of network services or a whole operating system. A honeypot can be an isolated part of operating system (e.g. SSH server) or a system itself (e.g. virtualised). It is

monitored in order to act as an early warning system or to gather information regarding attacks trends.

Honeypots can be combined to create a so-called honeynet, which allows one to monitor several networks at once.

Evaluation

Timeliness

Honeypots can be configured to report in real time.

Accuracy of results

Due to their nature honeypots (depending on whether they are low-interaction or high-interaction) may be prone to a number of false positives. They can provide valuable information about network, but only as an additional source.

Ease of use

Proper installation and configuration require some effort. After this process, monitoring is simple and straightforward.

Coverage

Honeypots can be placed in many locations in network, both physical and logical, to cover a whole constituency. Honeypot can cover different types of services, depending on honeypot configuration or type chosen. There is no guarantee that a malicious host on a network will connect to it. Honeypots usually detect threats that use scanning to propagate: bots, scanning worms or mass rooting tools.

Coverage is better when honeypots are placed in separate logical and physical locations.

Resources required

A possibility to assign (public) IP addresses or at least forward single UDP/TCP ports to a honeypot may be required. Additional hardware may be also required.

Examples

Most popular server honeypots include:

- honeyd⁵³
- nepenthes⁵⁴
- dionaea⁵⁵
- kippo⁵⁶

⁵³ <http://www.honeyd.org/>

⁵⁴ <http://nepenthes.carnivore.it/>

⁵⁵ <http://dionaea.carnivore.it/>

⁵⁶ <http://code.google.com/p/kippo/>

- kojoney⁵⁷
- VoIP Honey⁵⁸
- amun⁵⁹
- argos⁶⁰

*dionaea*⁶¹

Dionaea is a nepenthes successor. It can emulate a number of network services using different network protocols. Dionaea can also detect shellcodes using libemu and it uses python as a scripting language. It supports IPv6 and TLS.

Dionaea's intention is to trap malware exploiting vulnerabilities exposed by services offered to a network; the ultimate goal is to gain a copy of the malware.

Supported network protocols include: HTTP, FTP, TFTP, SMB, Microsoft SQL and SIP (VoIP).

Installation and configuration might look complicated at the first sight, but it is documented at the project website.

Project is actively developed, so there is a great chance that other services will be supported and more functions implemented in the future. It is also possible to write your own modules.

⁵⁷ <http://kojoney.sourceforge.net/>

⁵⁸ <http://voiphoney.sourceforge.net/>

⁵⁹ <http://amunhoney.sourceforge.net/>

⁶⁰ <http://www.few.vu.nl/argos/>

⁶¹ <http://dionaea.carnivore.it/>

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Data are collected in real time.
Accuracy of results	Good	False positives depending on whether honeypots are low- or high-interaction. In general high-interaction honeypots (or ones that emulate services in a more advanced manner) will generate few false positives. Low-interaction client honeypots may generate more, but these can usually be filtered out easily over time.
Ease of use	Good	Proper installation and configuration require some effort.
Coverage	Good	Connections to honeypot may be established from the whole network. Coverage will depend on exact implementation, such as placing of nodes, amount of IP addresses assigned, as well as supported services. Honeypots usually detect threats that use scanning to propagate: bots, scanning worms or mass rooting tools.
Resources required	Good	A possibility to assign (public) IP addresses or at least forward single UDP/TCP ports to honeypot. Additional hardware may be required.
Scalability	Good	One can create honeypot network (i.e. honeynet) to monitor several networks at once. More hardware may be required.
Extensibility	Good	Honeypots are usually designed to be able to use modules in order to emulate different network services.

5.3.3 Sandboxes

Description

A sandbox is an environment in which a suspicious code or application can be run in isolation without affecting a real operating system (environment is separated from OS). All behaviour of analysed software is logged including network connections. From the information gained in the analysis process, specialists can deduce if the software is malicious or not. In case of active incident detection, all network traffic captured in sandbox is interesting. If software is malicious it is likely that it will download other binary executables, or connect to servers on the Internet. These servers could be compromised or dedicated hosts that have malware or are C&Cs of the botnet. As a result of such analysis IP addresses and URLs are obtained. If they belong to a CERT constituency's network an incident should be created. Otherwise sharing of such information with relevant CERTs is recommended.

Sometimes running a web browser with a suspicious URL in an isolated and controlled environment can be considered a sandbox mechanism, but it might also be categorised as high-interaction client honeypot. The distinction is usually based on the fact that the purpose of a sandbox is more in-depth analysis of (usually) malware, while the goal of a client honeypot is to determine whether something is malicious in the first place, and thus of interest to analyse further.

Evaluation

Timeliness

Information is available after analysis, which means almost in real time.

Accuracy of results

Sandboxes may produce false positives and false negatives. False positives are especially possible if an analysed binary is not malicious in the first place. Furthermore, sandboxes are normally not meant to be used as a sole source of information about incidents but rather an incident would be generated as a side effect of another case (i.e. deep analysis of the malware or botnet).

Ease of use

Difficulty of install and setup process depends on the provider of the sandbox mechanism. Advanced technical knowledge could be necessary to understand an analysis result. Probably a full-time specialist is required.

Coverage

Analyzing applications or piece of code does not guarantee finding malicious hosts in CERT's constituency's network. Rather an incident would be generated as a side effect of another case (i.e. deep analysis of the malware or botnet).

Required resources

Requires technical knowledge (to install, configure, operate, analyse the results). A high performance computer or server is needed (depending on the scale – how many samples are analysed simultaneously). For large scale observations many machines are needed, along with maintenance. However, not many resources are needed for ad-hoc analysis from time to time. These can be also provided in a virtual environment.

Example (open-source) tools that can be used:

- Cuckoo⁶²
- Minibis from CERT.at⁶³
- Zero Wine Malware Behavior Analysis⁶⁴
- Buster Sandbox Analyzer⁶⁵

⁶² <http://www.cuckoobox.org/>

⁶³ http://cert.at/downloads/software/minibis_en.html

⁶⁴ <http://zerowine.sourceforge.net/>

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Information is available after analysis, which means almost-real time.
Accuracy of results	Fair/Good	Sandboxes could generate false positive data and additional analysis/verification is necessary. Sandboxes are not meant to be used as a sole source of information about incidents but rather an incident would be generated as a side effect of other case.
Ease of use	Fair	Advanced technical knowledge could be necessary to understand and analyse the result. Probably a full-time specialist is required. Additional analysis/verification for exclusion of false positives could be necessary.
Coverage	N/A	
Resources required	Fair	Requires technical knowledge (to install, configure, operate, analyse the results), high performance computer or server is needed.
Scalability	Fair-Excellent	Depends on vendor.
Extensibility	Fair-Excellent	Depends on vendor.

5.3.4 Firewall

Description

A firewall is a device or application designed to filter (permit or deny) network connections. Nowadays firewalls are sophisticated and can analyse protocols over layer 4 in the OSI model as well as provide stateful packet inspection. These tools can be used in proactive incident detection and create an incident when suspicious (denied) traffic from a CERT's constituency's IP addresses (both incoming and outgoing) is detected. It is recommend to filter and analyse also outgoing traffic, not only incoming. For example, bulk connection attempts from IP address in constituency network to typical Microsoft Windows service TCP or UDP ports (like 138, 139 or 445) are suspicious and could be a result of worm infection. Bulk connections to one or limited number of IP addresses could be a result of DDoS being prepared/carried out by infected hosts in constituency network (bots connected to some botnet). In similar way it is possible to detect hosts sending spam (bulk connections to the SMTP servers – typically port 25/TCP). Detection of connections to the known C&C server IP addresses (if such list is provided) from constituency network is another way to proactively detect incidents.

⁶⁵ <http://bsa.isoftware.nl/>

There are two main ways to use firewalls in proactive incident reporting – depending on firewall features. One is to create an alert directly in a firewall. Second is to use additional tools (i.e. some script) to parse/analyse firewall's logs.

Evaluation

Timeliness

Typical firewall alerts are real time or almost-real time. If an external log parser or alert generator are used the delay will depend on how often and for how long these tools run. It is possible (and recommended) to make a firewall log in real time using 'syslog' straight to the server, where data are then analysed.

Accuracy of results

A lot of false positives and false negatives are possible. Some additional algorithms/analyses are recommended to lower the amount of both.

Ease of use

If a firewall has alarm generation feature, usage is not challenging. If a log parser and/or analyser is required such method requires more advanced technical knowledge, but usually logs are easy to parse.

Coverage

Potentially network coverage can be very large, as firewalls quite commonly are used for all connections between internal networks and Internet (and also to separate local network segments). However the actual coverage will depend a lot on specific implementation – from which parts of the network logfiles or alerts can be collected. Also it is worth to remember that using firewalls it is possible to detect incidents using patterns mostly on Level 3 and 4 of the OSI model – higher level protocol patterns are out of reach of Level 3 and 4 firewalls.

Required resources

A specialist with technical knowledge is required to interpret information, filter false positives and configure a firewall. Firewalls are usually present in any network infrastructure. If not, then the required resources (dedicated firewall) will be relatively high. Parsing of the logs may also require additional hardware, depending on the amount of traffic and what is actually being logged (for instance, are we looking for all connections or just blocked ones).

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Typical firewall alert in real time or almost-real time. The rest depends on implementation, depending on how often analysis is actually run.
Accuracy of results	Fair	A lot of false positives and false negatives are possible.
Ease of use	Good	Usually logs are easy to parse.
Coverage	Fair/Good	From good: if constituency's network is fully controlled by firewalls and their logfiles are used for log analysis and alerts. (Excellent rating is not applied as higher-level protocols are usually poorly covered.) To fair: if firewalls are not used comprehensively or their logfiles/alerts are not accessible to the CERT team (which is true for many national CERTs).
Resources required	Good	A specialist is required to interpret information, filter false positives and configure firewall. For most cases, the hardware resources required will not be much, but for larger networks this may become a significant problem.
Scalability	Excellent	Usually firewalls are focused on analysis of large amount of traffic.
Extensibility	Fair–Excellent	Depends on vendor.

5.3.5 IDS/IPS

Description

An Intrusion Detection System is a software component (often integrated with device, especially in the case of commercial solutions) that monitors and analyses network traffic or operating system behaviour for unauthorised or malicious activities. An IDS system typically works in passive mode: detects threat, logs information and triggers an alert. An Intrusion Prevention System is similar to IDS, but typically works in active mode: it is able to block malicious behaviour.

Detection mechanisms in both IDS and IPS systems are classified into two main categories: signature-based (activity is compared with predefined attack patterns), or anomaly-based (activity other than learned by the system earlier is treated as suspicious). These systems typically can analyse full network traffic with payload.

IDS or IPS can be used in proactive incident detection – an incident will be created when IDS/IPS alerts are generated with a CERT constituency's IP address involved.

Evaluation

IDS/IPS systems are natural sources for incident reporting. All reported IP addresses are suspicious by default and related with attack type (or even name in case of signature-based systems).

Timeliness

Typical IDS/IPS alert is in real time or almost real time.

Accuracy of results

A lot depends on configuration and tuning. Some false positives and false negatives are possible; however, if signatures are good, output data are reliable.

Ease of use

Installation, setup, configuration and fine-tuning process of IDS or IPS require advanced technical knowledge and considerable time. Daily usage is less challenging.

Coverage

If IDS or IPS is used, typically all traffic in a LAN is analysed, so this method could potentially cover large network space. Sometimes it is limited due to a licence or number of sensors. Coverage could be excellent if constituency's network is fully monitored by IDS/IPS, or fair if constituency's network is only partially monitored by particular IDS/IPS (i.e. national CERTs).

Required resources

Commercial IDS and IPS systems usually are supplied in hardware. Free/non-commercial systems are delivered as software – hardware must be obtained separately. Significant technical knowledge and experience is required.

Examples (open source, free):

- Snort IDS⁶⁶
- Suricata IDS⁶⁷
- Bro IDS⁶⁸

Note also the rise of specific bot-hunting tools, commonly focusing on passive IRC bot detection. While not strictly IDS, they operate on similar principles. Some examples are:

- Bothhunter⁶⁹
- Rishi⁷⁰

⁶⁶ <http://www.snort.org>

⁶⁷ <http://www.openinfosecfoundation.org/index.php/download-suricata>

⁶⁸ <http://bro-ids.org/>

⁶⁹ <http://www.bothunter.net/>

⁷⁰ <http://sourceforge.net/projects/rishi/>

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Typical IDS/IPS alert in real time or almost-real time.
Accuracy of results	Good	Some false positives and false negatives are possible; however if signatures and configuration are good, output data are reliable.
Ease of use	Good	Installation and setup/configure process of IDS or IPS requires advanced technical knowledge and considerable time. Daily usage is less challenging.
Coverage	Fair–Excellent	From fair: if constituency’s network is not fully monitored by particular IDS/IPS (i.e. national CERTs). To excellent: if constituency’s network is fully monitored by IDS/IPS.
Resources required	Fair/Good	Commercial IDS and IPS systems are usually supplied in hardware. Free/non-commercial systems are delivered as software – hardware must be delivered separately. Significant technical knowledge and experience is required.
Scalability	Good	Usually tool is focused on analysis of large amount of traffic, but requires more performance (hardware) than firewall. Depends on hardware.
Extensibility	Fair-Excellent	Depends on vendor.

5.3.6 NetFlow

Description

NetFlow is a network protocol designed for collecting, monitoring and analysing IP-based traffic. While NetFlow is not exactly a tool, but a mechanism/protocol for collection, monitoring and analysing IP-based traffic, it is included here in the tools chapter because the CERT community uses it and related tools extensively to gather and analyse traffic flow data in order to identify potential incidents.

This protocol is used by active L3 (layer 3) network devices such routers and managed switches. NetFlow data (named NetFlow records) contains information about layer 3 (IPv4 or IPv6) headers and layer 4 (only if TCP or UDP protocol was used) headers. No information about payload and higher level (5 and over) protocols is included. NetFlow requires devices that support this protocol and software that can operate on collected data (i.e. analyse and generate statistics). One of the most popular free software for collecting NetFlow data – often used by CERTs, as shown in the survey – is nfdump (with CLI – Command Line Interface) and NfSen, which is a graphical (WebGUI) frontend for

the nfdump⁷¹. Other commonly used tools include Team Cymru's Flow Sonar,⁷² flow-tools,⁷³ argus⁷⁴ and SiLK.⁷⁵

The protocol was developed by Cisco Systems, but is supported by other vendors of active network devices. Many vendors created similar solutions but used different names, i.e. jFlow (Juniper), sFlow (HP and others), NetStream, nProbe.

Netflow could be used in anomaly and abnormal traffic detection, and is especially helpful in scanning and DDoS attack detection and mitigation process. This makes this technology useful in proactive incident detection. Main target is detection of compromised (or problematic in any other way) hosts in LAN.

All these things are also possible in case of other similar protocols (jFlow, cFlow, IPFIX, etc.)

Evaluation

Timeliness

NetFlow provides data in real time.

Accuracy of results

Accuracy of results can be good but dependent on quality of scripting to detect suspect activity. It may require additional filtering. Because NetFlow data can be computationally expensive and can overload network devices, sometimes not every packet is analysed. It is sampled instead. This can potentially cause some false negatives. A solution in this case may also be to use a dedicated flow generator off a tap instead. Examples include: fprobe,⁷⁶ YAF⁷⁷ or argus.

Ease of use

Using NetFlow is quite complicated and requires technical and practical knowledge in network administration. A dedicated person to administer and operate NetFlow analyser is probably required.

Coverage

All traffic passing through network devices supporting NetFlow can be monitored. Both IPv4 and IPv6 (additionally TCP and UDP) traffic is supported. The coverage can be very good if the entire constituency's network is covered by NetFlow, but could be fair if constituency's network is not fully covered by NetFlow data (i.e. national CERTs). As in case of many firewalls, NetFlow can be used to detect incidents using patterns only on Level 3 and 4 of the OSI model.

⁷¹ <http://sourceforge.net/projects/nfdump/> and <http://sourceforge.net/projects/nfsen/>

⁷² <http://www.team-cymru.org/Services/FlowSonar/>

⁷³ <http://www.splintered.net/sw/flow-tools/>

⁷⁴ <http://www.gosient.com/argus/>

⁷⁵ <http://tools.netsa.cert.org/silk/index.html>

⁷⁶ <http://fprobe.sourceforge.net/>

⁷⁷ <http://tools.netsa.cert.org/yaf/index.html>

Required resources

Network devices that support NetFlow or similar protocols are required. Enabling NetFlow on network device may require upgrading its hardware due to additional load imposed by processing of NetFlow data. Additional server to collect data and run data processing software is also required. Disk space is a major issue. Advanced technical knowledge and skills are required.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	NetFlow provides data in real time.
Accuracy of results	Good	Depends on the quality of interpretation. Different scripts to detect suspicious activity can be found in some tools.
Ease of use	Fair	Using NetFlow is quite complicated and requires technical and practical knowledge in network administration. A skilled person, who administers and operates NetFlow analyser is required.
Coverage	Fair/Good	From fair: if constituency's network is not fully covered by NetFlow data (i.e. national CERTs). To good: if constituency's network is covered by NetFlow.
Resources required	Fair	Network devices that support NetFlow or similar protocols are required. Additional investment in enhancing routers capabilities may be necessary. Additional server to collect data and run data processing software is also required. Advanced technical knowledge and skills are required.
Scalability	Good/Excellent	Every new network device sending flows is easy to add. Only hardware performance could be an issue.
Extensibility	Good	Depends on vendor. Nfdump/nfsen allow for easy extension through plugins.

5.3.7 Darknet

Description

A darknet (aka network telescope) is used to monitor traffic targeting unused IP addresses with no interaction with the observed traffic whatsoever (this is the main difference between a darknet and honeypot). All traffic to the unused IP addresses is suspicious by default. Normally this traffic would be dropped in the edge routers. Information obtained through a darknet describes large automated scanning, worm activities, DDoS attack backscatter and misconfigured network devices. Due to its nature, usage is limited only to the large ISPs, academic networks or large enterprises (large blocks/amount of unused IP addresses is required).

Because all incoming traffic is suspicious, a darknet is a great tool for proactive detection of incidents. All connections from CERT's constituency's IP addresses are malicious or suspicious. False positives in this case are usually the result of misconfiguration of the device generating the traffic.

Evaluation

Timeliness

Data are collected in real time. Processing time may be dependent on the size of the monitored IP addresses. Analyses can be fully automated and should give information such as: source IP addresses, ports used in conversation, number of connection attempts. Additional information depends on the protocol used (for example flags in TCP).

Accuracy of results

False positives are possible but relatively easy to filter out.

Ease of use

There is no publicly available dedicated 'darknet' solution. Therefore deployment and setup of the darknet may be difficult and requires skilled technical staff. In many cases using a darknet will not be possible due to the requirement of having a large amount of unused IP addresses. However, when deployed, the daily usage of darknet is not challenging. Analyses should be fully automated, with additional scripts filtering out obvious false positives.

Coverage

In this solution only traffic targeting unused IP addresses is detected. This means scanning, scanning worms and bots and DDoS backscatter. However no limitation to the source origin is applied – connections to a darknet can be established from the whole Internet.

Required resources

Potentially a separate server with adapted software is required. Qualified technical staff may be necessary. In many cases using a darknet may not be possible due to requirement of large amount of unused IP addresses.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Information is available after analysis, which may mean (almost) real time for smaller installations.
Accuracy of results	Good	False positives are possible but relatively easy to filter out.
Ease of use	Fair	Because there is no known dedicated solution, deployment and setup of the darknet idea may be difficult and require skilled technical stuff. However after deployment, daily usage of darknet is not challenging.
Coverage	Fair–Excellent	Scanning activity will be detected. Much is dependent on the number of IP addresses assigned to the darknet. ISPs that have hundreds of thousands of IP addresses still free, can assign them to a darknet, ensuring excellent coverage.
Resources required	Fair	Potentially a separate server with adapted software is required. Qualified technical staff could be necessary. In many cases using a darknet would not be possible due to the requirement for a large number of unused IP addresses.
Scalability	Good	Usually darknet is focused on analysis of large amount of traffic. Only slow hardware or network bandwidth could be a problem.
Extensibility	Fair	Probably writing own piece of code (analysers) is required.

5.3.8 Passive DNS monitoring

Description

It is possible to analyse DNS traffic to detect malicious activity. Deep investigation of the DNS lookups might be useful to track down a botnet and or even track the origin of malware and or a botnet. One of the spectacular cases when DNS query logs were used in post-mortem analysis and investigation of the attack was the Google ‘Aurora’ case.^{78,79}

In case of proactive incident detection analysis of DNS logs could be used to detect infected IP addresses. These hosts have to use DNS servers under CERT’s control or with access to logs.

⁷⁸ Damballa's white paper, ‘Report: The Command Structure of the Aurora Botnet: History, Patterns, and Findings’, 2 March 2010, <http://www.damballa.com/research/aurora/>

⁷⁹ Interview with Heather Adkins, information security manager in Google Inc., ‘For Google, DNS log analysis essential in Aurora attack investigation’, 15 June 2010, <http://searchsecurity.techtarget.com/news/1514965/For-Google-DNS-log-analysis-essential-in-Aurora-attack-investigation>

Blacklists can be used to match DNS queries about known malicious domains. In this case false negatives are expected. There are techniques to evaluate the maliciousness of the domain (the name does not necessarily have to be on the predefined blacklist)^{80,81} – in that case fewer false negatives are expected. For example an algorithm can detect unexpected spikes of domain queries (in some circumstances this could be suspicious), another can detect hard-to-pronounce (randomly generated) domain names or typo squatter ones. Similarly, searching of not-recommended (in RFC) characters in DNS names could lead to detection of fraud or phishing pages, etc.⁸² Other abuse could be detected in anomalous DNS records (for example: low TTL or of bulk associated A records). Other mechanisms can detect fast-flux domains, or analyse reverse-DNS lookups (compare to original domain name).

Evaluation

Timeliness

DNS logs can be analysed in real time.

Accuracy of results

Accuracy of results depends on quality of algorithms and techniques using for analysing DNS logs. Google's Aurora case is an example of how this method could be excellent in proactive incident detection. If blacklists are used, accuracy is dependent on their quality.

Ease of use

Technical skills and knowledge are necessary to set up, configure and adjust the tool. Daily usage is not so challenging, but sometimes could require technical knowledge.

Coverage

Coverage may be excellent if a CERT has access to all constituency's DNS logs and has excellent blacklists. However, this is a rather unlikely scenario. Only attacks that use DNS entries in some way may be detected. Hence coverage is estimated to be fair to good.

Required resources

Passive DNS log analysis requires technical knowledge (to set up, configure and analyse the results). For large number of DNS logs a high-performance server is required.

Examples

Services providing passive DNS data:

- ISC pDNS⁸³

⁸⁰ Leyla Bilge, Engin Kirda, Christopher Kruegel, Marco Balduzzi, 'EXPOSURE: Finding Malicious Domains Using Passive DNS Analysis', <http://www.iseclab.org/papers/bilge-ndss11.pdf>

⁸¹ Bojan Zdrnja, Nevil Brownlee, and Duane Wessels, 'Passive Monitoring of DNS Anomalies', http://www.caida.org/publications/papers/2007/dns_anomalies/dns_anomalies.pdf

⁸² Security Monitoring of DNS traffic, Bojan Zdrnja, 2006: http://www.caida.org/~nevil/Bojan_Zdrnja_CompSci780_Project.pdf

⁸³ <https://dnsdb.isc.org/>, https://sie.isc.org/wiki/Passive_DNS

This service cannot be used directly as a source of incident since it does not provide information regarding maliciousness of a domain. Despite this pDNS can help one to deliver more data on a particular case. DNSDB is currently in a closed beta.

- BFK Passive DNS replication⁸⁴

The service is free of charge and available publicly on the Internet. Information contained in its database is very useful for conducting investigations and allows for correlation and enrichment for already owned datasets in an automated manner.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	DNS logs could be analysed in real time.
Accuracy of results	Good/Excellent	Excellent results require good algorithms for analysing logs. For blacklist based detection, dependent on quality of entries.
Ease of use	Good	Technical skills and knowledge is required to set up, configure and adjust the tool. Daily usage is not so challenging.
Coverage	Fair/Good	Coverage is limited by access to constituency logs, blacklist entries and quality of detection algorithms.
Resources required	Good	Technical knowledge is required. For large numbers of DNS logs a high-performance server is required.
Scalability	Good/Excellent	Only hardware performance could be an issue.
Extensibility	Fair	Probably writing own piece of code (analysers) is required. Scripts to match blacklists may be needed.

5.3.9 Antivirus programs

Description

An antivirus program is used to prevent, detect, and remove malware from computers. In this evaluation only server-side antivirus software is taken in consideration.

These are often installed on the email gateway (SMTP server) in order to scan both incoming and outgoing email messages.

Another possibility is to install it on the web proxy and look for malware in HTTP traffic (works both for downloading and for uploading content).

⁸⁴ http://www.bfk.de/bfk_dnslogger_en.html

Most antivirus software is signature-based; thus they detect only known malware. However, nowadays all major vendors of antivirus programs use additional heuristics for detection purposes.

Evaluation

Timeliness

Antivirus detects suspicious content and inform administrator in real time.

Accuracy of results

False positive rate is low. Some vendor mistakes may occur in signatures.

Ease of use

Installation of this kind of software is straightforward. Configuring antivirus to work with other services (such as filters on MTA) requires additional effort. Signatures need to be kept up to date (this process can and should be automated).

Coverage

Antivirus program can, theoretically, scan all incoming and outgoing traffic.

Resources required

SMTP server or web proxy. May require additional disk space or computing power.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Detects malware in real time.
Accuracy of results	Good	Low rate of false positives.
Ease of use	Good	Only installation and initial configuration might require some effort.
Coverage	Fair–Excellent	Depends on where it is placed – how big share of constituency traffic is analysed by antivirus.
Resources required	Good	May require additional disk space or computing power.
Scalability	Good	One antivirus program can be used in the whole network. Hardware performance may be an issue.
Extensibility	N/A	N/A

5.3.10 Spamtrap

Description

In short, a spamtrap is a tool (i.e. honeypot) used to collect spam. Basically it is an email mailbox, which is advertised e.g. on web pages for the sole purpose of being harvested by spammers who then add them to their database and send unwanted messages.

Another possibility to deploy a spamtrap is to intercept all messages. A spamtrap can also be 'fed' by sending or forwarding to them spam messages received by regular users.

Evaluation

Timeliness

Given how much spam is being sent, one can conclude that spamtrap provides information in real time. Of course reporting about messages being received depends on implementation.

Accuracy of results

There is a chance that a legal (i.e. not spam) message will be received by spamtrap. Due to that a spamtrap can occasionally generate false positives. To minimise false positives rate it is a good idea to deploy dedicated spamtrap.

Ease of use

There are lots of tools and libraries to access email data; thus it is easy to use information gathered by a spamtrap. Nevertheless the process of parsing spam messages requires technical knowledge and some programming skills. Deploying a complete spamtrap solution from scratch might be complicated.

Coverage

Coverage depends on how well propagated the spamtrap email addresses are. One should make an effort to 'advertise' spam domains to catch more spam messages. Nevertheless, even if a spammer is in the constituency network there remains a high probability he will not fall into the spamtrap (which will probably remain a small part of the world email address space).

Resources required

In order to use a spamtrap one has to use the existing MTA or configure one dedicated exclusively for this purpose. There are filters (i.e. mail filters) available which allow the MTA to specify email addresses or full domains to be defined as honeypot ones. Analysing spam messages and/or logs from spam trap requires technical knowledge.

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Data are provided in real time.
Accuracy of results	Fair/Good	Depends on specific implementation. There is a chance that legitimate email will be classified as spam and received by a spamtrap. It also depends on quality of filters and parsers.
Ease of use	Fair	Initial setup requires creating filters as well as parsers.
Coverage	Fair	Depends on spamtrap domain 'popularity'.
Resources required	Good	Mail Transfer Agent and a message filter/log analyser are required.
Scalability	Good	Modern MTA can handle very large amount of data. The limitation is the number of available domains, which can be deployed as spamtrap.
Extensibility	Good	Spamtrap can be extended, but it requires knowledge regarding MTAs. One can easily add different mail filters for spam classification.

5.3.11 Web Application Firewall

Description

A web application firewall (WAF) is an appliance, server plugin, or filter that applies a set of rules to an HTTP conversation. Generally, these rules cover common attacks such as Cross-site Scripting (XSS) and SQL Injection. By customising the rules to your application, many attacks can be identified and blocked. The effort required to perform this customisation can be significant and needs to be maintained as the application is modified.⁸⁵

It is also possible to implement WAF in web application itself.

Evaluation

Timeliness

Web Application Firewall provides information in real time.

Accuracy of results

With well-suited configuration false positives rate is very low.

Ease of use

Good setup and adaptation to specific application is needed, which can be complicated.

⁸⁵ https://www.owasp.org/index.php/Web_Application_Firewall

Coverage

It covers just WWW services, often specific applications.

Resources required

Good knowledge of how a particular web application works in order to configure WAF. Extensive tests may be required to eliminate false positives.

Examples

- ModSecurity⁸⁶
- AQTRONIX WebKnight⁸⁷
- TrustWave WebDefend⁸⁸
- FortiNet FortiWeb⁸⁹
- Cisco ACE⁹⁰
- Imperva SecureSphere⁹¹
- f5 BIG-IP⁹²

ModSecurity⁹³

ModSecurity is an open-source and free web application firewall Apache module. It provides protection from a range of attacks against web applications and allows for HTTP traffic monitoring and real time analysis with little or no change to existing infrastructure.

Developers provide an extensive documentation on project website. Since ModSecurity is probably the most popular free WAF, there are many configuration examples or how-tos available online.

⁸⁶ <http://www.modsecurity.org/>

⁸⁷ <http://www.aqtronix.com/?PageID=99>

⁸⁸ <https://www.trustwave.com/web-application-firewall/>

⁸⁹ <http://www.fortinet.com/products/fortiweb/>

⁹⁰ <http://www.cisco.com/en/US/products/ps6906/index.html>

⁹¹ <http://www.imperva.com/products/products.html>

⁹² <http://www.f5.com/products/big-ip/>

⁹³ <http://www.modsecurity.org/>

Evaluation table

Criterion	Evaluation	Reason
Timeliness	Excellent	Real time.
Accuracy of results	Good/Excellent	Depends on configuration, e.g. built-in WAF can render excellent accuracy of results, while one with wrong configuration produces lots of false positive alarms.
Ease of use	Fair	Might be complicated to achieve good results.
Coverage	Fair	Covers only specific web application.
Resources required	Fair	Good knowledge of how a particular web application works in order to configure WAF. Extensive tests may be required to eliminate false positives.
Scalability	Good	It is possible to have one WAF that monitors many web applications. Tweaking the configuration may be a continuous process as new applications are added.
Extensibility	Good	One can add new filters to detect different threats.

5.3.12 Application logs

Description

The category 'Application logs' includes analysis of system logs, database logs and proxy logs. Log monitoring and further analysis can be used as a part of early warning system. Various queries that generate errors or incorrect login attempts can be gathered from a wide variety of logs. Furthermore, for example, proxy logs could be used together with blacklists to find if known malicious hosts/domains/URLs have been contacted by a constituency.

Evaluation

Timeliness

Timeliness depends on configuration of a particular application. It is possible to gather data in real time.

Accuracy of results

With proper configuration number of false positives should be on a fairly low level.

Correlating logs from different systems reduces the number of false positives as well as providing a broader look at what is happening in the monitored system.

Ease of use

Access to logs may depend on the particular application. Interpretation may be difficult. Formats will also be different and will require dedicated regular expressions for parsing.

Coverage

Different logs cover different parts of constituency services and applications

Resources required

Knowledge of logging capabilities of specific application is required.

Evaluation table

Since this category includes many different logs from different tools, it is impossible to evaluate it as one.

5.4 Summary of the evaluation of services and tools/mechanisms⁹⁴

This section presents summarised results of the desktop research. Services and categories of tools recognised as useful for gathering information for incident detection and handling were evaluated in a similar way as in the survey. The table below presents the evaluation of services and categories of tools. The evaluation was made by the researchers in regard to a common feature list for both types of incident data sources.

Service	Timeliness	Accuracy of results	Ease of use	Coverage	Resources required
DNS-BH Malware Domain Blocklist	Fair	Good	Excellent	Excellent	Excellent
MalwareURL	Good	Good	Excellent	Excellent	Excellent
DSHIELD	Excellent	Fair	Good	Excellent	Excellent
Google Safe Browsing Alerts	Good	Fair	Good	Excellent	Good
HoneySpider Network (as a service)	Excellent	Fair	Good	Fair	Excellent
AusCERT	Good	Good	Good	Good	Excellent
Cert.br data feed	Good	Good	Fair	Good	Good
FIRE	Good	Good	Fair	Good	Good
Team Cymru - TC Console	Excellent	Good	Good	Excellent	Excellent
EXPOSURE	Good	Good	Excellent	Good	Excellent
AmaDa	Excellent	Good	Excellent	Fair	Excellent
Malware Domain List	Excellent	Good	Excellent	Good	Excellent
Zeus/SpyEye Tracker	Good	Excellent	Excellent	Fair/Good	Excellent
The Spamhaus Project Datafeed	Excellent	Good	Good	Excellent	Good
Shadowserver Foundation	Good	Good	Excellent	Good/Excellent	Excellent
SGNET	Good	Excellent	Good	Fair	Good
ARAKIS	Good	Good	Excellent	Good	Excellent
Malc0de database	Excellent	Good	Excellent	N/A	Excellent

⁹⁴ Evaluation valid as of September 2011. Note that services and tools can change over time.

Service	Timeliness	Accuracy of results	Ease of use	Coverage	Resources required
ParetoLogic URL Clearing House	Excellent	Good	Good	N/A	Good
SpamCop	Excellent	Good	Good	Excellent	Good
Arbor ATLAS	Good	Good	Excellent	Excellent	Excellent
CBL (Composite Blocking List)	Excellent	Excellent	Fair/Good	Excellent	Good
Cert.br Spampots	Excellent	N/A	Good	Fair	Fair
Team Cymru's CAP	Good	Excellent	Excellent	Excellent	Good
Project Honeypot	Good	Good	Excellent	Excellent	Good/Excellent
Malware Threat Center	Good	Fair	Excellent	Fair	Good
Smart Network Data Services	Good	Good	Excellent	Excellent	Good
Malware Patrol	Excellent	N/A	Excellent	N/A	Excellent
Zone-H	Excellent	Excellent	Good	Good	Fair–Excellent
Cisco IronPort SenderBase	Excellent	Good/Excellent	Excellent	Excellent	Good

Table 1: Evaluation of services – summary

Category	Timeliness	Accuracy of results	Ease of use	Coverage	Resources required	Scalability	Extensibility
Client honeypot	Excellent	Fair–Excellent	Fair/Good	Fair/Good	Good	Excellent	Fair
Server honeypot	Excellent	Good	Good	Good	Good	Good	Good
Firewalls	Excellent	Fair	Good	Fair/Good	Good	Excellent	Fair–Excellent
IDS/IPS	Excellent	Good	Good	Fair–Excellent	Fair/Good	Good	Fair–Excellent
Netflow	Excellent	Good	Fair	Fair/Good	Fair	Good/Excellent	Good
Sandboxes	Excellent	Fair/Good	Fair	N/A	Fair	Fair–Excellent	Fair–Excellent
Darknet	Excellent	Good	Fair	Fair–Excellent	Fair	Good	Fair
Passive DNS monitoring	Excellent	Good/Excellent	Good	Fair/Good	Good	Good/Excellent	Fair
Spamtrap	Excellent	Fair/Good	Fair	Fair	Good	Good	Good
Web Application Firewalls	Excellent	Good/Excellent	Fair	Fair	Fair	Good	Good
App logs	-	-	-	-	-	-	-
Antivirus	Excellent	Good	Good	Fair–Excellent	Good	Good	N/A

Table 2: Evaluation of categories of tools – summary

6 Services and tools recommended for proactive detection by CERTs

New and developing CERTs often have trouble acquiring useful data about incidents occurring on their networks. The main sources of information for such CERTs are usually incident reports submitted by users or clients in their constituency. But this information can be scarce, especially when the CERT has not yet established contacts and is not widely known. In such cases there is a strong necessity to supplement the (reactively) received data with additional sources of information.

This section covers recommendations for specific services (external data feeds) and tools that can be used by a CERT for the proactive detection of network security incidents. An inventory of such tools and services can be found in Chapter 5, along with an evaluation based on a specific set of criteria. The idea behind this section is to suggest the top 5 services and top tools that can be used by a CERT – particularly a national/government CERT – to jumpstart its services and get a good coverage of security incidents occurring in its constituency. This by no means implies that these are the best services or tools in any situation or for every CERT; nor does it mean that other services or tools are worse or should not be used. The rationale behind the selection is described in section 6.1, and additionally elaborated for tools in 6.3.

Note that in the process of the study, as shown in the shortcomings section (Chapter 7), it was discovered that many of the problems are not directly in the services and tools that give CERTs access to incident data, but in its subsequent processing.

6.1 *Explanation behind the rationale for selection*⁹⁵

Selection to the top 5 lists of services and top tool categories is influenced by a number of different factors:

- Results of the survey conducted among CERTs are taken into account. In this survey, we asked CERTs what services and tools they used most often. Thanks to this, we obtained a list of the most commonly used services and tools for proactive detection. In the survey we also asked the CERT community to rate these tools and services, based on a certain set of criteria. Finally we asked CERTs to mention explicitly what they thought were the best services and tools.
- Results of the desktop research. In this research, we reviewed many services and tool categories according to a similar set of criteria as in the survey and attempted to rate them. These ratings are of course subjective. Also, sources/tools may change over time, so the choice reflects opinions only at the time of the study.
- Results of comments from experts engaged in this study.

⁹⁵ Recommendations reflect the situation as of the time of the study – September 2011. Note that services and tools change over time; hence those presented here may become out of date.

- We stress the coverage requirement for services. That is, a service that has sensors worldwide as well as offering multiple incident types is preferable to a service that has limited sensors and offers just one incident type.
- We stress the openness requirement – particularly for a service – i.e. how easy it is for a national/government CERT in particular to gain access to the source (providing of course that they can prove their credentials to receive the data in the first place). This is related to a common problem for national/government CERTs: that often they can only receive data for AS or networks that they own.
- We take into account services that one has to subscribe to, by actively registering for the service – i.e. one does not get alerts automatically simply because one is listed in whois records as responsible for incident handling for a particular network/domain (as is for example the case with SpamCop). This is because a team that is listed as an abuse contact for a specific network will get incidents from such services automatically anyway, and thus become aware of such a source.
- For tools, we discovered that the situation is more complex than for services. The selection of tools is greatly influenced by the technical capabilities and networks that a CERT has access to. We therefore do not recommend a top 5 selection as is the case for services, but instead introduce three categories of tools/mechanisms that should be considered for implementation (for more details see section 6.3: Top ‘must have’ tools (and mechanisms) for proactive detection of network security incidents).

In the end, recommendations are subjective and ultimately based on the experience and opinion of the authors of this study and members of the expert group. It may be the case that depending on a situation the selection of choices may be different.

6.2 Top 5 services (data feeds) for detection of network security incidents

Based on the factors introduced above, the following 5 services are suggested for priority subscription for national/government CERTs as well as other CERTs/abuse-teams. By implementing these services first, a national/government CERT will be able to quickly gain insight into what is happening in their constituency. Note: this does not mean that other services should not be implemented.

6.2.1 Shadowserver Foundation⁹⁶

Rationale: The Shadowserver Foundation is a known and respected non-profit security organisation in the CERT community. It offers data on a daily basis that include amongst others: botnet, C&C, and DDoS reports. The data are of high quality and cover multiple types of incident reports, making it a very useful source of information. The service is free of charge and available to CERTs. As shown in the survey, it is the most often used and highly rated service by CERTs.

⁹⁶ <http://www.shadowserver.org>

For more complete details see 5.2.15.

6.2.2 Zeus/SpyEye Tracker⁹⁷

Rationale: One of the biggest threats on the Internet today is spyware – malicious software which spies on a user’s activities, attempting to steal sensitive information such as login credentials or credit card data. The current premier examples of such trojans are Zeus and Spyeeye, which focus on bank account information. Zeus/Spyeye tracker is a service that allows CERTs to obtain samples and gain an overview of the current behaviour of these trojans, as well as providing blacklists of IP addresses and domains used for C&C. The Zeus/SpyEye Tracker is the second most often used service by CERTs, scoring a lot of excellent marks in the survey. Note that this service is prone to changes in the threat landscape, as it is focused on two specific threats, which may disappear in the future.

For more details see 5.2.11.

6.2.3 Google Safe Browsing Alerts⁹⁸

Rationale: The most popular service dedicated to the discovery of malicious URLs is a must-have for any CERT. While subscribing to receive alerts about a network that is not owned by a CERT may be difficult – and thus a problem for national/government CERTs (but not for ISPs) – the scale and processing power of the biggest search engine enables relatively quick detection of malicious code on websites in a constituency. Ranked number 4 as the most often used service by CERTs.

For more details see 5.2.4.

6.2.4 Malware Domain List⁹⁹

Rationale: The Malware Domain List provides information about malicious domain names responsible for propagating malware on the Internet. The service also provides information about the classification source of each blacklisted domain. It is number 5 on the most often used service by CERTs list according to the survey, being one of the few services that scored no poor marks at all.

For more details see 5.2.13.

6.2.5 Team Cymru’s CSIRT Assistance Program¹⁰⁰

Rationale: While not mentioned in the survey because it was officially launched after the survey was completed, general experience with the datasets offered, along with their variety, position it as one of the top sources of information in the opinion of the authors of the study. Having access to high-

⁹⁷ <https://spyeyetracker.abuse.ch> , <https://zeustracker.abuse.ch>

⁹⁸ <http://safebrowsingalerts.googlelabs.com>

⁹⁹ <http://www.malwaredomainlist.com/>

¹⁰⁰ <http://www.team-cymru.org/Services/CAP/>

quality data covering different types of incidents is a must for a CERT. Team Cymru is a known and respected security organisation in the CERT community. Under the CSIRT assistance programme it offers daily lists of compromised or abused devices for the ASNs and/or netblocks within a CERT's constituency. It covers such incident types as bot-infected hosts, command and control servers systems, open resolvers, malware URLs, phishing URLs, and brute force attacks. The service is free of charge and available to national/government CERTs.

For more details see 5.2.23.

6.3 Top 'must have' tools (and mechanisms) for proactive detection of network security incidents

It is not an easy task to recommend the best tools for proactive detection of network security incidents. This is because access to tools that can be used by a CERT for direct monitoring and proactive detection of network incidents is highly dependent on the specific situation of a CERT and its constituency. CERTs may be faced with a situation where they lack direct access to a larger network and are often limited to the one in their office. Other CERTs may have many more possibilities at their disposal – including the monitoring of a core ISP network.

In order to address the above issues, for the purpose of this study recommended tools and mechanisms are split into three categories, based on their complexity and potential additional resources required – both network and human resources:

- **Standard tools/mechanisms** are those that are essentially part of the network infrastructure by design and thus at the disposal of nearly every CERT. Example: routers, firewalls, antivirus systems, IDS/IPS systems, netflow and various kinds of logs.
- **Advanced tools/mechanisms** are those that go beyond the standard networking tools available, usually requiring additional resources in the form of extra hardware, access to a core ISP network or DNS service or some additional IP space to operate. Examples include: darknets, server honeypots, spamtraps and networks of sensors.
- **Upcoming tools/mechanisms** require even more additional resources and skills to set up, often being still in an immature or unstable phase. Examples include: client honeypots, sandboxes, passive DNS analysis techniques.

The premise behind the above advice is that a CERT should start with using standard tools/mechanisms first, then move on to advanced ones, and finally to upcoming. There may of course be some limitations on specific tools which reduce their usefulness in a particular situation – or in fact, given the situation, they may not be available for use at all. Nevertheless, CERTs are advised to first explore all the possibilities related to a 'lower' category before moving to the more sophisticated one.

6.3.1 Standard tools and mechanisms

The survey itself very clearly showed that the most popular tools (and mechanisms) for the detection of incidents are those that are part of the network infrastructure in the first place, such as firewalls, antivirus systems, system logs and IDS/IPS. These tools are the most effective at detecting incidents directed at a particular network, whether it is a small office network or a corporate one. If a CERT can gain access to these logs – especially if they belong to the CERT network itself or to a larger host organisation – it is highly recommended that it uses them to discover attacks against itself. These tools can also be extended to monitor attacks against a wider constituency by observing malicious IP addresses being detected and blocked.

However, using logs from devices in the host infrastructure has its limitations. For one, depending on the scale of the network and variety of devices it may take quite some effort – and financial resources – to set up an adequate monitoring system. This includes potentially scripting a SIEM solution (or using an open source one) or buying a commercial SIEM solution for correlation. Maintaining it and also tackling false positives from potentially very different sources may generate a high cost for very large networks. Nevertheless, ignoring information from these ‘standard’ devices would be a mistake – in some cases a costly one, ending in the compromise of a CERT network.

Tools to consider in this category:

- Firewalls (see 5.3.4). According to the survey, firewalls are the most used category of tool for the detection of incidents by CERTs. This is probably due to the fact that firewalls are present in almost any network, so they can be relatively easily adapted for the detection of network security incidents. This type of tool could potentially generate many false positives according to the desktop research (the accuracy of results is considered fair). As with every raw data source, a firewall’s logs require some additional processing – such as working on parsers, whitelisting potential sources of false positives, etc.
- Antivirus (see 5.3.9). The second most popular source of detection of incidents. As with firewalls, these can be found everywhere – on gateways, servers and client desktop machines. Unlike most firewalls they focus on the malware. Compared to many other discussed solutions, they have a low false positive rate. Nevertheless, they often report incidents in different formats and attribute incidents to malware under different names, making parsing a potential challenge if multiple vendors are involved.
- IDS/IPS (see 5.3.5). Number five on the most popular list of sources. These focus primarily on attacks launched against a network. They have a higher false positive rate than antivirus solutions. Attack details are possible only if a connection handshake was not blocked by a firewall in the first place. When placed on a connection facing the Internet they may generate a lot of alarms, thus requiring tuning. Require additional work on parsers.
- NetFlow (see 5.3.6). NetFlow is one of the top mechanisms that can be used by a CERT for the proactive detection of incidents. In fact it was ranked number four in the survey in terms of usage. Not only can it be used for detecting new threats and infections in a network but also to perform post-compromise network forensics. Tools such as *nfsen/nfdump* allow for a relatively easy setting up of a monitoring framework, although scripting skills are required. On the

downside, NetFlow is most useful when collecting information within an ISP's infrastructure – which may be beyond the reach of many national/government CERTs. It also requires significant storage and processing power. Nevertheless, regardless of whether in an ISP environment or not, placing NetFlow sensors at strategic locations within an organisation's infrastructure helps increase situational awareness.

- Other tools that can prove useful: various kinds of logs (system, database, router, proxy) (see 5.3.12), WAFs (see 5.3.11). System logs are the third most popular source of information. Nevertheless, for the purpose of proactive detection of incidents they do require additional interpretation and parsing. It is recommended that a centralised location is used for collecting all logs and that logs are subsequently monitored. Web Application Firewalls may also be useful sources of information. They are relatively less common in a network infrastructure and require additional expertise in configuration.

6.3.2 Advanced tools and mechanisms

Once adequate monitoring of logs and alerts from the above solutions are implemented, there are a number of tools that can be additionally deployed by CERTs to pick up further information on incidents that may otherwise be missed by using only tools from the previous standard group. In some cases they can greatly enhance coverage of a particular incident type, make detection of certain malicious activities (like scanning) much simpler and allow more information relating to a particular incident be collected.

Tools to consider:

- Darknets (see 5.3.7). Darknets – or all sinkholes in general – are a very useful source of information about incidents for CERTs. One of the advantages is that they have a low false positive rate. A disadvantage is that a darknet requires some free IP space that can be used for monitoring purposes. Logistics may differ from easy to hard, depending on the scale of the darknet. However, once set up, log files will be very easy to parse as a CERT can decide on uniform format.
- Server honeypots (see 5.3.2). Server honeypots are in many ways similar in functionality to darknets. However, unlike darknets they are usually configured to interact with attackers. This means they can deliver more detailed information about an attack. Nevertheless – depending on the type of honeypot – they do require additional skills and resources to set up and manage. They also do not scale as well as darknets.
- Spamtraps (see 5.3.9). These can be a good source of information about problems in a constituency network and outside. Nevertheless, gaining good coverage is a challenge. The cost of setting up can be significant.
- Others: Networks of sensors. CERTs that do not have access to IP address space (or, in the case of NetFlow, a core network) necessary to deploy any of the above solutions should consider negotiating the possibility with their constituency to install a sensor in their network in order to improve their coverage. The type of sensor (presumably based on one or more categories of

tools mentioned in this report) can be negotiated with the constituency – if monitoring production environment traffic is a problem, consider using server honeypots or darknets.

6.3.3 Upcoming tools and mechanisms

This section covers even more advanced and elaborate technologies that can be used by CERTs to detect incidents proactively. These either require additional skillsets or very extensive resources to be effective, or are still relatively unstable. They are less understood because of a higher complexity. As shown by the survey they are not as commonly used by the CERT community as the tools mentioned in the other two categories. They are also the tools that CERTs most often cite when asked what tools they plan to use in the future.

Nevertheless, correct usage of these tools can result in valuable incident information that is often difficult to detect with the previous tools.

Tools to consider in this section include:

- Client honeypots (see 5.3.1). As shown in the survey, this is the least used tool category of all, tried in the past by numerous CERTs but abandoned. However, they are top of the ‘must try again’ list. This may be because in general they are complex to set up, prone to instability and need a lot of oversight – but are still seen as potentially very useful. They allow a CERT with limited IP space to monitor a constituency network for malicious activity – primarily malicious URLs. None of the previously mentioned tools allow for the direct detection of malicious URLs (though proxy logs or spamtraps, for example, can be used as input for client honeypots), which remain a very popular attack vector. Expanding the capabilities of client honeypots is an ongoing area of research in the CERT community.
- Sandboxes (see 5.3.3). The functionality of sandboxes overlaps somewhat with client honeypots. However, they mostly focus on identifying actions of a piece of malware. This can be a part of a proactive system of detection of incidents as well. A CERT can set up a farm of machines which are deliberately allowed to be infected with malware and subsequently monitored. IP addresses, domains and URLs contacted can be monitored and C&C servers and malware download sites and drop zones identified. Any hits on a constituency can be detected. However, setup of such a system requires considerable resources to be efficient.
- Passive DNS monitoring (see 5.3.8). Passive DNS remains one of the least understood technologies for the detection of incidents. Nevertheless, the DNS has shown itself to be one of the most useful sources of such information. Not every CERT has the possibility of setting up a passive DNS sensor – only access in some form to a busy recursive DNS server makes it worthwhile. Setting up a sensor may allow a CERT to gain access to worldwide networks of passive DNS servers, and hence more information.

7 Identification of shortcomings in the proactive detection of incidents

This section provides an overview and discussions of shortcomings identified during the study in the area of proactive detection of network security incidents by CERTs. It is split into two main categories: technical issues and legal/organisational issues. Each issue is described in the form of an OBSERVATION, identification of a SHORTCOMING(S), followed by RECOMMENDATIONS. Recommendations are aimed at the three identified broad groups of stakeholders: the data providers, data consumers, and EU and national government organisations, such as ENISA.

Overall 16 shortcomings were identified. Note that there may be some overlap between the different issues discovered. Some problems are more general than others, focused on weaknesses in particular technologies used.

7.1 Technical issues

Thirteen shortcomings of a technical nature were discovered during the study.

7.1.1 Data quality and reliability concerns

7.1.1.1 Observation

One of the main concerns of many CERTs is not the lack of access to data sources of network security incidents, but the reliability of the many available data sources. CERTs are expected to act on the received data, but this is often not possible if the data are of low quality. Furthermore, poor information quality inhibits automated handling of incidents because it requires manual verification.

7.1.1.2 Shortcomings

Data quality is influenced by a number of different factors. **Detection of incidents is inherently subject to both false negatives and false positives.** A false negative occurs when a security incident takes place but remains undetected. False positives on the other hand happen when an event is diagnosed as a security incident but in fact is not.

Assessment of false negative rates for external services that supply security information is a difficult exercise, beyond the scope of this study. CERTs can get a vague idea of false negative issues only if multiple external services that provide similar types of incident information and offer similar coverage report different events or when an incident is reported by the constituency but not by an external service even though its purported coverage should have allowed the incident to be detected.

It is easier for CERTs to rate false positives of a data source. False positives can occur as a result of weaknesses of the technology (tools) used to detect incidents. For instance, antiviruses may mistakenly trigger on general heuristics when assessing benign code. Thresholds may be set too low on an intrusion detection system that is configured to detect scanning activity – or the scanning in fact was due to a configuration error. A sandbox report may identify software that is attempting to

connect to a domain in our constituency and the domain thus reported as suspicious – but in fact the software that is submitted to the sandbox turns out to be benign. Some datasets are inherently more prone to false positives than others; for instance, data gathered through an experimental technology such as a form of shellcode detection. Others, like a sinkhole dataset, will probably generically contain few false positives.

Providers of information usually do not reveal from where and under what conditions the data were gathered. **Therefore the context in which an incident is detected by a data supplier is often unclear from the acquiring entity point of view. This influences the perceived reliability of data.** For example, the recipient may be unable to assess the likelihood of a reported IP being spoofed or not.

Another important factor is time. As one of the experts at the workshop formulated it: ‘time is quality’. **Timeliness of delivery to the receiving entity contributes to the perceived false positive rate of a data source.** Information about an incident is often a long time coming when using external services for proactive detection. A report dating back a few days, sometimes even hours – depending on the incident type – may be useless. Hence time is an extremely important part of quality. Even if the information is accurate, reliable and exhaustive, an overlong delivery time means that it may become out of date. As an example, when one receives information about some compromised host and the information reaches the recipient with some delay (perhaps caused by checking the information) it may turn out that the host no longer exists at all. Similarly, if malware is removed before an incident report is received, it may be counted as a false positive.

Time plays a role in another context too. **Data aging is another element of quality.** Data suppliers often provide information about detected incidents in the form of blacklists. Sometimes such lists are published on forums or blogs and then forgotten. Many providers of information concentrate on adding entries to their published blacklists, but remain unclear on how they age this information. Quite often the information on the lists quickly becomes out of date – thus influencing the perception of accuracy of the data provider. It is therefore very important to find out how the gathered data are aged: for example, how long an IP address, domain or URL remains on a blacklist and what procedure is used for its removal.

7.1.1.3 Recommendations

While achieving perfect data quality may not be possible, certain actions by data providers can help improve data quality and perceptions of their reliability:

- Data providers should carefully screen their data for potential false positives. The exact mechanisms applied here depend on the type of data and collection technology used. Therefore, apart from ensuring high quality, a data source should deliver information about the technology/method used to acquire information about the incident. This contextual information is extremely useful for interpreting information and assigning priorities to incidents by receivers of information. Such information can consist of types of systems used to discover the incident (e.g. honeypots, firewall logs), information about the definition of an incident (for instance, how many events blocked by a firewall constitute a scan?), whether the incident was verified with other major data sources, etc. These actions will greatly

improve the trust in the data of recipients, allowing for a wider adoption of a service and greatly improving the incident handling processes.

- Data providers should be encouraged to assign confidence levels (validity indicators) to information streams. What these confidence values mean should be clearly explained, together with recommendations on whether they think additional verification is needed by recipients. Development of community guidelines or standards in this area would be beneficial.
- Preferably the incident data should be distributed as soon as information about it is complete and available. Data providers should consider the trade-offs: the time it takes for them to verify their data vs. time of delivery to the recipient. A real time data feed is not always possible to deliver and is also harder to implement and handle on the client's side. Bulk data are much easier to maintain for both service and client but should not be delivered in periods greater than 24 hours. Otherwise data quality may be considered diminished and in some cases data are no longer actionable. Note that the probability of false positives is higher for raw, unanalysed information. However, usage of raw information makes it more probable that a new malicious activity is discovered in the first place. This can be tied in with the previous bullet point – raw data should be delivered with a different, lower, confidence level, rather than not being shared for fear of false positives.
- When maintaining blacklists, data providers should implement a data aging and removal process for the list entries and explain these procedures to the recipients. A suggestion here may be an algorithm that uses a decay rate. For an individual data element, considerations could include false positives and false negatives (changing over time) combined with feedback mechanisms from CERTs (see below). These would increase or decrease the data aging process.
- It would be useful if data feeds were also enriched with additional statistical information about incident types, including aggregated information about other constituencies. Along with data profiled for the specific client, the service should provide some statistical information on registered incidents. Such information is useful when assessing a threat scope in a global and local perspective.
- The data feed provider should create a process for data enrichment with information either from the service users or external services used as incident verification mechanisms. Data enrichment with external services may be troublesome mainly because of unspecified delays between posting data for classification and obtaining its results. Verification of data with other services is an important aspect but should not delay the data delivery. Additional information about incidents can be provided in a parallel data feed.
- Data providers are encouraged to adapt existing standards for the sharing of incident information (see 7.1.4).
- Data providers should provide easy-to-use mechanisms for recipients of the service to provide reports on detected false positive incidents in the data feed. In turn this will improve the data feed itself and reduce the number of false classifications. The service should also

provide information on threats that are no longer active. This data should be delivered either as soon as possible or in batches together with a standard data feed. This is a partial solution for the data aging problem and allows for the building of daily updated blacklists with the assurance that only active threats are taken into consideration.

Data receivers on the other hand are encouraged to develop methods and criteria for the evaluation of the quality of a data source. These should answer questions such as:

- Does a dataset contain entries that are there clearly due to error, and if so, how often?
- Does a data provider clearly explain how information about an incident type was collected and verified?
- How often were reported incidents found to be outdated on verification?
- Is a dataset suitable for automated processing without verification?

If the intention is to use a data source for blocking of some resources special care needs to be taken in order to evaluate quality of the data, as well as the appropriateness of data aging and correction/removal procedures.

The opinions of data users are of great value to the data providers, and should therefore be provided as feedback. Furthermore, data receivers are also encouraged to employ correlation techniques (see 7.1.2) to lower false positives and gain a better understanding of the context of an incident.

7.1.2 Correlation is still limited

7.1.2.1 Observation

‘Correlation, correlation, correlation’ (one of the experts responses when asked what tools are missing for the detection of incidents). As was pointed out, CERTs are often focused on detecting and remediating a single compromise rather than identifying and understanding larger events that encompass small individual attacks. Furthermore, even in the case of simpler incidents, it is worthwhile to correlate to gain better insight or eliminate false positives. Data duplication may also be detected in such cases.

7.1.2.2 Shortcomings

Incident correlation is the process of comparing different events, often from multiple data sources, in order to identify patterns and relationships enabling identification of events belonging to one attack or indicative of broader malicious activity. It allows us to better understand the nature of an event, to reduce the workload needed to handle incidents, and to automate the classification and forwarding of incidents only relevant to a particular constituency. It also allows analysts to identify and reduce potential false positives.

Correlation is useful both in the case of processing data from multiple tools on a monitored network and in the case of using multiple different external services that supply incident data. Based on the survey results, it is noticeable that most CERTs still do not deploy tools that enable correlation in an automated manner.

SIEM tools that perform correlation on the enterprise level have been available on the market for many years already, both commercially and as open source. These allow additional value to be derived from very varied datasets. However, the commercial solutions are often beyond the reach of CERTs due to their high costs. The open-source solutions on the other hand are often harder to manage. Some are being developed by CERTs themselves.

There is still no standard framework that defines how to get to the root cause of an incident by fully utilising all data feeds available to a CERT team. Emerging solutions that enable correlation of external services that provide incident data, such as Megatron or AbuseHelper (see 7.1.2.3) are becoming available now, but are still not in the mature phase. The need for such tools is recognised by many CERTs but they remain underdeployed.

Without correlation from multiple sources, both at the enterprise and external data source level, a CERT, especially working at the government and national level, might be unable to identify and eliminate the common problem behind the separate incidents affecting its constituency. This can lead to resources of a CERT and potentially involved stakeholders (e.g., vendors, law enforcement, ISPs) being spent in dealing with consequences and not the root cause.

7.1.2.3 Recommendations

Data providers are encouraged to employ correlation to remove false positives and duplication of data (see 7.1.1). The data consumers are encouraged to implement their own solutions for verifying datasets to help improve quality of data before forwarding them to their constituencies.

Aside from verifying the correctness of information, stakeholders should try to implement mechanisms correlating events in received datasets and output from their own monitoring solutions. This step can be performed when data are to be consolidated and stored in an internal database. Extracting common behaviour patterns and relations between incidents is no trivial task. Fortunately there are both open-source and commercial solutions available on the market which can aid the process (open-source solutions are listed below).

If data providers implement feedback mechanisms then information about detected false positives should be posted back. Such information can be used by the data provider to evaluate its algorithms and improve them, making the dataset more reliable.

Some open-source solutions that can be used for correlation (by no means an exhaustive list):

- Generic tools:
 - SEC¹⁰¹
 - Loghound¹⁰²
- SIEM tools:

¹⁰¹ <http://simple-evcorr.sourceforge.net/>

¹⁰² <http://ristov.users.sourceforge.net/loghound/>

- Cyberoam iView¹⁰³
- AlienVault Open Source SIEM¹⁰⁴
- Tools specifically for incident handling or analysing larger datasets:
 - AbuseHelper¹⁰⁵,
 - Megatron (contact SITIC for availability¹⁰⁶),
 - BGPrank¹⁰⁷
 - CIF¹⁰⁸

7.1.3 Lack of automation

7.1.3.1 Observation

‘We are focusing now on tools to automate as much of the incident analysis process as possible. We believe too much valuable analyst time is wasted performing tasks that can be automated’ (A comment provided in the survey).

7.1.3.2 Shortcomings

In the survey that was part of this study, five teams had indicated that they cannot handle even the number of incidents that they currently have. At the same time all these teams recognise that there is room for improvement in terms of the correlation of incidents and automation of this process. Many teams indicated in the survey that there is a huge need for automation of processes because ‘manual processing does not scale’. Some teams also pointed out that lack of automation can be the most important obstacle in effective incident handling, because plenty of information is readily available from many external sources or can be relatively easily collected from internal monitoring systems. Thus, it becomes crucial to effectively filter this information to be able to identify and prioritise security incidents and focus on those that pose a larger security threat.

At the same time, there are a number of inhibitors preventing effective automation, including:

- Lack of appropriate tools. This has led to many teams using in-house developed tools and scripts for data processing and correlation. In some cases, this is attributed to the specific needs of a given CERT or its constituency. In other cases teams claim that they are unaware of existing solutions and it is easier to write their own tools than to properly research the market and test ones that might be available. In fact, many tools for correlation of security events are commercial closed solutions that do not adapt well to the CERT environment beyond large enterprises. In answer to this, several tools have been developed recently by

¹⁰³ <http://www.cyberoam-iview.org/>

¹⁰⁴ <http://alienvault.com/products/open-source-siem>

¹⁰⁵ <http://www.abusehelper.be/>

¹⁰⁶ <http://www.sitic.se>

¹⁰⁷ <https://github.com/CIRCL/bgp-ranking>

¹⁰⁸ <http://code.google.com/p/collective-intelligence-framework/>

the CERT community with the specific goal of meeting the needs of incident analysis and correlation. AbuseHelper and Taranis (for details of this tool, please contact Govcert.nl) should be mentioned among these few tools. Their success depends heavily on their reception and adoption by CERTs and may impact teams' capability for data processing and correlation.

It should be noted that development of tools is not the core job of CERT teams (or at least, should not be). In reality, due to the facts mentioned above, many teams need to use their resources for development, which obviously creates an additional cost to their host organisations.

- Significant maintenance cost of existing tools. The most common weakness of tools used for correlation, as indicated in the survey responses, is their maintenance cost. While the benefits of correlation and automation are clear in the long run, some teams do not have the resources to set up and maintain tools properly in the beginning. This is also impacted by the fact that there are very few if any tools that are widely adopted in the CERT community and that have good level of community support.
- Incompatibility between tools. 71% of the responders made some custom changes to the tools they use. Only 41% of those changes were done in order to provide some additional functionality to meet specific constituency needs. The remaining 59% were drawn by the need to integrate tools into correlation systems, generate adequate reports, tweak output formats, etc. An obvious consequence of the fact that so many teams invested their efforts into customisation of tools is a loss of resources that would be needed elsewhere. Many tools were not made with integration into more complex systems in mind. Also, teams vary greatly in their expectations and needs. Even such fundamental concepts as what is an incident and what is not are not common to all teams. Obviously, priorities are also different for the same types of events in different CERTs – depending on the type of their constituency, resources available and other factors. The lack of adapted common format for exchange of information about security incidents is also a possible problem affecting possibilities of integration of tools and growing need for customised parsers and plugins. See 7.1.4.

7.1.3.3 Recommendations

Data providers are encouraged to adopt common formats for exchange of incidents (see 7.1.4), while data consumers are encouraged to deploy correlation tools (see 7.1.2). They should attempt to add plugins to handle new sources of information rather than building new tools from scratch.

7.1.4 Lack of common formats

7.1.4.1 Observation

There is no standard adopted format for exchange of incident-related data.

7.1.4.2 Shortcomings

An existing IDMEF¹⁰⁹ format for intrusion detection messages is commonly used by IDS/IPS sensor vendors, but is not adequate to handle data related to incident response in a more general context. Another proposed XML-based format, IODEF,¹¹⁰ developed between 2001 and 2008, turned out to be too complex and never actually went into the mass-scale adoption phase because development of tools that would fully support IODEF was beyond the reach of most CERTs. Until now, most teams have relied on free text email exchange for inter-team data exchange. Some other protocols and messaging formats such as X-ARF¹¹¹ or XMPP¹¹² may replace or complement it if widely adapted in the near future.

There are also many text-based file formats for different automated data feeds. Although this may be seen as an issue, because each file format generally requires a different parser to process and unify the data with other sources, it may not necessarily cause big problems. Some of the observations during the expert group meeting were that if a data source is stable (e.g. it does not go up and down or change formats often) and valuable (e.g. if a lot of data are delivered and they meet high standards), cost of creating a custom parser for this feed is bearable when compared to benefits. This is especially the case when the file format is simple. For example, CSV files can be easily parsed even with command line tools and data can be easily accessed even if stored in flat text files. On the other hand, using complex XML-based formats will call for more high-level parsers and probably database type of storage for data.

A real-life example of a reasonable approach to data sharing is a feed that provides a set of CSV text files with raw data, updated every 15 minutes, available for download from a defined location. Limited access is granted based on client certificates. This model allows for easy setup and almost no running cost for the client using the feed.

An example of an approach that is not recommended (but sometimes found) is a data source that sends emails every few days with text files in a custom format. The description of content is provided in the first few lines of the file, which is then compressed with tar and gzip and encrypted with public GPG key of the recipient. While not the worst possible scenario, this one requires much more effort to implement and handle (e.g. automated decryption of messages can be an issue on the client side) and the infrequency of data deliveries may render the source not worth the effort, especially if the format gets changed.

7.1.4.3 Recommendations

It is recommended that data providers stick to lightweight file formats (e.g. CSV-based) and authentication mechanisms that can be easily implemented in an automated workflow. If this is matched with valuable and frequently updated content, the recipients are very likely to adapt to any

¹⁰⁹ The Intrusion Detection Message Exchange Format (IDMEF): <http://www.ietf.org/rfc/rfc4765.txt>

¹¹⁰ The Incident Object Description Exchange Format: <http://www.ietf.org/rfc/rfc5070.txt>

¹¹¹ x-arf - network abuse reporting 2.0: <http://www.x-arf.org/>

¹¹² Extensible Messaging and Presence Protocol (XMPP): Core: <http://www.ietf.org/rfc/rfc3920.txt>

delivery method proposed by the vendor as well as invest resources to create interfaces to their own systems.

The format of data should be characterised by a set of features which allow easy reprocessing and correlation with other sources. Basic information the data source should deliver are Autonomous System Numbers, IP addresses, domain names, timestamps of incident and of course the category of incident or a set of tags/labels describing it.

The data should be distributed by a standard and secured communication protocol such as HTTPS, SFTP or SCP. Using standardised tools will ensure that none of the registered users will have major problems with the underlying technology and the shared data will remain secured when in transport. Note also that it may be possible to use the existing infrastructures of established data providers to distribute information – new data providers are encouraged to explore such an option.

The data receivers are encouraged to try to consolidate information into one consistent internal format. Having such a database will ease the data management process and integration with incident handling systems. It will also allow changes to be tracked in the whole dataset, provide an overview of the whole observed network and act as a research base on evolution of threats in the constituency network.

Development of such systems as the one above may be beyond the reach of smaller CERTs. Fortunately, solutions that attempt the above are already being developed as open-source solutions by CERTs. One such example is Megatron (introduced 7.1.2.3).

7.1.5 Lack of own monitoring (sensor networks)

7.1.5.1 Observation

Some CERTs do not operate their own sensor networks (or perform any sort of independent detection of incidents).

7.1.5.2 Shortcomings

The survey revealed that some CERTs do not operate their own sensor networks. The lack of own monitoring by many CERTs makes them entirely reliant on external services or incident reports (reactive). This means that CERTs may be late at detecting some types of network incidents and may lack the ability to verify through their own systems what others are reporting. This issue is sometimes related to lack of resources – lack of manpower or simply financial. Quite often it is related to the fact that CERTs do not manage their own network or AS or have access to the networks of their constituency, which is often the case for government/national CERTs and/or small CERTs.

7.1.5.3 Recommendations

The organisations that have the resources and means to implement internal system for monitoring network for incidents are encouraged to do so. Comparing information received from data providers with data from sensors distributed in a constituency network is the best way to verify quality of given

data. And in case the sensors miss some important information, their configuration can be improved to produce data more consistent with the received ones, thus improving the detection rate of internal systems. CERTs should try to deploy their own sensor networks even if they do not manage their own network or AS.

This can be achieved through partnerships with their constituency or with stakeholders at the national or government level. Networks of sensors can take various forms. They need not always consist of sensors that monitor production level traffic or use existing network device infrastructure (an issue that often raises privacy concerns). Some may take the form of honeypots and therefore not monitor traffic at the production level (may require legal consultation). Free open-source solutions such as SURFids¹¹³ allow for easy deployment of IDS sensors based on honeypots with a low false positive detection rate. CERTs can also consider joining existing worldwide sensor deployment initiatives – often at the cost of deploying a sensor in their network they gain access to a vast array of results from other sensors worldwide (CERT.br runs a spampots project, for example).

7.1.6 Lack of client honeypot deployments and sandbox analysis capabilities

7.1.6.1 Observation

CERTs lack client honeypot deployments and have few sandbox analysis capabilities of their own.

7.1.6.2 Shortcomings

While not exactly the same, client honeypots and sandboxes are two complementary technologies – one more focused on mechanisms used for infection, the other on the actions of actual malware itself. In fact, sandboxes often have client honeypot functionality.

Having the ability to independently and automatically scan websites in a constituency for malicious content is a useful functionality for a CERT team. While search engines such as Google scan the web regularly for malicious content and alert web page owners of such content, they may not be as fast at detecting malicious content on critical servers. Commercial companies can also offer such a service, but not all CERTs are able to meet the financial costs. Automated scanning of websites for malicious content is possible through the deployment of client honeypot technologies. In these solutions, various browser emulation techniques or browsers driven in real environments are used to make an assessment on the maliciousness level of a web page. As the survey showed, CERTs are interested in using this technology but find it difficult to deploy and manage. Part of the reason is that these technologies remain complex and are sometimes unstable.

Similarly, having the ability to utilise sandbox technologies in-house to study malware can greatly enhance a CERT's operational capabilities, including the enablement of proactive detection: infected hosts can be observed for sites they attempt to connect to for instructions or attack. This is a functionality not normally possible when using online services (such as Anubis¹¹⁴) (it would require complex additional parsing, would not provide long-term monitoring and would be reliant on

¹¹³ <http://ids.surfnet.nl>

¹¹⁴ <http://anubis.iseclab.org/>

processing times offered by a service). Unfortunately sandbox technologies, as the survey showed, are rarely deployed by CERTs.

7.1.6.3 Recommendation

Detection of malicious URLs and discovery of malware in a constituency is possible through the deployment of client honeypot and sandbox technologies. They are a useful solution for CERTs that do not have vast IP addresses that they can monitor for attacks. Nevertheless, the deployment of these requires dedicated IP addresses. Relevant URLs and collected binaries can be fed to such systems and potentially compromised sites, exploit sites and C&C servers identified – without being reliant on external services only. Since these solutions still have multiple shortcomings, CERTs are encouraged to actively participate in their development.

7.1.7 Visualisation still underutilised

Security data visualisation is an important aspect of any system which aims at providing useful content in an easy-to-interpret way. The ease of interpretation is the major key because it directly reflects time needed to assess the severity of a threat and provide a proper reaction. Practical visualisation of data produced by security systems is even more difficult to achieve because the information is in most cases sensitive and private. The legal aspects of sharing data directly reflect its representation, which needs to be prepared for different audiences – it will be different when shown to the general public than when viewed by analysts with proper clearance to access classified data.

7.1.7.1 Observations

A problem faced by many organisations nowadays is the poor quality of data and lack of standard form allowing easy correlation. In most cases information needs to be verified and sometimes enriched to allow meaningful and easy usage. Lack of standard form also hinders visualisation. Many tools come with almost no graphical interface and allow just simple logging of events. This can impose difficulties and usually forces an organisation to direct its resources on finding effective way to represent gathered data. Very often data need to be reprocessed and analysed before visualising. This is especially true when detecting trends or enriching dataset, for example with geographical information. Effective visualisation techniques are particularly needed when dealing with incidents which need to be correlated with many heterogeneous data sources. This allows quick assessment of a threat and proper mitigation.

7.1.7.2 Shortcomings

There is no single good way to visualise security data. The knowledge needed to build a proper solution is interdisciplinary and often comes not only from a security expert but is extended with input from social studies and research about human perception. Such an approach gives great results but the devised method is usually applicable to a particular dataset and when it changes the visualisation also needs to be revised and changed. Creating a meaningful visualisation of information usually takes considerable resources but in exchange provides a powerful tool for analysts and incident handlers.

7.1.7.3 Recommendations

A consistent internal repository of data as recommended in 7.1.4.3 is also a great base for developing a visualisation solution tailored to the needs of the organisation. Such software can help isolate threats found in the dataset and assign priorities, thus helping in quick reaction and mitigation of incidents. Incident handlers using such software will spend less time gathering information about the threat from the database and can focus on removing it as fast as possible.

Organisations willing to build visual representations of their datasets should try to focus on visualisation tools and methodologies prepared especially for security and network traffic analysis data. One of the good places to start research is the 'Security Visualization – secviz.org' web page. It is a place where experts exchange their knowledge and thoughts on applied visualisation techniques. Another place to start is FlowingData.com, where one can find helpful articles and exchange information on forums.

There is no ready-to-use tool which will be versatile enough to allow visualisation of any dataset, but there are some ready-to-use tools which can be used as a reference or a starting block when building a specialised solution. One of the tools is a Linux distribution called DAVIX. It is a Slackware-based LiveCD equipped with many tools helpful in data analysis and visualisation.

7.1.8 Lack of incident reporting tools integrated with users' desktop software

7.1.8.1 Observation

Although there are many incident reporting tools, there is still a lack of solutions that give an average user the capability to report an incident in the simplest way – through his/her desktop software (this was pointed out in the survey).

7.1.8.2 Shortcomings

Most of the data shared by CERTs come from automated data feeds. Operators of these data feeds (identified in this study), collect information in various ways, often also automated – for example through darknets or honeypots. Other reports by these operators are collected manually, either using their own observations or from reports coming from individual users. One should keep in mind that these users are a valuable group who can provide incident data in a quick and efficient way. A user visiting, for example, a web page which suddenly turns out to be malicious can report it quickly, providing that this action does not cost him much effort.

The average Internet user's possibilities to report the accident are limited because of their lack of knowledge about the right contact to report the incident to (such as a CERT in their country) and the inconvenience of this action (many are simply not interested or consider it futile).

One-click report sent through a plugin integrated with most commonly used software could change the way the incidents are currently reported and encourage users to do so because of the simplicity and convenience of this feature. This would allow collectors of this information to build new datasets which in turn could be shared with CERTs in an automated manner. Note that community-driven reporting has its downsides too, including a potentially high number of false positives.

The question of which entity should provide this kind of tool, and gather and analyse the data coming from this, remains open for discussion. Another question that needs to be answered is: would CERTs be ready to process this kind of information, which could possibly be of a high volume and low quality?

7.1.8.3 Recommendations

The possibilities of using incident reporting tools by end users, and ways to develop and deploy such tools requires further analysis. As mechanisms for this functionality are in general available in AV/IDS tools, vendors are encouraged to share such data gathered with relevant CERTs/abuse teams.

7.1.9 Lack of long-term trend analysis of incident data

7.1.9.1 Observation

There is a lack of long-term trend analysis of incident data by CERTs

7.1.9.2 Shortcomings

CERTs do not publish long-term trend analysis of incident data. While trend reports are published by many vendors, analysis of trends based on incident data received by CERTs, especially by national/government ones, is generally missing. This is a significant gap in the security community as these kinds of reports would be viewed as much more objective than vendor publications. It would also allow other CERTs to adapt their operations to improve incident handling.

However, long-term analysis of data is a problem for CERTs for various reasons: there is often a lack of resources (both in terms of manpower and financially) and tools necessary to carry out such research. These tools usually have to be developed from scratch by a CERT.

Furthermore, to make the research results more useful, it would be worthwhile comparing trend reports published by various CERTs. This is problematic, however, as there is no common understanding of the term 'incident', not to mention incident types. Other issues, such as different collection methods and capabilities, may mean that statistics are skewed in some manner – leading to comparisons of 'apples' with 'oranges'.

7.1.9.3 Recommendations

CERTs are encouraged to perform long-term analysis of their datasets and publish the results. Convincing all CERTs to use a single taxonomy for incidents does not seem achievable and projects that had this as one of their goals, such as EuroCERT or eCSIRT.net, have mostly failed in this regard. However, keeping the incident 'types' as generic as possible and the framework for collective reporting as lightweight as possible could still yield very informative content at a very low price. As an incentive to CERTs, ENISA could publish collective trend reports based on data that national, government and other CERTs collect about their constituencies, naming participating teams. These trend reports could influence the future objectives of ENISA, giving first-hand objective information about state of security and saving resources on costly research.

7.1.10 Targeted attacks underreported

7.1.10.1 Observation

In the experts' opinion, there is little or no detailed incident information available or shared actively on attacks targeted at specific organisations or groups of organisations.

7.1.10.2 Shortcomings

While there is a lot of information and data feeds available on general attacks on a constituency, there are no feeds dedicated to information on more specific attacks against organisations. The goal of these attacks is usually different from the goal of mass exploitation attacks, even though the underlying techniques themselves may be similar. While the general attacks are motivated by cybercrime and the underground economy (often resulting in bots sending spam and harvesting financial credentials), targeted attacks are conducted as part of cyber-espionage activities (which may also include financial gain as an ultimate goal) and long-term invigilation operations known as APT (Advanced Persistent Threat).

Lack of such information is of serious concern to CERTs – particularly national/government ones – which will remain unaware of them or unable to react early to a potential threat.

7.1.10.3 Recommendations

Research should be conducted into finding ways to report targeted attacks affecting a select group of organisations. As the nature of these attacks is such that they normally have to be reported by the targeted organisation (it is less likely that such attacks will be noticed by the outside observer, as is the case of mass exploitation attacks) and have more serious repercussions than mass exploitation attacks, additional barriers for information exchange are introduced. These barriers – often legal and political – need to be overcome before automated redistribution of such attack information becomes possible. For instance, experts pointed to the following issues: if an organisation does provide such information, is it admitting to a data breach? If reported, does a formal complaint to law enforcement have to be filed? Will the fact that the incident was disclosed encourage repeat or copy-cat attacks? Potentially trusted information brokers could be set up to allow such attacks to be reported anonymously, but this needs further analysis (see also 7.1.13).

7.1.11 DDoS attacks underreported

7.1.11.1 Observation

Some survey respondents felt that DDoS attacks are being underreported.

7.1.11.2 Shortcomings

The survey showed that the DDoS attack category is one of those most reliant on reactive reporting (i.e. a constituent – the victim – first detects the incident and then informs a CERT). Therefore, if a constituent does not report an incident and the responsible CERT does not have tools deployed on the path to their network that can potentially signal an attack, these types of events go unnoticed.

There are closed data feeds that report DDoS attacks related to a particular constituency. However, it is noticeable that such reports are much more rare compared to other incident categories. Based on these data feeds, they would appear to be noticeably less popular than media reports or conventional wisdom would suggest. Indeed, the CERT Polska Annual Report for 2010¹¹⁵ mentions receiving just 11 such incident reports concerning Poland from automated external data feeds (out of over 12 million automated reports), and 11 different incident reports reported directly by its constituency (out of 674 incidents in total). This may indicate that there are in fact many more DDoS incidents going underreported.

One of the reasons that DDoS attacks may be underreported is the difficulty of detecting them by an observer. The major incident data feed providers detect such attacks by monitoring botnet channels for botmaster commands initiating DDoS attacks against IP addresses or domains. Alternatively, such attacks can be detected by analysing backscatter traffic to darknets. They then match the IP addresses or domains against CERT constituencies and forward the information to the responsible CERTs.

Even when actively monitoring a network of a constituent (hence being able to detect many volume-based attacks) it is often difficult to identify Layer 7 DDoS attacks. This is because they attack weaknesses specific to an application at the other end, and hence are not understood by an external observer.

7.1.11.3 Recommendations

Data providers should focus on improving their methods for detecting DDoS attacks in an automated manner. A good starting point, at least for detecting volume attacks, would be NetFlow and nfsen/nfdump.

7.1.12 Passive DNS monitoring underused

Passive DNS (pDNS) monitoring or passive DNS replication is a technique used to collect DNS data for the purpose of discovery and analysis of threats. The technique is based on monitoring and storing DNS queries and responses. Information is put into a database which can be used to find resources used by malicious domains and track how threats are changing. Gathered information can also be used to analyse sophisticated threats which use DNS system as an element of their malicious behaviour. One of the examples could be the Morto worm which queried the DNS for TXT records of particular domains in order to receive and execute new commands.

7.1.12.1 Observation

The survey showed that Passive DNS replication is still underused, despite being very helpful in tracking malicious domains, worm and virus propagations and post-mortem analysis of attacks.

¹¹⁵ http://www.cert.pl/PDF/Raport_CP_2010.pdf

7.1.12.2 Shortcomings

Despite being one of the building blocks of the Internet and literally everywhere, DNS remains rarely used for network security monitoring (as shown in the survey).

There are two ways of implementing passive DNS replication which rely on monitoring traffic either inside the organisation's network (pDNS sensor is located on the local interface of recursive DNS server) or outside it (the sensor is located on the outbound interface of recursive DNS server). Inside monitoring assures that all resolution attempts will be logged and gives great results when tracking infections, especially for workstation computers or detecting sudden spikes in number of DNS queries. Unfortunately the volume of logged traffic can quickly become overwhelming, which makes resource requirements fairly high to sustain data for long enough periods of time for analysis it to give feasible results. The other concern is privacy of users from whom the traffic originated. Inside monitoring makes it possible to attribute each DNS query to a specific internal host, which of course is not always needed. Outside monitoring of DNS traffic can also be useful and resolves privacy concerns, because observed traffic is aggregated for the whole constituency network. The outside sensor unfortunately cannot observe the true volume of traffic generated by the users because of caching mechanisms on the recursive server. Due to the much smaller traffic volume, fewer resources are required and data management is much easier than in the case of the internal monitoring approach.

An organisation willing to make full use of data gathered by pDNS sensors needs to implement proper analysis methods and take into account legal considerations, which always create additional demands on resources.

7.1.12.3 Recommendations

Passive DNS traffic monitoring can be a powerful weapon in identifying malicious activity on the network. Organisations willing to share pDNS data are advised to consider joining a DNS traffic monitoring project, for example Secure Information Exchange managed by ISC.¹¹⁶ It involves setting up a sensor which forwards data to the project's database for analysis. The raw data gathered in the database are accessible for approved parties. Joining such an initiative can gain access to much larger datasets, enabling better coverage.

7.1.13 Lack of services for data leak reporting (data repatriation)

7.1.13.1 Observation

Leaks of private data are becoming an increasing problem. They may occur as a result of a database compromise, a disgruntled employee, or a number of other issues including trojan horses, key-loggers and APT attacks. Data that are subject to such leaks include access credentials, email and postal addresses as well as financial data. Examples of such leaks are the database dumps published by Anonymous in early 2011 or collections of data found during occasional takeovers of botnet command and control servers. In the latter scenario, data are often shared with appropriate

¹¹⁶ <https://dnsdb.isc.org/>, https://sie.isc.org/wiki/Passive_DNS

administrators (e.g. banks, webmail services, e-store owners) who in turn notify their customers. This can be done via the national/regional CERT or directly. Direct contact may require too many resources because, for example, a single leaked document can contain data belonging to many different organisations and it would be necessary to find contact details and inform them one by one. Quite often data are simply found published on an open forum or a board like Pastebin and thus remain uncontrolled. Site owners usually just provide a means to publish leaked data but because of resource issues do not sort out the published data and do not inform involved parties, which remain unaware if not monitoring specific sites themselves.

7.1.13.2 Shortcomings

While data leaks can have serious consequences to both data owners and administrators, there are very few monitoring services that look for this kind of information and allow data owners to be notified when non-public data that possibly leaked from their organisation is found. An example of such a service is the BIN feed¹¹⁷ provided by Team Cymru to vetted and verified financial institutions to inform them whenever compromised accounts' credit card data are found. As a consequence, when stolen data are discovered, there are very limited possibilities to report this other than directly to data owners. As mentioned before, this can be very bothersome and is not a very encouraging perspective. Several survey respondents have identified data leaks as a type of incident that is heavily underreported.

7.1.13.3 Recommendation

Exactly who could set up a trusted information broker to handle such data and distribute it to relevant parties is unclear – and something that could be further investigated (see also 7.1.10). National and government CERTs are usually ready to handle bulk data stolen from services in their constituency. They should have appropriate contacts or be able to reach out to appropriate persons promptly. Obviously, such data must be handled carefully, confidentially and on a need-to-know basis only. Because in many cases data which were subject to theft are clearly private, CERTs should be allowed to legally handle them for this specific purpose and in this context.

7.2 Legal and Organisational issues

The following 3 legal and organisational issues were identified in the study:

7.2.1 Legal issues impede data sharing

7.2.1.1 Observation

The study has identified legal concerns as one of the major obstacles for data sharing. Data such as IP addresses, URLs and timestamps are inherently included in the operational set of information about incidents. Other examples of sensitive data are mentioned in 7.1.13.

¹¹⁷ <http://www.team-cymru.org/Services/BINFeed/>

7.2.1.2 Shortcomings

In some jurisdictions and some contexts, the data mentioned above are considered personally identifiable information and thus their processing is subject to certain legal requirements. This impacts not only possibility of sharing data with others, but in some cases means that CERTs are unable to receive data feeds from third parties. In fact, during the expert group meeting some vendors reported that they have experienced cases where CERTs refused to receive incident data concerning their constituency because of legal considerations. Even in countries where laws do not clearly regulate processing and sharing of such data, exchange of information including IP addresses, URLs and timestamps (when used in the IP context) is often considered by CERTs to be a 'grey area'.

7.2.1.3 Recommendations

During the workshop, many experts expressed the opinion that a balance between privacy protection and security provision needs to be reached. In particular it was mentioned by experts that ENISA could potentially play a role in helping to clarify how operational data including IP addresses should be handled by the CERT community. European and/or national law should make it as clear as possible what (sub)sets of information, in what contexts and under what conditions, can be legally processed and exchanged by actors such as:

- data providers – parties that collect information from network traffic, servers and/or user machines and are willing to share parts of their data for the benefit of Internet security;
- data consumers – including, but not limited to:
 - ISPs that use these data in their operations, in order to help in providing the security of the service they deliver to their customers, which may include limiting access to harmful resources from their networks, inform infected customers, help to clean their machines
 - AV and other security products vendors
 - Vendors of web browsers and other software that actively processes web resources;
- Intermediaries, such as national or government CERTs which may facilitate relevant data sharing between different communities and their local constituencies.

Unfortunately, the identification of specific gaps in European and national laws as well as recommendations for concrete regulations is beyond the scope of this study.

ENISA in 2011 is running another study on the Legal and regulatory aspects of information sharing and cross-border collaboration of national/governmental CERTs in Europe, which will focus more on the legal issues.

7.2.2 Lack of human resources

7.2.2.1 Observation

Many teams lack enough human resources to deal with all incoming incident reports as well as information from data feeds. 56% of the responders to the survey indicated that they can handle only higher priority incidents or are unable to process all incoming information.

7.2.2.2 Shortcomings

Lack of resources negatively impacts the ability of CERTs to deal appropriately with all incidents. Even when prioritisation is applied to deal with high-priority incidents, there is always a risk of losing valuable information. Lack of resources to deploy tools and communication channels was also indicated in the survey as preventing teams from sharing data.

7.2.2.3 Recommendations

The most natural way in which CERTs adapt to lack of resources is prioritisation of incidents. Based upon type, severity, scope and many other factors, decisions can be made about which incidents should be handled immediately. Prioritisation is a routine part of incident triage and it is discussed in training courses such as TRANSITS,¹¹⁸ ENISA Exercises¹¹⁹ for CERTs and many others.

Another way to deal with lack of resources is through automation. It can be observed from the responses to the survey that teams deploying automated methods for information processing and correlation have fewer issues with resources. As creation and maintenance of tools for automation is itself a resources issue, it would be of great benefit to make a collective use of tools that are developed in the CERT community and dedicated for incident handling purposes. This can involve recently created tools such as AbuseHelper or Megatron. CERTs should also stress the need for appropriate staffing, including people with not only incident handling skills but also the capability to discuss system administration, programming, etc., when communicating with the managements of their hosting organisations.

7.2.3 Obstacles in reaching closed groups

7.2.3.1 Observation

Many sources of incident-related information that are considered valuable can only be accessed by members of closed groups, upon recommendation or in exchange for own data. This is a problem for new teams which have not developed own monitoring infrastructure and are not well known (and trusted) by any community.

7.2.3.2 Shortcomings

While it is important that the data are shared with trusted parties and only with those who can use it to benefit their network and Internet security, trustworthiness is not something that can be easily

¹¹⁸ <http://www.terena.org/activities/csirt-training/>

¹¹⁹ <http://www.enisa.europa.eu/act/cert/support/exercise>

evaluated. At the same, for newly established teams, lack of external information often means that there are no incidents to act upon, less visibility in own constituency and less success stories, which translates into less visibility also in the security community. This can cause a 'chicken and egg' scenario. In particular, the cost of joining some membership-based groups may be too high for new teams or teams from developing countries.

7.2.3.3 Recommendations

There are many public and open information sources that can be a good starting point to build upon. Those which push out data (such as SpamCop, rely upon current full contact data published in appropriate places. Therefore, creating an IRT record in RIPE database and registering it with appropriate netblocks, entering contacts in abuse.net and similar databases is essential. Some of the databases, such as Trusted Introducer list of 'Listed teams' involve a vetting process, where two or more trusted members are required to vouch for a new entry. Therefore, it is essential to gain recognition in own constituency and neighbourhood as early as possible. Membership in CERT groups and forums can be a way to meet this goal as well as to be introduced to data feed vendors and exchange contacts with other teams. While some of these groups involve membership fees, there are also forums such as TERENA TF-CSIRT as well as many national-level ones which are free to participate in (apart from possible travel and lodging costs).

8 Summary of recommendations

This section provides a summary of recommendations based on presentations and discussions in the previous section. The idea behind this section is to provide readers with a quick lookup guide to the main recommendations for three main stakeholder groups: data providers, data consumers and institutions at the European/national level.

8.1 Recommendations for data providers

Proactive network incident data sharing between CERTs is one of the most important elements allowing successful cooperation in the fight against Internet threats. Despite attempts to develop standards (such as IODEF or X-ARF) there are no widely established rules on how CERTs should exchange information about various types of detected incidents in a secure and common way (see 7.1.4). The survey and the workshop meeting provided valuable insight on expectations demanded of a high-quality data source, of which the most important ones are discussed here.

8.1.1 Identification and vetting of data consumers

To meet privacy concerns about shared information, the data providers should follow the set of recommendations defined below.

- potential data recipients can be screened for eligibility; each screening process should take into account demands of data receivers regarding access to information and determine whether such access should be granted to full dataset or just part of it (see 7.2.3)
- data providers should establish contacts with security institutions and communities such as FIRST; becoming a member of such community often involves establishing trust relationships with other members (see 7.2.3)
- data providers should create an easy process of registration; email or web page interface are the most common and convenient ways of interacting between data provider and clients (see 7.2.3).

8.1.2 Data format and distribution

Data that are shared between a publisher and its clients need to be delivered in one of the commonly accepted ways and formats. Below are the recommendations describing data sharing principles for high-quality data sources:

- data providers are encouraged to adapt existing standards for the sharing of incident information (see 7.1.4). Providers should use one of the standard data transportation methods such as HTTPS, SCP or SFTP ensuring that technology would not be a barrier for clients (see 7.1.4)
- delivered data should contain information which allows the receiver to perform correlation procedures between various data sources; such information should at least contain:

timestamps of events (with time zone information), ASN, affected IP addresses or domain names and preferably type of incident/exploit/malware (see 7.1.1, 7.1.2, 7.1.3, 7.1.4)

- information should be delivered to clients as soon as it becomes available, taking into account time needed for data enrichment and filtering false positives; bulk delivery of data is acceptable but should not be done in periods longer than 24 hours (see 7.1.1)
- data providers should include detailed descriptions of methods used for acquiring information about incidents, thus putting context on classifications (see 7.1.1).

8.1.3 Data source quality enrichment

The enrichment process allows incident information to be extended with additional data, often providing new insight into observed events. The recommendations below aim at helping data providers to increase the quality of their service.

- data providers should try to decrease the number of false positive classifications by means of filtering, verification and correlation as long as it will not delay the data sharing process significantly; if delays are inevitable, parallel data feed can be provided with corrected or extended information on incidents (see 7.1.1)
- a service should keep data in some form (possibly aggregated) for historical reference and research purposes; offline analysis can further increase quality of data by finding trends and patterns in observed malicious behaviour (see 7.1.9)
- data providers should deliver a feedback mechanism allowing users to post updates to the data feed; such updates, when verified, should be immediately incorporated in the data feed or be delivered as a separate one (see 7.1.1)
- when maintaining blacklists, data providers should implement a data aging and removal process for the list entries and explain these procedures to the recipients (see 7.1.1)
- data providers are encouraged to assign confidence levels (validity indicators) to information streams (see 7.1.1)
- it would be useful if data feeds were also enriched with additional statistical information about incident types, including aggregated information about other constituencies (see 7.1.1).

8.2 Recommendations for data consumers

New CERTs usually seek ways of extending their knowledge about incident information related to their constituency. Application for access to external data sources can be a troublesome process if the organisation has not yet established trust relations in the community. Recommendations presented in this chapter try to provide guidelines which make the process easier and prepare applying organisation for requirements imposed by data providers. Suggestions are also made on automation, correlation and ways of improving quality of datasets. Furthermore, some recommendations on technologies are made for data consumers to develop their own analysis capabilities.

8.2.1 Acquiring access to datasets

Registration and proof of eligibility is the first step an applying organisation has to take to gain access to a data source. This section focuses on recommendations for data feed users to make the application process quick and easy.

- organisations should review the list of sources evaluated in this study and select the most appropriate for their situation; the TOP 5 list for prioritisation may be useful to start with in this regard (for examples see 6.2)
- organisations are encouraged to develop their own monitoring capabilities taking into account their specific situation, as suggested in 7.1.5
- organisations seeking new information sources should establish relationships with security communities such as FIRST, TF-CSIRT, APCERT; being accepted in such a community helps them gain trust and can quicken the process of verifying eligibility to access restricted data feeds (see 7.2.3)
- in some cases the data provider requires that the client installs a sensor in its network to gather data for the use of the service; young CERT organisations should consider the benefits coming from participation in such venture, especially when they do not deploy network monitoring solutions themselves (see 7.1.5)
- CERTs should be aware of potential legal issues concerning data sharing when applying for access to services which require setting up a sensor or sharing back data (see 7.2.1).

8.2.2 Integration of feeds with internal incident handling systems

Automation systems are one of the key elements allowing successful processing of incident information. The recommendations presented below aim at delivering best practices for developing such systems.

- internal system of data processing should be prepared to handle data in many different formats; there is no single commonly accepted standard for the exchange of security data, therefore internal systems need to reprocess data into one consistent form (see 7.1.2, 7.1.4)
- received data after reprocessing should be stored in a database which allows offline analysis and correlation to be performed and can be used as a base for incident handling and visualisation software (see 7.1.2, 7.1.7, 7.1.9)
- organisations which are able to deploy network monitoring systems should integrate data from them with data from external sources; the integration should be achieved through correlation and removing duplication of data feeds; verification of internal data with external feeds permits monitoring of the quality of data produced by monitoring systems as well as assessment of the quality of external data feeds, which can be very useful especially for newly established CERTs (see 7.1.2, 7.1.3, 7.1.5)

8.2.3 Verification of quality of data feeds

Incoming data feeds need constant verification to ensure correctness of identification of incidents. Recommendations in this section provide guidelines on how data receivers can verify data feeds and contribute their findings to community.

- data consumers are encouraged to develop methods and criteria for the evaluation of the quality of the data source (see 7.1.1, 7.1.2)
- incident information, if not verified by data producers, should be verified by client software before submitting information to database or incident handling software; verification should not delay information flow in a manner that could render the data unusable (see 7.1.1)
- aside from verification, data should be correlated with external services, enriched with additional data and filtered to reduce number of duplicated events; again delay imposed by the operation is one of the key aspects which should be taken into consideration (see 7.1.1, 7.1.2)
- if the original data producer implements some feedback mechanism, data consumers should make use of it in order to give the data provider information necessary to improve the quality of provided data in the future (see 7.1.1)

8.2.4 Deployment of rising technologies

The study has shown that a number of technologies that can be useful for CERTs remain underused. Having access to these technologies may allow a CERT to develop a service of its own for others, enhancing its position in the community and gaining leverage for exchanging information with others. In short, a data-consuming CERT may become a data provider. Data-consuming CERTs are thus encouraged to:

- deploy their own sensor networks (see 7.1.5)
- deploy client honeypot technologies (see 7.1.6)
- deploy sandbox technologies (see 7.1.6)
- implement passive DNS monitoring technology if possible, perhaps starting a CERT-wide passive DNS monitoring project (see 7.1.12)

8.3 Recommendations for further activities on the EU and national level

These recommendations are a summary of activities and areas of research that can be carried out on the EU and/or national level to eliminate many of the shortcomings discovered during this study. ENISA can play a role in all of the activities mentioned below.

- Facilitating wider usage of underused technologies for network monitoring mentioned in this report.
- Investigating how to improve reporting of data leaks to the affected organisations (see 7.1.13).
- Encouraging the adoption of common standards for the exchange of incident information (see 7.1.4).

- Integration of statistical incident data gathered on wider scale from national CERTs and organisations willing to share data, performing long-term analysis and correlation, producing reports, research materials, advisories and predictions (see 7.1.2, 7.1.9).
- ENISA could, as part of its tasks, advise the relevant EU and national bodies on how to reach a balance between privacy protection and security provision needs and clarifying how sensitive security data can be shared between data providers, consumers and intermediaries such as national CERTs (see 7.2.1).

9 Conclusions

Efficient detection and handling of security incidents is a demanding and responsible task. The threat landscape evolves constantly, forcing security teams to keep up with new threats. The broad spectrum of different incident types that can affect a constituency makes it difficult for just one CERT to detect and mitigate every single case. However, as the results of the survey carried out among European CERTs and some selected CERTs worldwide showed, the proactive detection of incidents is still used to a limited extent compared to reactive actions. Additionally, a large number of survey respondents expressed dissatisfaction with their current sources of information and are considering trying new ones. Fortunately – as shown in the study – the diversity of today’s security community has led to the creation of multiple data feeds that use different methods and technologies to monitor different types of attacks, most often affecting multiple constituencies. Information presented in this report can be used by CERTs interested in improving their operations to gain insight into available data sources – both external and ones that they can develop internally.

Moreover, the analysis of the survey results, research and discussions with experts allowed for the identification of a set of shortcomings of various kinds. Most of these are technical issues, where CERTs expressed concerns about data quality and reliability, many and incompatible formats of delivered data or underreported incidents of various types. Other identified shortcomings are legal and organisational in nature and mainly concern data sharing between CERTs in different countries and cooperation problems. Based on the identified shortcomings a set of recommendations for data providers, consumers and EU and national-level governmental bodies is proposed to mitigate the identified shortcomings. The end goal is improving data sharing and cooperation in proactive detection and incident handling between CERTs – an essential element for the successful mitigation of cyber-attacks.

10 Annex I: Abbreviations

API	Application Programming Interface
APT	Advanced Persistent Threat
AS	Autonomous System
ASN	Autonomous System Number
AV	Anti-Virus
C&C server	Command and Control server
CAPTCHA	Completely Automated Public Turing test to tell Computers and Humans Apart
CERT	Computer Emergency Response Team
CERT/CC	CERT Coordination Center
CIDR	Classless Inter-Domain Routing
CSIRT	Computer Security Incident Response Team
CSV	Comma-Separated Values
DDoS	Distributed Denial of Service
DNS	Domain Name System
DoS	Denial of Service
ENISA	European Network and Information Security Agency
EU	European Union
FIRST	Forum of Incident Response and Security Teams
FTP	File Transfer Protocol
GPG	GNU Privacy Guard
GPL	General Public License
GUI	Graphical User Interface
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
HTTPS	HyperText Transfer Protocol Secure
IDS	Intrusion Detection System
IODEF	the Incident Object Description Exchange Format
IP	Internet Protocol
IPS	Intrusion Prevention System
IRC	Internet Relay Chat
ISP	Internet Service Provider
LAN	Local Area Network
MS	Member State
MTA	Mail Transfer Agent
MUA	Mail User Agent
MX record	Mail eXchanger record in DNS
OSI model	Open Systems Interconnection model
PGP	Pretty Good Privacy
RSS	Really Simple Syndication
SCP	Secure Copy
SFTP	Secure File Transfer Protocol
SIEM	Security Information and Event Management
SIP	Session Initiation Protocol
SMB	Server Message Block
SMS	Short Message Service
SMTP	Simple Mail Transfer Protocol
SQL	Structured Query Language
SSH	Secure Shell

TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
TSV	Tab-Separated Values
TTL	Time To Live
UDP	User Datagram Protocol
URL	Uniform Resource Locator
UTC	Coordinated Universal Time
VoIP	Voice over IP
WAF	Web Application Firewall
WWW	World Wide WEB
XML	Extensible Markup Language
XSS	Cross-site scripting

11 Annex II: CERT survey analysis

This annex is published as a separate document on ENISA

website: <http://www.enisa.europa.eu/act/cert/support/proactive-detection/survey-analysis/>



P.O. Box 1309, 71001 Heraklion, Greece
www.enisa.europa.eu