Alerts, Warnings and Announcements

Best Practices Guide

November, 2013

www.enisa.europa.eu

About ENISA

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

Authors

Cosmin Ciobanu from ENISA, with support from Michael Potter and Don Stikvoort from S-CURE, The Netherlands, Mirosław Maj and Tomasz Chlebowski from ComCERT, Poland, Roeland Reijers from Rubicon Projects, The Netherlands and Mirko Wollenberg from DFN-CERT Services, Germany who produced this document as consultants.

Contact

For contacting the authors please use CERT-Relations@enisa.europa.eu.

For media enquiries about this document, please use press@enisa.europa.eu.

Acknowledgements

ENISA wants to thank all institutions and persons who contributed to this document. A special 'Thank You' goes to the following contributors:

1. NCSC-NL is thanked and acknowledged for their contribution to this good practice guide, in the form of providing information about their TARANIS system, and various discussions with NCSC-NL experts on the topic of this guide.
2. Countless other experts contributed to the content of this document. These contributions happened via email and in live talks, especially in the TF-CSIRT meeting in Bucharest and the FIRST conference in Bangkok in 2013. The contributions from CERT-EU, CERT Societe Generale, CERT.LV and DFN-CERT colleagues need special mention.
3. CERT Société Générale for providing information about their documents covering their Incident Response Methodologies.

Legal notice

Executive summary

This guide complements the existing set of ENISA guides[1] that support Computer Emergency Response Teams (CERTs, also known as CSIRTs). It describes good practices and provides practical information and guidelines for the process of preparing and issuing alerts, warnings and announcements to a CERT's constituency.

The main focus area of the guide is the process of informing the CERTs and their constituencies about threats and ways to contain threats – a core service carried out by most CERTs – which involves having an identified and reliable set of information sources and a well structured process of assessing and processing the incoming information, enabling the CERT to get the right information at the right places in the most timely fashion.

Other topics covered by the guide include incident response methodologies and recommendations on how to improve the process of alerts, warnings and announcements.

The primary target audiences of this guide are CERT technical staff and management. Secondary target audiences are IT security vendors, universities and CERT training institutions.

For a CERT in the set-up stage this guide will provide very valuable input on how to shape the process of alerts, warnings and announcements. For existing CERTs, it can serve as a means to re-design their current processes and further improve them. For established CERTs this document contains recommendations on how to improve this process together in cooperation (there is considerable potential here which is being under-used at the moment of writing this guide).

[1] http://www.enisa.europa.eu/activities/cert/support

Table of Contents

1 Introduction

The majority of CERTs warn their constituents about pending dangers and upcoming threats. The methodology by which this is done is often unchanged since the early 1990s. However, since then the landscape has changed. Threats and dangers have multiplied, the required reaction time is decreasing and the stakes are high. CERTs all over the world are working hard to meet challenging requirements to perform in a more professional manner in all procedures. That includes improving and streamlining the processes, systems and functions of the warning and alerting services. This guide aims at helping CERTs in the task of improving and streamlining those processes.

Goal

This good practice guide aims to:

(1) inform about the current best practices in preparing and issuing alerts, warnings and announcements for the CERT's constituency, and
(2) suggest ways in which alerting processes can be improved, not only inside an organisation but also in cooperation with others, especially the CERT community and IT security vendors.

Target audience

The target audience is primarily all those involved in information security incident prevention and response – in other words mainly the CERT community. ENISA's focus is on EU Member States and their CERTs, but this document can be applied for any CERT worldwide.

A secondary audience are IT security vendors, universities, training institutions, and in general all those who interact with the CERT community or train people in the subject matter of CERT work.

Structure of this document

This document provides information on all aspects related to the process of preparing and issuing alerts, warnings and announcements.

Chapter 1 Introduction

This chapter provides the background and introduction to this guide.

Chapter 2 Alerts, Warnings and Announcements: the Concepts

This chapter discusses the concepts of alerts, warnings and announcements, starting form the CERT services portfolio and going on to a basic process idea and existing challenges in this area.

Chapter 3 Alerts, Warnings, and Announcements: Best Practices

This chapter describes the best practices in alerts, warnings and announcements. It covers a detailed process, sources of information, risk assessment, dissemination, feedback, standardised formats and useful tools.

Chapter 4 Gap Analysis and Recommendations for the Alerting Process

This chapter offers an analysis of the existing gaps in the processes for alerts, warnings and announcements and follows up with recommendations how to approach these challenges. Most of these challenges need to be dealt with on the level of the community of CERTs and their stakeholders.

Chapter 5 Incident Response Best Practices

This chapter describes existing good practice in incident response and incident classifications.

Annex A: CERT Survey Results

This annex highlights results from the survey done within the CERT community prior the compilation of this guide.

Annex B: Public security news feeds sources

This annex provides an extensive list of public sources of security information, used by many CERTs in their alerting process.

Annex C: Relevant ENISA Documents Cross-reference

This annex places this guide inside the context of other relevant ENISA guides and other documents.

2 Alerts, Warnings and Announcements: the Concepts

This chapter outlines the concepts behind alerts, warnings and announcements. All notions defined are placed in the context of the traditional CERT services portfolio and from there we take a practical approach by looking at the common information sources used by CERTs, the severity rating of incoming and outgoing information and the dissemination channels to reach a team's constituency.

The definitions we give are compatible with common CERT service definitions but do contain some clear elements of improvement, especially in the inclusion of the concept of 'risk'.

2.1 Methodology behind the study

As a first step the concept of alerts, warnings and announcements as services was defined.

After this a survey among CERTs was prepared and conducted to collect their opinions about relevant aspects of this report, e.g. tools used by teams for their alerting, working and announcements activities.

Next, stocktaking of existing categories, and the most common types and channels of alerts, warnings and announcements, was conducted. Various sources were evaluated, as well as relevant standards and common data formats. This allowed recommendations to be drawn up based on the lessons learnt.

Next, incident response methodologies were evaluated and mapped to a classification by type of incident in order to help find and/or improve ways to mitigate attacks.

Finally, a report was prepared including findings, analysis and recommendations.

2.2 CERT services portfolio

The classification of CERT services as originally introduced around 1998 is still common today.[2]

CERT services can be grouped into three categories:

- Proactive services, which are aimed at improving the infrastructure and security processes of a CERT's constituents before any incident or event occurs or is detected. By providing proactive services, CERTs help to avoid incidents and minimise their impact and scope when they do occur.
- Reactive services, which are aimed at responding to requests for assistance from a CERT's constituency, reports of incidents, and tackling threats made or attacks against the CERT's systems.
- Security quality management services, which consist of services that improve an organisation's overall security. Any CERT should provide these services by leveraging its experiences of providing proactive and reactive services to its constituency and applying these experiences to quality management services.

A table that lists various services in these three categories is available on the website of the CERT Coordination Centre.[3] This is referred to as the (traditional) CERT services portfolio:

[2] See ENISA CERT baseline capabilities (updated 2012 version): https://www.enisa.europa.eu/activities/cert/support/files/updated-recommendations-2012 (p. 41); or the original from the CSIRT Handbook: http://www.cert.org/archive/pdf/csirt-handbook.pdf (p. 24).
[3] See http://www.enisa.europa.eu/activities/cert/support/guide2/annex/services or the source of that information: http://www.cert.org/csirts/services.html

Reactive Services	*Proactive Services*	*Security Quality Management Services*
Alerts and Warnings *Incident Handling* *– Incident analysis* *– Incident response on site* *– Incident response support* *– Incident response coordination* *Vulnerability Handling* *– Vulnerability analysis* *– Vulnerability response* *– Vulnerability response coordination* *Artifact Handling* *– Artifact analysis* *– Artifact response* *– Artifact response coordination*	***Announcements*** *Technology Watch* *Security Audits or Assessments* *Configuration and Maintenance of Security Tools Applications, and Infrastructures* *Development of Security Tools* *Intrusion Detection Services* *Security-Related Information Dissemination*	*Risk Analysis* *Business Continuity and Disaster Recovery Planning* *Security Consulting* *Awareness Building* *Education/Training* *Product Evaluation or Certification*

Table 1: CERT services

In **bold** in this table: alerts, warnings and announcements – the main topic of this best practice guide.

2.3 Alerts, Warnings and Announcements

'Alerts and Warnings' and 'Announcements' are detailed as follows in the references given in the previous paragraph:

> ***Alerts and Warnings (reactive)***[4]
>
> *This service involves disseminating information that describes an intruder attack, security vulnerability, intrusion alert, computer virus or hoax and providing any short-term recommended course of action for dealing with the resulting problem. The alert, warning or advisory is sent as a reaction to the current problem to notify constituents of the activity and to provide guidance for protecting their systems or recovering any systems that were affected. Information may be created by the CERT or may be redistributed from vendors, other CERTs or security experts or other parts of the constituency.*
>
> ***Announcements (proactive)***[5]
>
> *This includes, but is not limited to, intrusion alerts, vulnerability warnings and security advisories. Such announcements inform constituents about new developments with medium- to long-term impact, such as newly found vulnerabilities or intruder tools. Announcements enable constituents to protect their systems and networks against newly found problems before they can be exploited.*

This has become the traditional and dominant definition of these services, found in many sources including ENISA[6]. There are very few if any competing definitions. NIST seems to differ in their

[4] See http://www.cert.org/csirts/services.html#alerts
[5] See http://www.cert.org/csirts/services.html#announcements

equally authoritative 'Incident Handling Guide' but their definition is very brief[7] and does not add new insights or different categories: NIST refers to 'advisory distribution' as a proactive CERT service, basically being the same as 'security advisories' under 'Announcements' above; and to 'information sharing' which is in the reactive 'Alerts and Warnings' category.

2.4 Challenges in the traditional definition of Alerts, Warnings and Announcements

The above CERT/CC definition has some challenges that we will address in this guide, in order to provide improved and updated good practice:

1. 'Alerts and Warnings' are always referred to together.[8] They appear to be the same thing, but why then use two words? Looking at various CERTs' use of the terms, it is clear that different people have given different meanings to these words, but there is neither clarity nor uniformity in those approaches.
2. 'Alerts and Warnings' are marked as reactive services, 'Announcements' as proactive. Alerts are sent out 'as a reaction to the current problem to notify constituents' – but some of those constituents may not have the problem (yet), and thus the alert is proactive for them. And by contrast, an announcement can be about 'newly found vulnerabilities' – but these are rarely found in a lab but rather in real life, and thus were 'someone's problem' already. This would make the announcement also reactive. In the context of this guide we consider the split between proactive and reactive services as not helpful for clearly defining the concepts.
3. In the definitions of both types of services, the words 'advisories', 'alerts' and 'warnings' are mentioned synonymously, which does not add to clarity.

Section 3 will propose a solution to these challenges by replacing the time-dependency inherent in proactive versus reactive, by a risk-based approach. We will use the same terms and essentially the same ideas, with the addition of the important concept of 'risk'.

2.5 Basic process for Alerts, Warnings and Announcements

The following process concepts need to be taken into account when assessing alerts, warnings and announcements:

1. Information collection: what sources of information are used and assessed.
2. Risk assessment: information and its sources need to be assessed before it may be sent out as alert or otherwise. The perceived risk to constituents will be essential in this process.
3. Dissemination: important information will need to be disseminated to the right constituents, using an effective communication mechanism.
4. Feedback: what do the constituents do with the information they receive? How effective is it? What lessons can be learnt?

[6] See https://www.enisa.europa.eu/activities/cert/support/guide
[7] See http://csrc.nist.gov/publications/nistpubs/800-61rev2/SP800-61rev2.pdf (pp. 18–19)
[8] There is no clear reason why they are always grouped together. This seems to have started with the CSIRT Handbook http://www.cert.org/archive/pdf/csirt-handbook.pdf and then continued from there.

These concepts form a logical process, as outlined in the figure below. In section 3 good practices are provided for all four steps of this process for alerts, warnings and announcements.

Figure 1: Process of assessing alerts, warnings and announcements

2.6 Scope limitation

In this good practice guide we limit ourselves to **non-automated** generation of alerts, warnings and announcements.

This means that the automated handling of e.g. the output of automated sensors is **not** included here. Examples of such sensors are IDS sensors,[9] Passive DNS Sensors[10] and Honeypots[11] – many CERTs also subscribe to managed sensor services.[12] While is true that many of the concepts and ideas in this guide will also apply to automated source handling and alert generation, these are separate topics and are handled in separate good practice guides[13].

[9] See e.g. http://en.wikipedia.org/wiki/Intrusion_detection_system
[10] See e.g. https://security.isc.org/Passive_DNS_Sensor_FAQ/
[11] See e.g. https://www.shadowserver.org/wiki/pmwiki.php/Information/Honeypots
[12] As for example from Team Cymru or Shadowserver Foundation
[13] For more best practice guides see: https://www.enisa.europa.eu/activities/cert/support

3 Alerts, Warnings, and Announcements: Best Practices

3.1 Main types of Alerts, Warnings and Announcements

In a survey among European CERTs, the following types of alerts, warnings and announcements were identified – and they are presented here in order of decreasing popularity, with security advisories being issued by 64% of all respondents, and security bulletins by 29%. In the second column we give some examples and additional information:

(Security) Advisories	http://www.cert.org/advisories/ https://www.ncsc.nl/dienstverlening/response-op-dreigingen-en-incidenten/beveiligingsadviezen (Dutch) http://cert.europa.eu/cert/newsletter/en/latest_Security%20Bulletins_.html https://www.cert.fi/en/reports.html http://tools.cisco.com/security/center/publicationListing.x
Early Warnings	In general used for automated warnings, see e.g. http://dashboard.arakis.pl/en/index.html http://www.carmentis.org (in German) Automated warnings are out of scope for this guide.
Warnings	In practically all cases not published under the banner 'warnings' alone but in conjunction with 'alerts'
Alerts	http://www.csirt.gov.sk/img/infobrochure-eng.pdf http://govcert.bg/EN/Pages/SecurityAlerts.aspx http://ics-cert.us-cert.gov/alerts
Notifications	None found online. Apparently teams use this word more in the generic sense than as an alert type.
Announcements	http://www.us-cert.gov/announcements http://www.restena.lu/csirt/EN-CSIRTservices.html (mentioned under 'incident coordination') Also this term is mentioned in many rfc-2350 CERT descriptions online, like e.g. http://www.dfn-cert.de/en/rfc2350.html (under 5.1.2)
Heads-up	None found online. Apparently teams use this word more in the generic sense than as an alert type.
Newsletters	http://www.qcert.org/services/security-newsletter http://cert-mu.gov.mu/English/Pages/NewsLetterSubscription.aspx http://www.ssa.gov.za/Portals/0/SSA%20docs/CSIRT/2012/ECS-CSIRT_Newsletter_Issue_2_2012.pdf
Security Bulletins	http://www.auscert.org.au/render.html?cid=1 http://technet.microsoft.com/en-us/security/bulletin Apple uses 'security updates' as a variety: http://support.apple.com/kb/HT1222?viewlocale=en_US&locale=en_US

Table 2: Types of alerts, warnings and announcements

The conclusions from the survey (see Annex A) and online research with regard to the types of alerts are as follows:

- The most commonly used alert types, which can also be found online, are[14]:
  - **Advisories;**
  - **Warnings;**
  - **Alerts;**
  - **Announcements;**
  - **Newsletters;**
  - **Security Bulletins**
- All types of alerts are used both in a proactive and a reactive sense. Therefore the classification into a proactive or reactive service is not useful to identify the type of alert.
- Many CERTs publish information in their native language. This means that the typology used here does not literally apply, although the terms used are usually direct translations.

3.2 Process for Alerts, Warnings and Announcements

The basic process for alerts as described in chapter 1 can be more detailed and looks like the figure below for most CERTs. The four basic steps of this process are described in subsequent paragraphs.

[14] 'Early warnings' have been left out as these are used for automated warnings which are out of scope here.

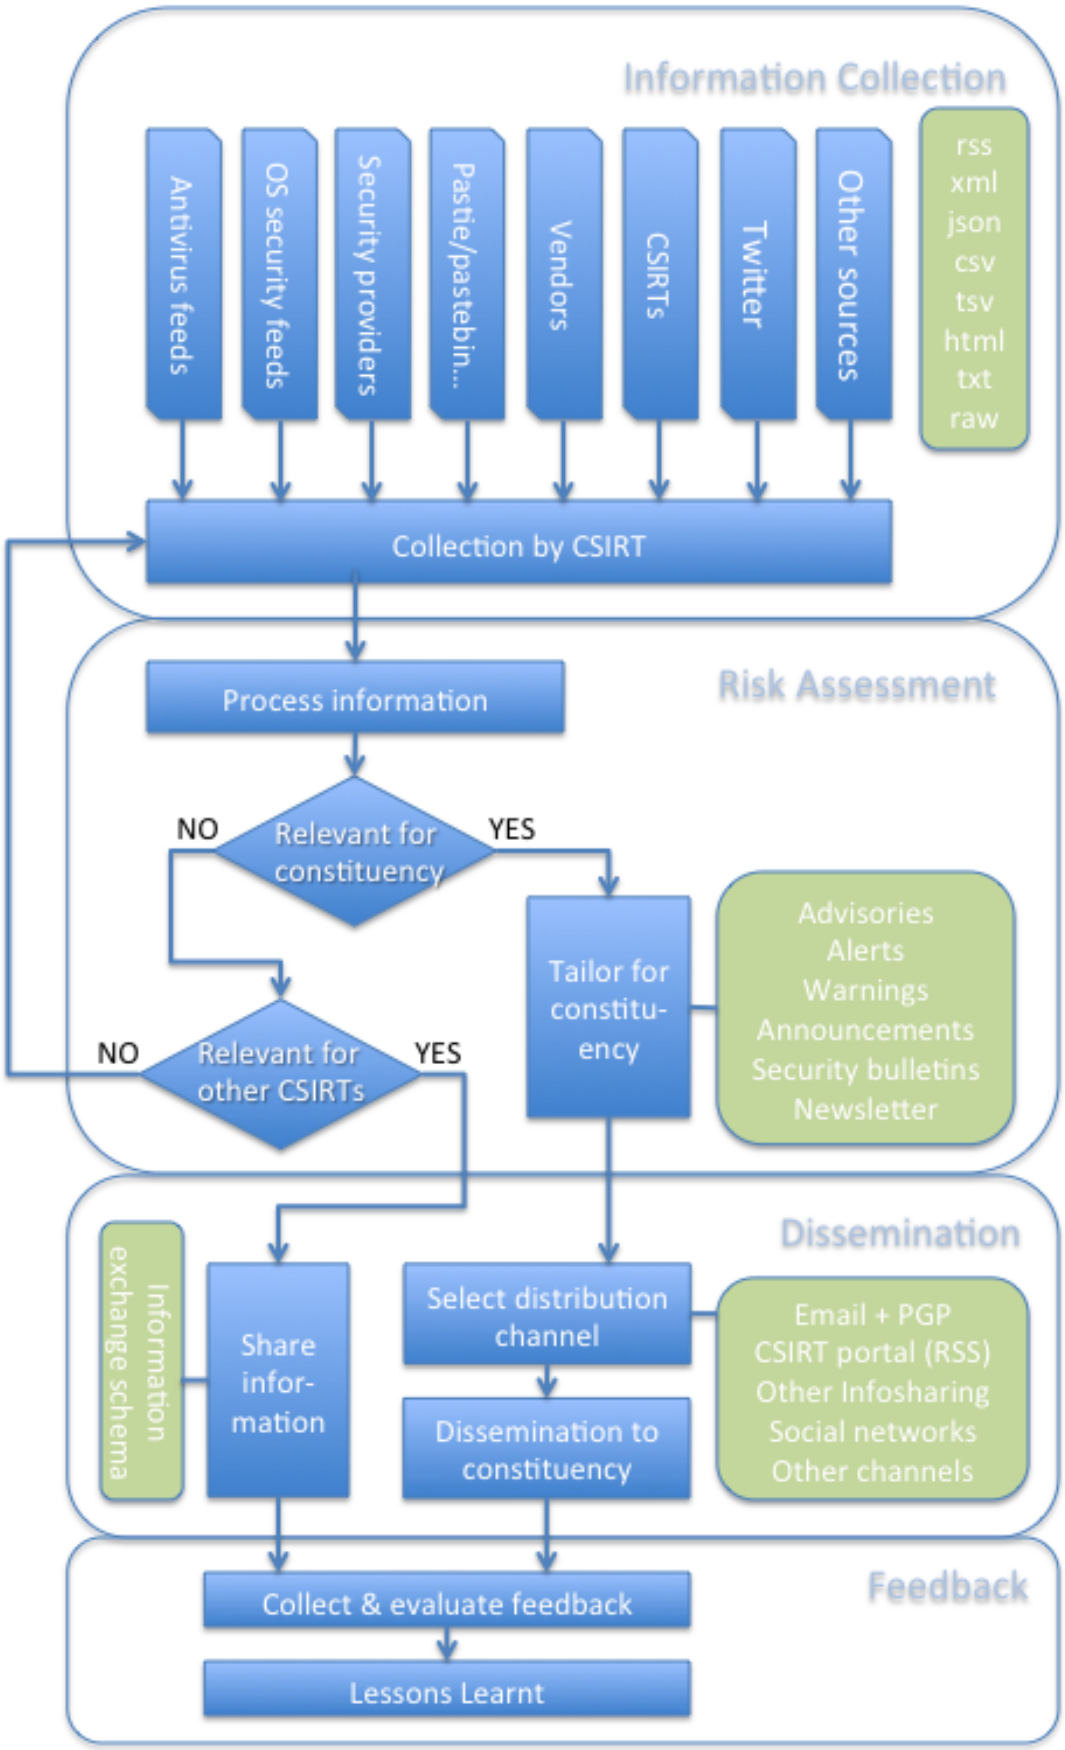

Figure 2: The four steps for issuing alerts

3.2.1 Information collection

Alerts, warnings and announcements all depend completely on the collection of useful information from good and reliable information sources. Each CERT needs to choose the type of information to collect and the sources of that information, and to assess the information gathered:

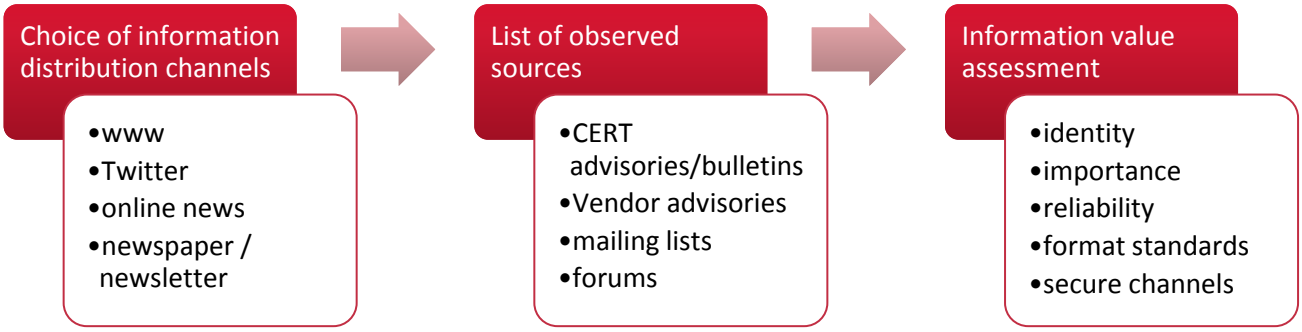

Figure 3: Information collection process

1. Choose the types of information channels the CERT wants to take into account: there are many electronic channels (including the Web, email, Twitter and other online news) but also paper ones (newspapers, magazines) and other media such as television.
2. List all the sources that need to be scanned from the information channels chosen. A partial list of sources includes:
   a. Other CERTs' advisories/bulletins;
   b. Vendor advisories;
   c. Mailing lists/forums (e.g. fulldisclosure, bugtraq);
   d. Security portals (e.g. thehackernews.com)
3. Assess the various information sources using the following steps for each source:
   a. verify identity;
   b. rate importance;
   c. rate reliability;
   d. rate default confidentiality;
   e. usage of communication/format standards;
   f. usage of secure channels if needed.

Step 3 is detailed below, followed by an example of a source list.

3.2.2 Verification of source identity

First, default sources need to be identified. This must be done properly and as this is not a time-critical service it can be done thoroughly. It is essential to know for sure that the content/ information provided by a particular source (vendor portal, IT security news portal, forum, mailing list) actually comes from that source: it has happened that alerts have turned out to be a hoax as the described hack or exploit never took place.

If a source is important, a relationship needs to be set up (meeting in person, PGP key exchange, etc.) and if possible, some form of MoU or subscription to their service should be made. A re-check must be conducted regularly, at least yearly or more frequently when necessary. These kind of relationships are especially important inside the CERT community or while communicating with or between experts.

For 'spontaneous' sources (not the set of default sources), use the following generic means to verify its identity:

- Ask peer CERTs (including the national or government CERT) and/or Law Enforcement nationally through existing national connections.
- Ask peer CERTs internationally (via FIRST and TI mailing lists, IRC channels of FIRST and TI).
- Use own information gathering (Google, whois etc.)
- If needed, call the source, using a generic phone number found on their website (don't use a number someone sends you in an email).

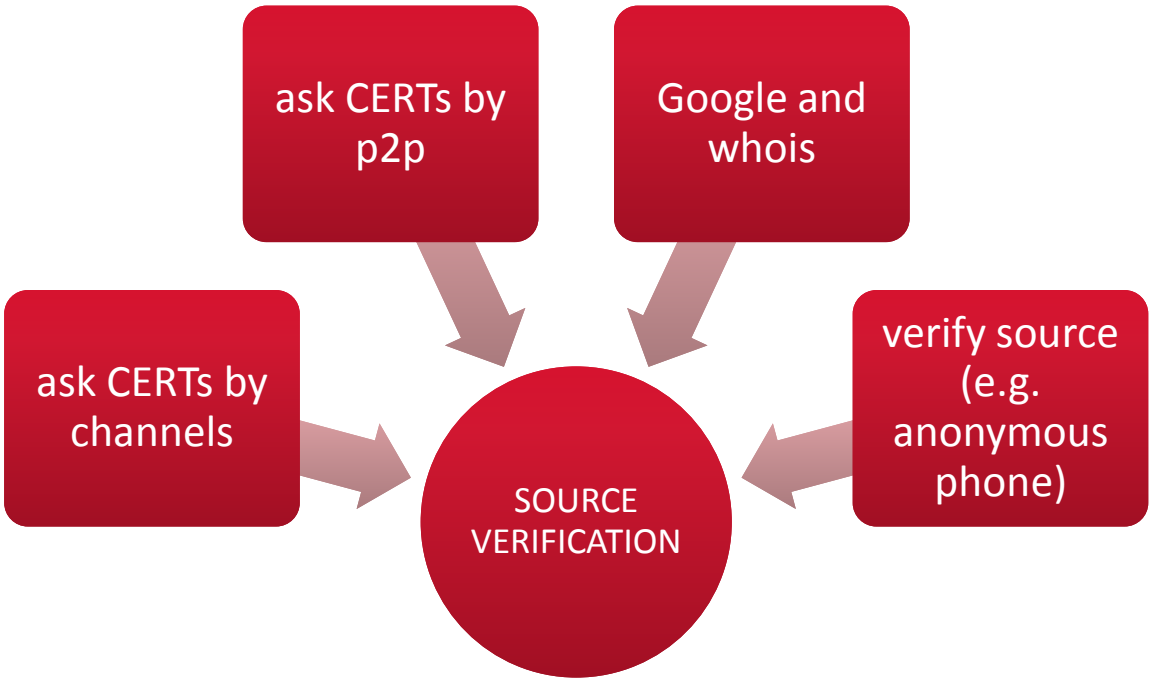

Figure 4: Activities for source verification

If the identity remains unclear, discard the information, unless it is of great importance. In that case, wait for another more trustworthy source or do your own research to find out if the information is true or false. You can do technical research – emulation etc. – and also get in touch with vendors, including Anti-Virus vendors and other security researchers.

3.2.3 Importance rating

The fact that CERTs usually have very limited resources requires them to a prioritise the incoming information (since there are not enough resources to monitor a large number of sources). Having a source importance rating makes it easier to focus only on important sources when under duress.

The importance of an information source for your CERT mostly depends on whether this source offers information that is useful for your constituency, and the main parameter here is the hardware and software which is used by the members of this constituencv. For instance, if your constituency

uses Microsoft products, Java-based or software from Adobe, then clearly the most important sources to monitor are Microsoft Security bulletins/ advisories Oracle Java [15]and Adobe security advisories[16].

If there is more than one source for specific products, another mechanism for rating of importance needs to be used. We propose the following additional method to rate the importance of information sources:

Source importance rating	Characteristics
Important	Most other CERTs receive and process information from this source
Fairly important	A fair number of CERTs are receiving and processing information from this source
Unimportant	Few or no CERTs are using this source

Table 3: Source importance rating example

3.2.4 Reliability Rating

The reliability of sources can be modelled as follows:

Source reliability rating	Characteristics
Reliable	The information from this source can be used and re-used without doubting its reliability.
Fairly reliable	This information source can usually be trusted, but a basic check needs to be done on each incoming report – such a check can either be done 'in the lab' by re-creating the situation described and comparing results, or by communicating with experts from other CERTs, security providers/vendors, or researchers.
Unreliable	The information from this source cannot generally be relied upon, so it needs to be thoroughly checked when of interest.

Table 4: Source reliability rating example

CERT work is not always routine work and the stakes can be high. So even when information from a reliable source comes in, you still need to do a quick plausibility check. If what you read doesn't make any sense, stop the process, discard the information or get in touch with the source first. The human factor always plays a role and thus errors and oversights do happen occasionally, even at the highest levels of professionalism and dedication. Another possibility is that a secure channel you use for this source may have been hacked. As a matter of principle in CERT work, it is good to always keep an eye out for the unexpected.

3.2.5 Default confidentiality rating

It is important to establish a default confidentiality rating for sources. Many CERTs use the 'Information Sharing Traffic Light Protocol' ISTLP[17] to rate confidentiality (see Table 5 for a summary in the context). A public news source like a website, open mailing list/ forum or Twitter is by default

[15] http://www.oracle.com/technetwork/topics/security/alerts-086861.html
[16] https://www.adobe.com/support/security/
[17] http://www.terena.org/activities/tf-csirt/publications/ISTLP-v1.1.pdf

WHITE in ISTLP terms and needs no further attention confidentiality-wise.[18] However a direct communication channel from IT vendor experts, peer CERTs, or from high up in a ministry, may be AMBER (or sometimes even RED) in ISTLP terms, and this is essential to know.

Default Confidentiality Rating (ISTLP codes)	Characteristics
WHITE	This information is **public**.
GREEN	This information may be **freely spread** within your constituency and to relevant others, **but not made public**.
AMBER	This information may only be spread on a **need-to-know** basis within a particular team.
RED	This information is **for your eyes only**.

Table 5: Default confidentiality rating – ISTLP codes

about the above table refers to **default** confidentiality ratings for a given information source. The confidentiality of a specific report can of course be different: in general the rating by a reliable source takes precedence; if a source gives out a report as AMBER, then it is not possible to distribute it as WHITE or GREEN unless the source allows it. A more detailed explanation of ISTLP will be given in chapter 3.5.3

3.2.6 Assess communication/format standards

Find out if your source uses any standard formats or communication methods, like the ones described in section 3.7 below. Make sure you can comply with them, or otherwise ask your source for alternative feeds, or easy tools/scripts to 'parse' the information they supply.

3.2.7 Secure channels

For critical sources,[19] a secure channel needs to be created. Protocols like https, imap and pop3 provide confidentiality and integrity, but you may need to rely on X.509 client & server certificates, PGP key exchange or other means to achieve sufficient end-to-end security. Remember that for secure communication you need to ensure the following three factors:

1. Availability: the channel works when you need it.
2. Integrity: the information passing the channel cannot be tampered with.
3. Confidentiality: the information passing the channel can only be read by the appropriate parties.

[18] It is, however, still useful to classify such sources for confidentiality – for instance not all mailing lists are WHITE, as some mailing lists have restricted membership. Also, not all web portals are WHITE – some can only be accessed via uid/pwd or via personalised X.509 client certificates.
[19] In general a source where importance, reliability and default confidentiality all rank high

3.2.8 Example of source list

The following serves as a proof-of-concept example for a source list including the three ratings. This list is subjective, incomplete and for illustration purposes only.

Information source	Type	Importance	Reliability	Default Confidentiality
https://www.adobe.com/support/security/	Application	Important	Reliable	WHITE
http://technet.microsoft.com/en-us/security/advisory	OS	Important	Reliable	WHITE
Dragon's Newsbytes	General Security News	Fairly Important	Fairly reliable	GREEN
Security bulletins from your n/g CERT	OS, Routing, Application	Important	Reliable	AMBER
Twitter feed from a renowned security researcher specialising in Linux	OS	Unimportant (but useful)	Unreliable	WHITE
Heads-up confidential information from your SCADA vendor	SCADA	Important	Fairly reliable (reliable but often still in test phase)	RED
Bugtraq mailing list	OS, Routing, Application	Fairly important	Fairly reliable	WHITE

Table 6: Some information sources and their ratings

Source lists can become quite extensive. NCSC-NL for instance scans over 1,000 sources, from closed mailing lists to public websites to Twitter feeds and other social media, using their TARANIS.[20] TARANIS is discussed below. A top-50 list of English public sources is provided in Annex B.

3.3 Monitoring emerging sources of information

Other sources of emerging information include Twitter[21], IRC[22], Pastebin[23], Internet forums, etc. These sources have become very popular for both sides of cyber conflicts – cyber criminals on one hand and security specialists including CERT staff on the other hand.

Very often this information is not based on IP addresses, which are a most important data source for CERTs to identify attacking or victim systems. More and more relevant information is context-specific, thus working with the constituency requires a better understanding of their technical environment as well as methods of attack against their systems. If, for example, a CERT provides services for a particular organisation that is an owner of a system 'ABC123', and the name of this system is specific and unique, then the CERT needs to start active monitoring for all information related to the system (like "We hacked ABC123" posts on pastebin). There are already many examples of successful usage of social media in tracking criminals:

[20] For more information on TARANIS: http:// https://www.ncsc.nl/english/services/incident-response/monitoring/taranis.html
To contact NCSC about TARANIS or other subjects, please use the contact form on their website:
https://www.ncsc.nl/english/organisation/contact/contactform.html
[21] See www.twitter.com
[22] http://en.wikipedia.org/wiki/Internet_Relay_Chat
[23] http://pastebin.com/

Figure 5: Hackers' forum screenshot presenting discussion about the types of data which hackers stole from Sony (New York Times online service: *http://bits.blogs.nytimes.com/2011/04/28/hackers-claim-to-have-playstation-users-card-data/?_r=0)*

- two men were identified as criminals who attacked Amazon, EBay and Priceline via DDos. They were carelessly bragging about this fact on an online hackers' forum and shared a lot of information about various attacks and stolen credit cards[24]
- hackers discussed break-in activity into Sony PlayStation Network and the credit card numbers they had stolen on an underground Internet forum[25].

3.3.1 Twitter channel

Currently many CERT teams maintain their own Twitter channels to reach out to their constituency.

Twitter is also a source of information from individuals or groups about hacking activity on the internet. These can be valuable sources of information for other CERTs as well. The main challenge is to set up an effective method to constantly monitor these sources.

Twitter channels become more and more important sources. Beside CERT Twitter accounts, hacking-related channels should be monitored. For determining these, keywords related to hacking activities need to be used. A good start is the following proposed list of keywords, which are of relevance at the time of writing of this document: 'anon', 'tango down', 'ops', 'corrupt', 'Cr3w', 'cyberwars'.

Another idea is to use geographical location names to determine information relevant to a specific constituency and/or country, like the 'AnonInPoland' user channel on IRC.

Another good attempt to find relevant information is following some specific tweets Prelated to periodic hacking activities, e.g.

Figure 6: Twitter channels of chosen CERT teams

[24] More: http://arstechnica.com/security/2012/07/hacking-duo-charged-for-amazon-ddos/
[25] http://thehackernews.com/2011/04/complete-irc-chat-of-playstation.html

'#Op[something]' type of operations. These names (e.g. #OpUSA, #OpBankseters, etc.) can be used to constantly monitor the most important facts of those activities on the Internet.

For more permanent and more effective Twitter monitoring the Twitter API is available on the Twitter website[26]; it can be used to manage tweets, as stipulated in the ENISA CERT training scenario 'Identifying and handling cyber-crime traces'.[27]

3.3.2 Monitoring emerging sources – keyword definitions rules

As mentioned, it is very important to develop and maintain a good list of keywords which will be used for monitoring and detection. In practice there are two sources of keywords:

- the set derived from parameters defining the constituency. These types of keywords are naturally very organisation-oriented and include for example names of systems or of particular persons. Ideally these keywords are proposed by representatives of the constituency.
- the set developed and maintained by the members of the CERT which pop up during regular monitoring activities as relevant.

To provide some simple examples:

- usage of the name of a particular organisation, for example 'ENISA' or 'ThisLargeCompany'
- translation of words into English from local language, like 'agency' (instead of 'agencja' in Polish)
- IP addresses or AS numbers
- domain names of the monitored organisation or part of the constituency, e.g. 'enisa.europa.eu' or 'europa.eu'
- words usually used when information about successful attack are issued, e.g.: 'tango down', 'p0wned', 'hacked'. If local language words are also often used in such situation, they should be added to the set.

3.3.3 Monitoring IRC channels

Another method of obtaining relevant information is monitoring of IRC channels. The problem in fulfilling this task is that it is very time consuming, which requires automation

In using automation, the most dangerous aspect is the possibility of 'false positives', of misidentifying a person or an organisation who are for example only present on the channel (undercover) to carry out monitoring to discover criminal activity.

The issue of working undercover is too complex to be tackled fully in this guide, but some basic ideas to cover the real identity of a person in an IRC channel are

- to use anonymisation of the network connectivity (e.g. with the TOR[28] service). The IRC channel can be reached anonymously by executing the 'torify' command which is a part of the 'tor' package (Ubuntu and Debian distributions). If we want for example to use irssi Linux client the following command should be executed: *torify irssi*,
- to periodically make a 'human interaction' on the channel in order to be recognised as a trusted party, and not as a monitoring bot
- to periodically share potentially valuable information (valuable from a criminal's perspective). This is tricky, as this information ('beacons') should not bring a real value and for example could be re-published from other public sources!

[26] https://dev.twitter.com/
[27] http://www.enisa.europa.eu/activities/cert/support/exercise
[28] https://www.torproject.org/

3.3.4 Internet forums and repositories

Other interesting sources of information are Internet forums and repositories. Services where anonymous users can post text-based information without experiencing problems subsequently are becoming more and more interesting for monitoring and investigations. The most widely used services have already been mentioned: pastebin.com. Originally created in 2002 for sharing of source code, it has also become a very popular repository for posting the result of hacking activities, such as the results of compromising popular services, of scanning activities, data leaked from attacks, etc.

It is possible to also automate monitoring of these services, by using functions built into the services or by using third party solutions like pastemon[29], which is basically a script which runs as a daemon on Unix-like systems and monitors pastebin.com for interesting postings, based on regular expressions. Found information can be sent to syslog.

3.4 Risk assessment

Once the information sources have been listed and rated, the information gathering can start. Each time a piece of information (referred to as 'report' below) comes in, it needs to be individually assessed to discover the relevance for the CERT's constituency, how important the report is and how urgent it is to act. This assessment is generally based on the following factors:

a. The report source ratings discussed in section 3.2.3.
b. The urgency of the report.
c. The initial risk assessment of the provider.
d. The severity in terms of direct potential impact.
e. The threat in terms of the loss of reputation, customers or money.
f. The type of constituency potentially impacted.

These factors can be taken into account by introducing the concept of risk assessment. Risk is generally defined as 'probability multiplied by impact'. This means that the risk involved with a certain event is the chance that that event will occur, multiplied by the impact of that event when it occurs. An example of a team which uses 'risk' as a deciding factor is NCSC-NL.[30]

Risk assessments are done in numerous ways, varying from down-to-earth and simple to advanced and complicated. For the sake of our argument we will first describe a straightforward example where both probability and impact are qualitatively assessed and given a value of 1, 2 or 3, with 1 being low and 3 high. This then leads to the following table, including a proposal for a simple severity rating:

[29] https://github.com/xme/pastemon
[30] https://www.ncsc.nl/

Probability (Chance)	Impact	Risk = Probability x Impact[31]	Severity
Low (1)	Low (1)	1 or 2 = Low	Low
Low (1)	Medium (2)		
Medium (2)	Low (1)		
Low (1)	High (3)	3 or 4 = Medium	Medium
Medium (2)	Medium (2)		
High (3)	Low (1)		
Medium (2)	High (3)	6 to 9 = High	High
High (3)	Medium (2)		
High (3)	High (3)		

Table 7: Severity rating example

The major strength of this approach is its simplicity! Much time can be wasted in highly detailed risk assessments and the results often do not justify the effort.

Both probability and impact do depend on the constituency – for instance, if a constituent does not use a specific piece of software, then the chance of a compromise of that software is clearly zero for that constituent. However, many CERTs, and especially national / governmental CERTs, teams for national research networks or for multinational companies often have a large variety of applications and operating systems in their constituency; many of them (33% of the teams in the survey we performed See Annex A) just send out all alerts to all constituents and leave the applicability to their constituency to decide. Of course it is necessary to make clear what products are concerned, so that such a decision can be fairly simple.

The next two sections give examples of commonly used, more elaborate risk assessment approaches. Section 3.4.3 then answers the question "How relevant is this, for whom, and how is the information tailored based on risk and target audience".

3.4.1 Common Vulnerability Scoring System (CVSS)

The intention of CVSS,[32] created by NIAC[33] in 2005 and now maintained by the FIRST community, is to create a global framework for disclosing information about security vulnerabilities. CVSS has since been widely adopted by vendors[34] who use it to rate their vulnerabilities.[35] CVSS calculators are available online.[36] The scores have a range of 0 (least severe) to 10 (critical).

The CERT community outside of vendors has been slower in adopting CVSS. In our survey only 1 in 10 teams used it, whereas over 80% mostly relied on human expertise for risk assessment. CERT-EU is an example of a team that uses CVSS scores in most of their advisories.[37]

[31] This definition of risk is common. See e.g. http://www.mitre.org/work/sepo/toolkits/risk/StandardProcess/definitions/occurence.html
[32] http://www.first.org/cvss
[33] http://www.dhs.gov/national-infrastructure-advisory-council
[34] http://www.first.org/cvss/eadopters
[35] see e.g. http://www.oracle.com/technetwork/topics/security/alerts-086861.html and click on any 'patch update'; or http://www.iss.net/threats/ThreatList.php
[36] see e.g. http://nvd.nist.gov/cvss.cfm?calculator
[37] http://cert.europa.eu/cert/newsletter/en/latest_Security%20Bulletins_.html

TARANIS is an information processing tool developed by NCSC-NL specifically to manage the process of creating useful alerts based on a multitude of information sources.[38] In our survey, 3 out of 10 teams used the TARANIS rating instead of CVSS. Now TARANIS is partly based on CVSS but the objection of the TARANIS developers towards CVSS is that the latter obscures the difference between probability and impact and puts the results in only one score. TARANIS, as you will see below, has chosen to keep chance and impact separately visible.

Even if you choose not to use CVSS in your team, you need to be aware of how it works and what the scores mean, as it is certainly a current practice for vulnerability rating, adopted by many vendors and security providers.

3.4.2 TARANIS risk assessment

The risk assessment used by NCSC-NL[39] in their TARANIS tool is also widely used, by teams in at least 20 countries as per September 2013. It serves as an excellent example how to approach this challenging area. The tool is available for CERTs on request from NCSL-NL

The TARANIS risk assessment matrix is a highly pragmatic combination of various approaches, taken from CVSS, US-CERT,[40] SANS Internet Storm Center[41] and Microsoft.[42]

They use two matrices for risk assessment, one for chance, and one for impact (which NCSC-NL refers to as 'damage' but we will stick with the term 'impact' here).

The **chance matrix** is as follows:

Question	Option 1		Option 2		Option 3	
Is the vulnerability present in the standard configuration/installation?	No	1	Unclear/yes	3	-	
Is exploit code available?	None	1	Proof of Concept (PoC)	4	Exploit	6
Are technical details available?	None	1	Somewhat	2	Full	3
Required access?	Physical	1	LAN/immediate vicinity	4	Internet	6
Required credentials?	Admin	1	User	2	None	4
How complex is it technically to exploit the vulnerability?	Complex	1	Average	2	Simple	3
Is user interaction needed?	Complex	1	Simple	3	None	4
Is the vulnerability being exploited in the field?	No	1	Limited	2	Large scale	3

[38] See https://www.ncsc.nl/english/services/incident-response/monitoring/taranis.html
[39] *https://www.ncsc.nl/binaries/nl/dienstverlening/response-op-dreigingen-en-incidenten/beveiligingsadviezen-toelichting/1/Inschalingsmatrix.pdf*: this is in Dutch and not available in English publicly, but we provide translated versions of the relevant parts here with the kind permission of NCSC-NL
[40] http://www.kb.cert.org/vuls/html/fieldhelp#metric
[41] http://www.sans.org/newsletters/risk/
[42] http://technet.microsoft.com/en-us/library/dd632949.aspx

Is the vulnerability expected to be exploited soon, or will an exploit come out?	No	1	Yes	3	-	
Availability of solution?	More than 2 months	1	Less than 2 months	2	None	3

Table 8: TARANIS risk assessment model – chance matrix

Points for all questions need to be added up. The 'chance rating' is set as follows:

- Low: 10–18 points
- Medium: 19–27 points
- High: 28–38 points

Second step is the **impact matrix** (referred to as 'damage' by TARANIS):

Question	Option 1		Option 2		Option 3	
Denial of service?	No	Low	Yes, client	Low	Yes, infrastructure service	High
Execute arbitrary code?	No	Low	Yes, user rights	Medium	Yes, admin/root rights	High
Remote rights (remote shell/root shell)?	No	Low	Yes, remote shell	Medium	Yes, remote root shell	High
Acquire local admin/root rights (privilege escalation)?	No	Low	Yes	Medium	-	
Information leakage?	No	Low	Yes, system information	Medium	Yes, data	High

Table 9: TARANIS risk assessment model – impact matrix

The impact follows directly from this table as the highest registered answer in terms of low, medium or high.

TARANIS characterises reports/vulnerabilities by a severity typology like low-medium, where the first term stands for chance and the second for impact, as stipulated in Table 10. This typology has the advantage that no 'averaging' occurs, and that anyone who reads disseminated reports immediately gets the idea of both the chance and the impact. NCSC-NL uses this probability-impact rating on their website with advisories[43] and in the advisories themselves,[44] by means of colour coding like the one used in Table 9.

In Table 10 we added a fourth column, which TARANIS does not use – but is compliant with Table 7 where both for probability and impact low equals to 1, medium to 2 and high to 3. This makes it easier to order the severity ratings in order of increasing risk (1 to 9). What we see then is reflected by some teams using TARANIS who choose to ignore issues with risk 1–2, that is low-low, low-medium and medium-low. Obviously issues with risk 9, severity high-high get top priority!

[43] https://www.ncsc.nl/dienstverlening/response-op-dreigingen-en-incidenten/beveiligingsadviezen: only in Dutch
[44] For instance this high-high advisory (in Dutch): https://www.ncsc.nl/dienstverlening/response-op-dreigingen-en-incidenten/beveiligingsadviezen/NCSC-2013-0285+1.01+Kwetsbaarheden+in+McAfee+ePolicy+Orchestrator+verholpen.html

Probability	Impact	Severity	Risk = Probability x Impact
Low	Low	low-low	1
Low	Medium	low-medium	2
Medium	Low	medium-low	2
Low	High	low-high	3
High	Low	high-low	3
Medium	Medium	medium-medium	4
Medium	High	medium-high	6
High	Medium	high-medium	6
High	High	high-high	9

Table 10: TARANIS risk assessment model – the final calculation

3.4.3 Relevancy and tailoring of the information

The first question associated with any risk assessment is: is this relevant for the CERT's constituency? Or, in other words, "Does this risk pose a threat to the constituency's assets right away or in the near future? How likely is the threat to have negative effects? And what is the impact of those effects?"

(The second question is: could this threat be relevant for other CERTs? If yes, then dissemination to those colleagues should be done immediately!)

If the risk is relevant for the CERT's constituency, then a report needs to be tailored corresponding with the risk.

As the survey shows most teams consider only two types of report. One is the alert, advisory or bulletin, which usually is the kind of information that needs immediate attention. The other is the non-urgent type of report, best characterised as 'newsletter'. It is perfectly in order for a new CERT to simply apply this approach, and use any of the report types identified as common in section 3.1.

However, we also propose a risk-based definition of various types of 'dissemination methods'. We combined the CERT/CC definitions of alerts, warnings and announcements with the risk assessment ideas presented above. The easiest way to explain this risk-based definition is by means of the following example, based on the various approaches presented above:

Simple Severity (2.4)	TARANIS Severity (2.4.2)	Report type
Low	low-low low-medium medium-low	IGNORE or NEWSLETTER
Medium	low-high high-low medium-medium	ADVISORY / BULLETIN
High	medium-high high-medium high-high	ALERT

Table 11: Recommendations based on the TARANIS assessment risk model

The 'metrics' used in this table present a choice – and the specific choice is up to you and your team.[45] However, the essential idea is to tailor the chosen report type or format according to the risk. High-risk urgent issues get an Alert – the name itself already suggesting urgency. (Security) Advisory or Bulletin is the most commonly used name and has been retained here for medium risk, medium urgency reports. Low-risk items are either ignored or distributed as newsletter.

3.5 Dissemination

The final step in creating alerts, warnings and announcements is to send them out to the stakeholders. This process is called dissemination.

Potential stakeholders are:

- Constituents/customers (external and/or internal to your host organisation).
- Other CERTs, usually subject to some agreed upon information exchange schema.
- In special cases: law enforcement / police.
- The world (many CERTs choose to make their reports available to the public).

For all stakeholders, effective dissemination channels need to be identified. This can vary from email and web-publishing, to Twitter or other social networks, RSS feeds, but also radio, television or newspapers.

Each dissemination channel has its own demands, and it is outside the scope of this guide to go into details here. In the next section we will, however, outline what the most common formats of dissemination (security advisory, bulletin or alert) sent out via email or made available on the web, should contain.

3.5.1 The dissemination channels

There are number of dissemination channels which can be actively used by CERTs. The most important are:

- CERT web portals – a traditional channel for security-oriented information which has effectively become the place where the most important and reliable information is issued.

[45] E.g. do you consider low-medium and medium-low to be ignorable or suitable for a newsletter – or does that already require an advisory? Is an alert only necessary for high-high, or also for medium-high and high-medium?

Very often references to postings on the portal are disseminated by other channels afterwards

- CERT mailing lists – probably the earliest method of disseminating information by CERTs. Nowadays mailing lists are still used especially by groups of teams which communicate with each other.
- CERT Twitter account – an increasingly popular method of disseminating CERT information. Twitter is especially effective if a team also wants to be reached by other parties like journalists, as twitter allows two way communication
- CERT Facebook pages which can be quite effective in reaching constituencies that need less technical information, so it is effective for ongoing basic awareness programmes.

3.5.2 Basic information inside advisories/bulletins/alerts

A written security advisory (or bulletin/alert) should contain the following type of information.[46] The shaded fields can be considered optional.

Title	Title of advisory
ID	Unique ID for advisory
Version & Date	Version and date of the advisory
ISTLP code	WHITE, GREEN, AMBER or RED (see section 3.5.3)
Risk	Risk e.g. in terms of simple rating (low, medium, high) or TARANIS rating (low-medium, high-high, etc.) or CVSS score
CVE-ID	CVE-ID(s)[47] 'Common Vulnerabilities and Exposure' tags associated with the vulnerability
Application	Vulnerable vendor application(s) / OS / app / etc.
Version	Version of the application(s)
CPE-ID	CPE-ID[48] 'Common Platform Enumeration' tag for application/version
Platform	Operating system(s) and version(s) where the vulnerability occurs
Update	Additional information regarding software updates or increased threat – this field would be empty in the first version of the advisory
Summary	Summary of the advisory
Consequences	Short description of the potential available to an attacker who exploits the vulnerability
Description	Detailed information on the vulnerability and how it can be exploited
Solutions	Software updates, patches, workarounds
Links	Links to more information

Table 12: Types of information in security advisory

[46] Inspired by https://www.ncsc.nl/dienstverlening/response-op-dreigingen-en-incidenten/beveiligingsadviezen-toelichting.html (in Dutch)
[47] http://cve.mitre.org/
[48] http://nvd.nist.gov/cpe.cfm

3.5.3 Information Sharing Traffic Light protocol (ISTLP) codes

The ISTLP 'Information Sharing Traffic Light Protocol' is widely used by CERTs for classifying their information,[49] and is supported and used by at least 100 CERTs in Europe.[50] In almost identical wording the ISTLP is also used increasingly outside Europe, especially by national/ governmental CERTs.

(Note: ISTLP was introduced in section 3.2.5 – but it was applied to the topic of the confidentiality of information sources. Fuller details are given in Table 13.)

ISTLP CODE	Characteristics
WHITE	Information that is for public, unrestricted dissemination, publication, web-posting or broadcast. Any member of the Information Exchange may publish the information, subject to copyright.
GREEN	Information can be shared with other organisations, Information Exchanges or individuals in the network security, information assurance or CNI community at large, but not published or posted on the web.
AMBER	Limited Disclosure and restricted to members of the Information Exchange; those within their organisations and/or constituencies (whether direct employees, consultants, contractors or outsource-staff working in the organisation) who have a NEED TO KNOW in order to take action.
RED	Non-disclosable Information and restricted to representatives participating in the Information Exchange themselves only. Representatives must not disseminate the information outside of the Exchange. RED information may be discussed during an Exchange, where all representatives participating have signed up to these rules. Guests & others such as visiting speakers who are not full members of the Exchange will be required to leave before such information is discussed.
NOTE that an 'Information Exchange' can be either in person, like a general FIRST or TF-CERT meeting of CERTs, or a meeting of a few teams together, but also an exchange in email or over the phone or fax. The rules below apply to all of those. It is not an absolute recipe, but needs to be applied thoughtfully – the ISTLP serves the purpose of bring more clarity with regard to the rules of information sharing, and is not a goal in itself.	

Table 13: ISTLP codes

3.6 Feedback

Whether information has been shared with a fellow CERT or disseminated to the team's constituency, it is important to ask for feedback from the recipient when possible!

For smaller teams, the best time to evaluate feedback is usually during regular team meetings. For bigger teams, evaluations may need to become more formalised and have their own process and manager.

However the evaluation works, make sure to draw 'lessons learnt' from it and to implement any recommendations based on that, right away or by adding them to the workplan for the next year.

This process is not only important for the CERT, but also to 'teach' the constituency how to react to the alerts and advisories they get from the CERT. If you disseminate advisories on software-patches,

[49] http://www.terena.org/activities/tf-csirt/publications/ISTLP-v1.1.pdf
[50] Supporting ISTLP is one of the accreditation demands of the TF-CSIRT Trusted Introducer: see https://www.trusted-introducer.org/processes/accreditation.html

you need to explain to your constituency how those patches should be applied. If you disseminate advisories on vulnerabilities, address how the constituency should plan their vulnerability mitigation process by providing workarounds or similar.

The following examples of useful feedback criteria is not exhaustive:

- Reaction time to the information.
- Speed of dissemination.
- Relevancy to the constituent.
- Content of the reports (clarity, wording, structure, risk assessment).
- Relevancy of the risk assessment.

3.7 Data formats and standards used in information collection & exchange

For this survey a wide range of information exchange formats for different stages of the process was proposed: receiving alerts, sharing & disseminating alerts, etc. However, most of the formats found were used very seldom or not at all.

Information exchange schemas and standards:

- EISPP/CMSI[51] not used;
- CAIF[52] not used;
- VulDEF[53] used by 1 team[54];
- Opensec ANML not used;
- OASIS AVDL[55] used by 1 team;
- VEDEF[56] not used;
- IODEF[57] **used by 1/4 of respondents;**
- IDMEF[58] not used;
- FIDEF not used;
- SFDEF not used.

Data representation formats (the information exchange schemas can be put into these representation formats, e.g. IODEF can be expressed either in XL or in JSON):

- JSON[59] used by 1 team;
- XML **used by most respondents;**
- CSV **used by most respondents.**

The survey shows that the most used exchange format by teams is IODEF. This is also due to the fact that CERT incident workflow tools like RTIR[60] adopt IODEF. However, it is not easy to use IODEF in a

[51] http://www.cert-ist.com/eispp/ and http://www.cert-verbund.de/projects/cmsi.html
[52] http://www.caif.info
[53] http://jvnrss.ise.chuo-u.ac.jp/jtg/vuldef/index.en.html
[54] The standard will be replaced with the new one – SECDEF; for additional information http://www.secdef.org/ site should be monitored.
[55] https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=avdl
[56] http://www.terena.org/activities/tf-csirt/vedef.html
[57] http://www.cert.org/ietf/inch/inch.html and http://www.ietf.org/rfc/rfc5070.txt
[58] http://www.ietf.org/rfc/rfc4765.txt
[59] http://www.json.org
[60] http://bestpractical.com/rtir/

generic way – the standard is so rich and complex that the only way to apply it is by defining specific use profiles. If for example RTIR uses profile A and another tool profile B, then it can safely be assumed that they will not be able to interoperate. Thus IODEF is a complex standard, most easy to use inside one platform, like RTIR. To use IODEF cross-platform requires defining the profiles in advance – or making translation filters per platform, with the risk of losing some information.

3.8 Supporting tools

There is a variety of tools for CERTs today which are used for the automated collection of data, like for example the output of sensor networks. Commonly known are Abusehelper,[61] Megatron, CarmentiS[62] and n6.[63] Automated collection is out of scope for this guide, however it is possible that for example Abusehelper will be further developed to work with non-automated source data (there are a few community projects ongoing in that respect). Therefore these tools are mentioned here.

The only tool which was fully designed for the purpose served by this good practice guide is TARANIS[64] from NCSC-NL.

TARANIS is used by an increasing number of teams in Europe and worldwide. It is well documented and maintained. NCSC-NL for example uses it to monitor more than 1,000 information sources, and it is the process tool used for the risk assessments and disseminations that follows the data collection. (TARANIS sources and the risk assessment process has already been described before.)

If a CERT needs to manage a considerable amount of information sources to use in their alerting process, they are strongly advised to see if TARANIS could help achieve their goals.

[61] http://abusehelper.be and https://www.clarifiednetworks.com/AbuseHelper
[62] http://www.dfn-cert.de/leistungen/forschung/carmentis.html
[63] http://www.cert.pl/projekty/langswitch_lang/en
[64] https://www.ncsc.nl/english/services/incident-response/monitoring/taranis.html

4 Gap Analysis and Recommendations for the Alerting Process

4.1 Gap analysis

4.1.1 Inefficient use of human resources in alerting process

The main weakness of the alerting process described in section 3 is the huge amount of information sources. As was mentioned before NCSC-NL, in comparison a rather big team (around 30 people), checks more than 1,000 sources regularly. Other, smaller teams will be able to process only a few hundred sources, or even less. Teams that consist only of 2–3 full-time equivalent staff members (the absolute minimum number of staff ENISA recommends for national / governmental CERTs[65]) it is practically impossibleto process 100+ sources. In addition to that most teams process the same information sources for rather similar constituencies – a huge overlap and inefficiency where definitely synergies could be found!

4.1.2 Standards are underused

Another gap we discovered during the work on this guide is the lack in utilisation of standards. ISTLP is widely known; generic data representation formats like xml (extensible markup language) and csv (comma separated values) are of course also widely used. CVSS, TARANIS, CVE and CPE are fairly popular and well known; IODEF is for instance used by those who use RTIR as incident management tool, as discussed above (see section 3.7). But none of these and other exisitng standards are being used to their maximum capability. In many cases rather than formalising information into categories free text and 'gut feeling assessments' are used. Still, some degree of format/description standardisation could really support interoperability and create a better platform for automated processes. With the increasing workload that CERTs face, this becomes increasingly important.

4.1.3 Lack of automation

Many CERTs still manage their core processes manually. Automation of processes has only just started to become interesting. Automated handling/ scanning of information sources, log files, sensor output, etc., will lead to significant time savings for CERTs, which are all handling increasing incident and threat volumes. However manual processing will always be necessary in many cases, because new and modified threats keep turning up, and automated processes are usually not good in spotting anomalies. However, using automation as much as possible frees up time for experts! They will have more time available to look for anomalies and more time to 'connect the dots'.

4.1.4 Lack of (uniform) education

Many people actually working in CERTs and most of those who join a team have almost never been educated for the work they are going to do. They are usually excellent system/ network engineers, security researchers, or IT students, etc. But there are few courses at universities or higher educational schools that prepare for CERT work, nor are there widely available courses specifically aimed at this community. If such courses were to exist, they would no doubt also deal with the aspects of information gathering and alerting including automated processes, tools, etc. It would be a great help if CERT members would come equipped with such knowledge. Initiatives in the area of CERT education are from ENISA which can deliver trainings to European national CERTs and from TERENA with TRANSITS training modules.

[65] According to ENISAs Baseline capabilities for n/g CERTs, refer to the operational capabilities https://www.enisa.europa.eu/activities/cert/support/baseline-capabilities

4.2 Recommendations

4.2.1 CERTs should have a shared alerting process

The single most important recommendation is to avoid overlaps and stop doing the same things a hundred times and more over in Europe, and create a shared alerting process instead.[66] This would ideally mean that a process like the one NCSC-NL maintains for more than 1,000 sources, can be distributed over a group of teams. Each team watches a subset of sources and does assessments on those – very important sources can even be covered by a few teams, to create second opinions. The results become available for the whole group of teams – and individual teams decide, based on that, if and to whom they will disseminate. Feedback and lessons learnt are collected per team, but also shared in the group. The following figure visualises this process.

Figure 7: Shared alerting process workflow

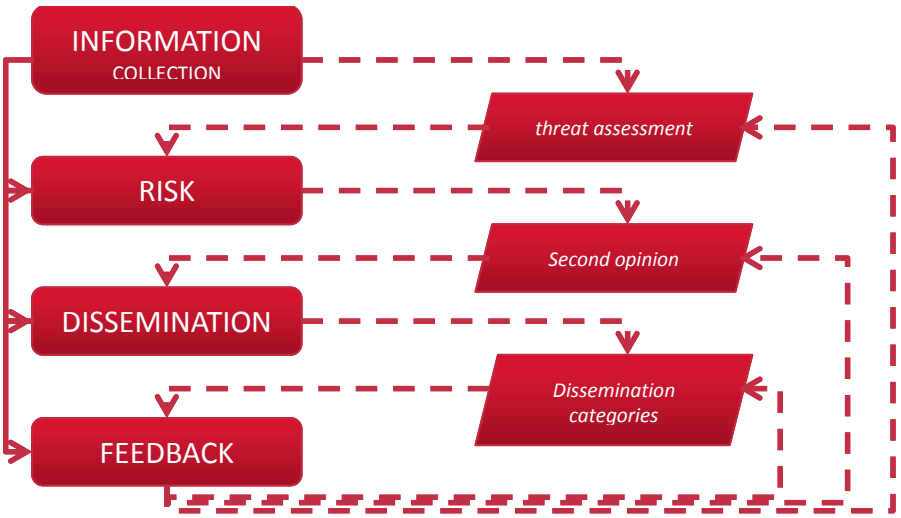

The idea of a shared alerting process is quite simple, but the implementation is not. The most important caveats are the following:

1. The group of teams need to cooperate coherently and constructively. Possible examples inside Europe are:
   1.1. TF-CSIRT accredited or certified teams.[67] The TF-CSIRT teams have a track record in achieving projects like this. An example was the eCSIRT.net project which ran around 2004. It used TF-CSIRT accreditation as starting point for participants, but added a Code-of-Conduct which is still of interest today[68] for inter-team projects. Since those days, TF-CSIRT has added an optional certification for their members, with much higher demands than for accreditation. Certification may therefore be an even more suitable starting point for cooperation agreements of the type suggested here.
   1.2. The European Government CERTs (EGC) group[69]
   1.3. ENISA national/governmental mailing lists and expert groups
2. Language. Many teams disseminate information in their native language – sometimes this is even a legal requirement. If however the information collection and risk assessment process – as well as the feedback process – could be done in English, then that would mean that

[66] We have not found proof that this has yet been undertaken.
[67] https://www.trusted-introducer.org/processes/overview.html
[68] http://www.ecsirt.net/service/coc.html
[69] http://www.egc-group.org

most of the 'engine' of this shared alerting process would be based on English as standard language. The decision whether to disseminate to constituents could be taken by a specific team based on the information available in English – and their own database of constituents' needs and equipment – and only after the decision was taken would translation would need to occur. We strongly recommend the use of English as common language whenever possible – this saves a lot of time and effort.

3. Tools and standards. ONE tool with standardised input- and output formats (like for example TARANIS) would need to be used for this process, and would probably need to be adapted in some ways to support such a shared alerting process. And accompanying standards like CVE and CPE need to be known and used in the same way by all participating teams.
4. Training. In order to do e.g. risk assessments in the same way in all participating teams, a shared system of relevant training would need to be put in place.
5. Restricted sources. Some local sources may be strictly limited to local or national use. These would need to be treated separately, or adequately shielded from shared sources, in order to be able to guarantee sufficient exclusivity.
6. Sources only available in one language. Some sources are only available in a language other than English. These sources should be monitored by a CERT native in this language, and other CERTs should get input from that team.

4.2.2 InfoSec community should promote use of relevant standards

Promotion of relevant standards should be preceded by the process of determining the most useful standards. As it is not easy to agree on what standards are the most useful, some simple methodology could be implemented to do this.

Methodology phases could include:

- This kind of process could start with the analysis of all tools used in terms of the standards implemented in them.
- Then the evaluation of the standards should be done together with the standards authors, tool developers and tool users.
- Thanks to such an approach, a list of the standards used, together with some simple evaluations can be created for all tools separately.
- Finally the list of the most valuable standards (understood as these are widely used and highly ranked) can be created. The possible criteria for such evaluation could be standard interoperability and readiness for automation process.

Figure 8: Promotion of standards process

4.2.3 CERTs should make increase use of automated processes

Regarding the automation of CERT processes, especially those relevant to alerting, warning and announcing activities, there is a wide gap between what particular teams are doing and what others know about it. A number of CERTs have developed their own tools and they use them actively and successfully in their environments working for their constituencies. From time to time some of these tools are presented to other teams. It is definitely a very useful and valuable activity but if we want to reach the next level of successful automation of CERTs work we probably need some regular and developed programme to be implemented.

The simple idea is to promote the exchange of information and best practices on a regular basis. For example the topic 'automation of CERT work' should become a regular part of CERT meetings, training, workshops and conferences. It is relatively easy to reach at least some specific groups like FIRST[70], TF-CSIRT[71] or AP-CERT[72]. This topic should appear in all 'calls for papers/ presenters'. Also it could become a part of team presentations – the formal (e.g. in the TI repository) and less formal (e.g. during CERT meetings).

Collecting many experiences from the CERT's work in automation could result in a new framework tool for CERTs, which will include all experiences and particular tools. If such solutions become interesting for teams, they can be developed to include the idea of stable and constant updating.

4.2.4 Improve CERT education and make it mainstream

As there is still a problem in implementing general ICT security aspects in education programmes, the idea of implementing specific CERT education seems even more difficult. But this does not mean it is not possible. Many aspects of ICT security topics related to incident handling are attractive and could get the interest of students. Just as the CERT concept is very often a good solution for implementing all security aspects in organisation, specific CERT education could become an important first step for students into the world of ICT security.

Thus it is recommended that CERT officers, specialists and managers promote the concept of CERT education whenever they are involved in education work at universities or other schools.

ENISA CERT training material[73] further improves the education in this area by providing more than 25 scenarios ready to be used by trainers and teachers. The scenarios cover a wide-range of topics like legal , technical and many more.

In practice the topic could be made relevant to students on almost all education levels wherever general IT aspects are taught – e.g. if the training is about code development topics such as vulnerability handling and vulnerability disclosure could be mentioned.

It is evident that deep analysis of ICT systems is a very interesting topic for a lot of young people who want to develop their technical skills. Unfortunately their skill development very often goes in the direction of black-hat activities. One reason for this is that there are no attractive alternatives for them to make their advanced skills practical. Their involvement in CERT-like work (e.g. computer forensics or a network investigation) could be just such an alternative for them.

[70] http://www.first.org/
[71] http://www.terena.org/activities/tf-csirt/
[72] http://www.apcert.org/
[73] https://www.enisa.europa.eu/activities/cert/support/exercise

5 Incident Response Best Practices

We have collected some incident response best practices in this section, plus accompanying measures like incident classification. As all CERT services are closely interconnected with each other (especially true for Incident Response and Alerting&Warning) this chapter aims at giving the reader more background to understand this connection and to properly react to incoming reports.

5.1 Incident response methodologies (CERT SG)

The most impressive collection of incident response practices that we uncovered are the IRMs[74] (Incident Response Methodologies) of CERT Société Générale[75] – inspired by SANS.[76] At the time of writing this report, the following IRMs were available, each in English, Russian and Spanish:

IRM-1: Worm Infection
IRM-2: Windows Intrusion
IRM-3: Unix Intrusion
IRM-4: DDoS
IRM-5: Malicious Network Behaviour
IRM-6: Website Defacement
IRM-7: Windows Malware Detection
IRM-8: Blackmail
IRM-9: Smartphone Malware
IRM-10: Social Engineering
IRM-11: Information Leakage
IRM-12: Insider Abuse
IRM-13: Phishing
IRM-14: Scam
IRM-15: Trademark Infringement

The example of IRM-13 on phishing is presented below. All IRMs are constructed in the same way.

5.1.1 Example of the incident response methodology – phishing

The methodology consists of six phases:

Preparation

In this phase the team prepares all relevant information regarding the mitigation of the phishing attack. Some important action during this phase are: establishing good relationship with relevant stakeholders, defining procedure, gathering information which will be necessary during the phishing attack, raising customer awareness.

Identification

Identification is the first operational phase in the phishing case response. The most important thing is to be able to detect the incident, determine its scope and start cooperation with all involved parties as well as those who can assist in resolving a problem .

[74] http://cert.societegenerale.com/en/publications.html
[75] http://cert.societegenerale.com/en/index.html
[76] http://www.sans.org

Containment

After proper identification, it is time for incident mitigation. During this phase there are number of actions which should minimise the effects of the attack.

Remediation

Stopping the attack is the next important action you should undertake. You contact the external parties (e.g. hosting company where the phishing site is hosted or/and local CERT) and ask for assistance in taking down the malicious site.

Recovery

During this phase you should try to return to the previous functional state. Analyse once more what has happened and treat this analysis as a checklist for your actions.

Aftermath

This is your 'lesson learnt' session. It is especially time to consider changes in your technical and organisational environment for better incident handling. Also it is time to collaborate with the legal department to decide whether the legal action is needed and what information you need for it.

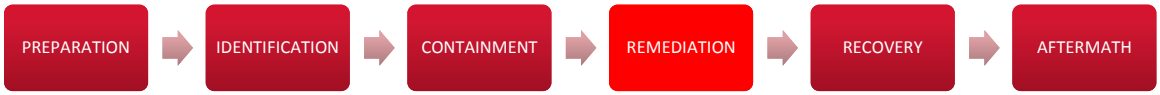

Figure 9: Phases of the phishing incident handling procedure

5.2 Other incident response best practices

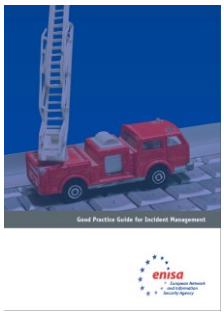

There are several incident handling procedures. One of them is presented by ENISA in its ENISA – Good Practice Guide for Incident Management.

Other incident response best practices worth mentioning are:

- National Institute of Standards and Technology – Computer Security Incident Handling Guide.[77]

The authors assumed that performing incident response effectively is a complex task and a successful incident response capability requires substantial planning and resources. The guide provides advice on how to establish computer security capabilities and how to handle incidents effectively.

- SANS Institute – A practical Social Media Incident Runbook.[78]

This guide provides advice on how to deal with incidents related to social media services such as Facebook, Twitter, LinkedIn and YouTube. These kinds of services have become new security risks for organisations and they are very often used by attackers. The result of these attacks are data breaches, phishing and DDoS attacks.

- New Zealand National Cyber Security Centre – New Zealand Security Incident Management Guide for Computer Security Incident response Teams (CSIRTs).[79]

[77] http://csrc.nist.gov/publications/nistpubs/800-61rev2/SP800-61rev2.pdf
[78] http://www.sans.org/reading-room/whitepapers/incident/practical-social-media-incident-runbook-34252?show=practical-social-media-incident-runbook-34252&cat=incident
[79] http://www.ncsc.govt.nz/sites/default/files/New%20Zealand%20Security%20Incident%20Management%20Guide%20for%20Computer%20Security%20Incident%20Response%20Teams%20(CSIRTs)_1.pdf

The governmental institutions issued a guide for all types of organisations in New Zealand on tackling computer incidents. The guide was developed in partnership with CERT Division of the Software Engineering Institute (SEI) at Carnegie Mellon University in Pittsburgh, USA, and is used as the part of supporting initiative for New Zealand National Cyber Security Strategy and the New Zealand Information Security Plan. It provides best practices and a basic framework for most organisations establishing a security incident management capability or reinforcing an existing one.

- Government of Canada – Cyber Incident Management Framework for Canada.[80]

The purpose of this framework is to provide a consolidated national approach to the management and coordination of potential or occurring cyber threats or incidents. Particular purposes are[81]:

- To clarify roles, responsibilities, authorities and capabilities of stakeholders in the cyber security community;
- To set expectations of all stakeholders on what they should be prepared to do, and what assistance they might obtain; and
- To serve as a vehicle for improving the management of cyber incidents and promoting coordination.

5.3 Incident classifications

As part of incident response best practices, CERTs need to be able to use a sensible and easily deployable incident classification – after all, it is important that incidents (and threats of incidents) can be put into at least some global categories in order to indicate what the incident/threat is about.

After studying the sources, and especially the ENISA Good Practice Guide for Incident Management and the recently improved eCSIRT.net taxonomy, we concluded that the following two classifications can be considered best practices:

- The classification of the Latvian national CERT CERT.LV.[82]
- The 10-year-old eCSIRT.net taxonomy which is still used by various teams in Europe, e.g. CERT Polska, however in a revised version of early 2013. This revised version is 95% backwards compatible with the 'old' taxonomy; however, some errors have been corrected and a few 'new' incident types like phishing have been added.

Both classifications are discussed in the next paragraphs.

CERT.LV Incident Classification

The CERT.LV Incident Classification consists of 11 types of Internet security attacks:

1. attacks on critical infrastructure,
2. attacks on Internet infrastructure, e.g. root or system-level attacks on any Server System, or any part of the backbone network infrastructure, denial of service attacks,
3. deliberate persistent attacks on specific resources, i.e. any compromise which leads or may lead to unauthorised access of systems,
4. widespread automated attacks against Internet sites, e.g. sniffing attacks, IRC 'social engineering' attacks, password cracking attacks,
5. threats, harassment, and other criminal offences involving individual user accounts,
6. new types of attacks or new vulnerabilities,

[80] http://www.publicsafety.gc.ca/cnt/rsrcs/pblctns/cbr-ncdnt-frmwrk/index-eng.aspx
[81] From the 'Scope of the Cyber Incident Management Framework' chapter of the document.
[82] This classification is no longer used by CERT.LV. It is mentioned as a good example of an approach to computer incident taxonomy.

7. botnets, i.e. activities related to network of compromised systems controlled by a party which is a source of the incident,
8. denial of service on individual user accounts, e.g. mail bombing,
9. forgery and misrepresentation, and other security-related violations of local rules and regulations, e.g. e-mail forgery, SPAM and etc.,
10. compromise of single desktop systems,
11. copyright violation.

Updated eCSIRT.net Taxonomy

Incident Class (mandatory input field)	**Incident Type (optional but desired input field)**	**Description / Examples**
Abusive Content	Spam	or "Unsolicited Bulk Email", this means that the recipient has not granted verifiable permission for the message to be sent and that the message is sent as part of a larger collection of messages, all having a functionally comparable content.
	Harmful speech	Discreditation or discrimination of somebody (e.g. cyberstalking, racism and threats against one or more individuals)
	Child/sexual/violence/...	Child pornography, glorification of violence, ...
Malicious Code	Virus	Software that is intentionally included or inserted in a system for a harmful purpose. A user interaction is normally necessary to activate the code.
	Worm	
	Trojan	
	Spyware	
	Dialler	
	Rootkit	
Information Gathering	Scanning	Attacks that send requests to a system to discover weak points. This includes also some kind of testing processes to gather information about hosts, services and accounts. Examples: fingerd, DNS querying, ICMP, SMTP (EXPN, RCPT,…), *port scanning*.
	Sniffing	Observing and recording network traffic (wiretapping)
	Social engineering	Gathering information from a human being in a non-technical way (eg, lies, tricks, bribes, or threats)
Intrusion Attempts	Exploiting known vulnerabilities	An attempt to compromise a system or to disrupt any service by exploiting vulnerabilities with a standardised identifier such as CVE name (eg, buffer overflow, backdoors, cross side scripting, etc).
	Login attempts	Multiple login attempts (guessing / cracking of passwords, brute force)
	New attack signature	An attempt using an unknown exploit
Intrusions	Privileged account compromise	A successful compromise of a system or application (service). This can have been caused remotely by a known or new vulnerability, but also by an unauthorized local access. Also includes being part of a botnet.
	Unprivileged account compromise	
	Application compromise	
	Bot	
Availability	DoS	By this kind of an attack a system is bombarded with so many packets that the operations are delayed or the system
	DDoS	

	Sabotage *Outage*	crashes. DoS examples are ICMP and SYN floods, Teardrop attacks and mail-bombing. DDoS often is based on DoS attacks originating from botnets, but also other scenarios exist like DNS Amplification attacks. However, the availability also can be affected by local actions (destruction, disruption of power supply, etc.) – or by Act of God, spontaneous failures or human error, without malice or gross neglect being involved.
Information *Content* Security	Unauthorised access to information Unauthorised modification of information	Besides a local abuse of data and systems the information security can be endangered by a successful account or application compromise. Furthermore attacks are possible that intercept and access information during transmission (wiretapping, spoofing or hijacking). Human/configuration/software error can also be the cause.
Fraud	Unauthorized use of resources	Using resources for unauthorized purposes including profit-making ventures (eg, the use of e-mail to participate in illegal profit chain letters or pyramid schemes)
	Copyright	Selling or installing copies of unlicensed commercial software or other copyright protected materials (Warez)
	Masquerade	Types of attacks in which one entity illegitimately assumes the identity of another in order to benefit from it
	Phishing	Masquerading as another entity in order to persuade the user to reveal a private credential.
Vulnerable	*Open for abuse*	Open resolvers, world readable printers, vulnerability apparent from Nessus etc scans, virus signatures not up-to-date,etc
Other	All incidents which do not fit in one of the given categories should be put into this class.	If the number of incidents in this category increases, it is an indicator that the classification scheme must be revised
Test	Meant for testing	Meant for testing

Table 14: eCSIRT.net classification schema

Annex A : CERT Survey Results

The survey took place in June 2013 and was targeted at the TF-CSIRT community. Some general characteristics of the results are as follows:

1. 21 CERTs reacted.
2. Types of CERT constituencies represented:
   - 43% research/NREN
   - 33% national/governmental/CIIP
   - 10% local government
   - 14% finance sector

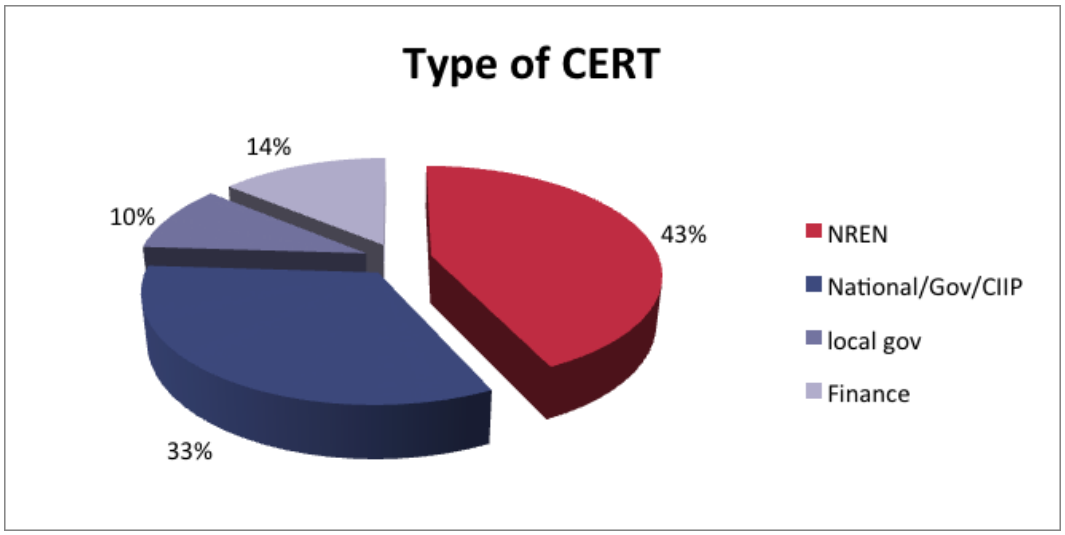

Figure 10: Types of CERT participating in the survey

3. Size of teams in FTE (full-time equivalents):
   - 32% teams have 1-3 FTE (<= 3)
   - 37% teams have 3-6 FTE (>3 <=6)
   - 21% teams have 6-12 FTE (>6 <=12)
   - the remaining teams are 36 and 70 FTE

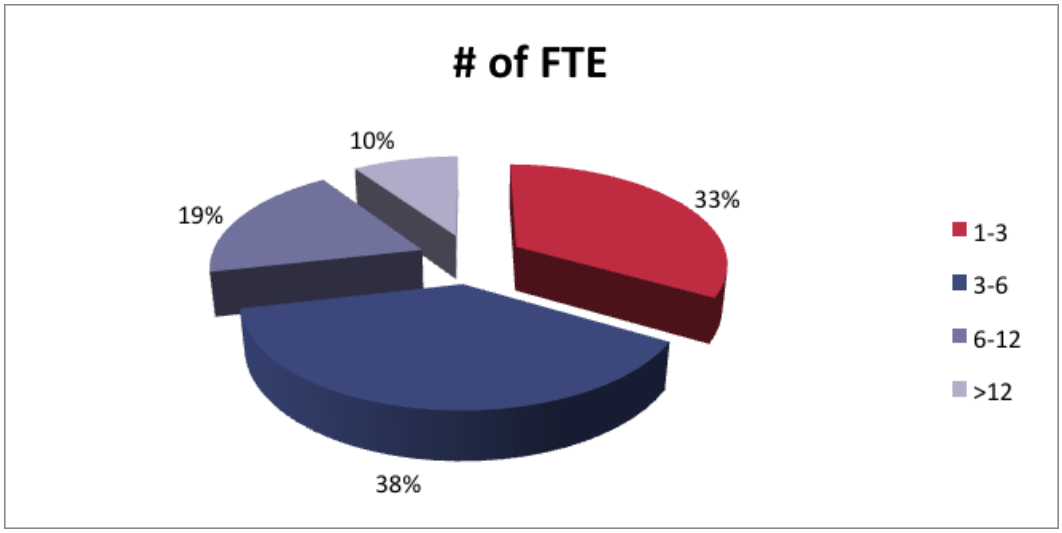

Figure 11: Level of employment in CERTs participating in the survey

4. The respondents have the following 'highest function' inside their team:
   - 32% is team member
   - 42% is the head or coordinator of the incident response team
   - 26% is the team's general manager

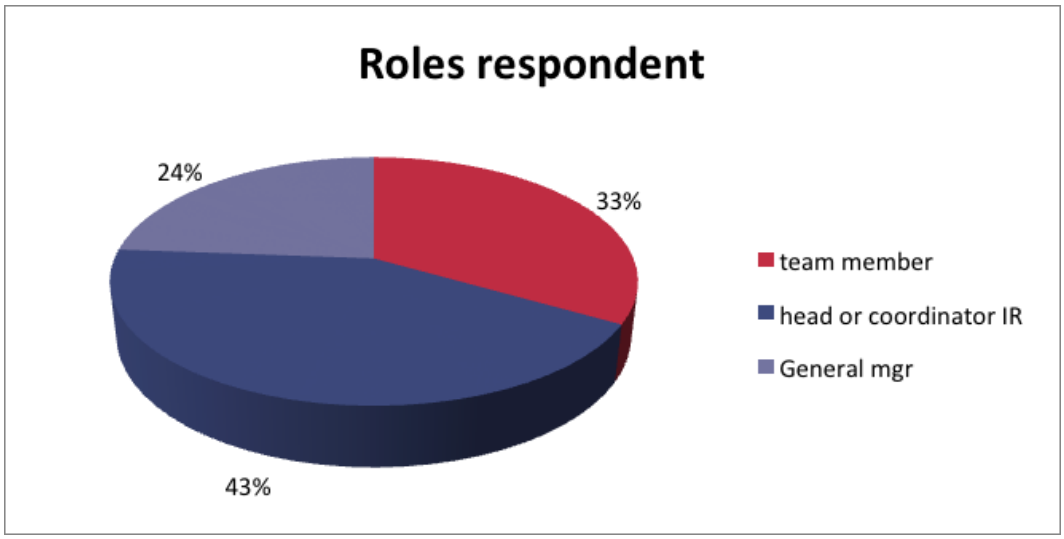

Figure 12: Roles of responders participating in the survey

In short, the survey has been filled out by respondents from government (57%), research/NREN (33%) and the finance sector (14%). One third of the teams are fairly small (3 FTE or less), but 58% of the teams are medium-sized (3-12 FTE) and 2 teams can be considered big (36 and 70 FTE respectively). Of the respondents, 68% are team leader or general manager – indicating that filling out this survey has been taken seriously by the teams, which is a good indicator for the quality of the results.

The survey results with regard to the content questions have been cited in various places in the main text of this guide. In all those cases, a referral to this Annex has been made.

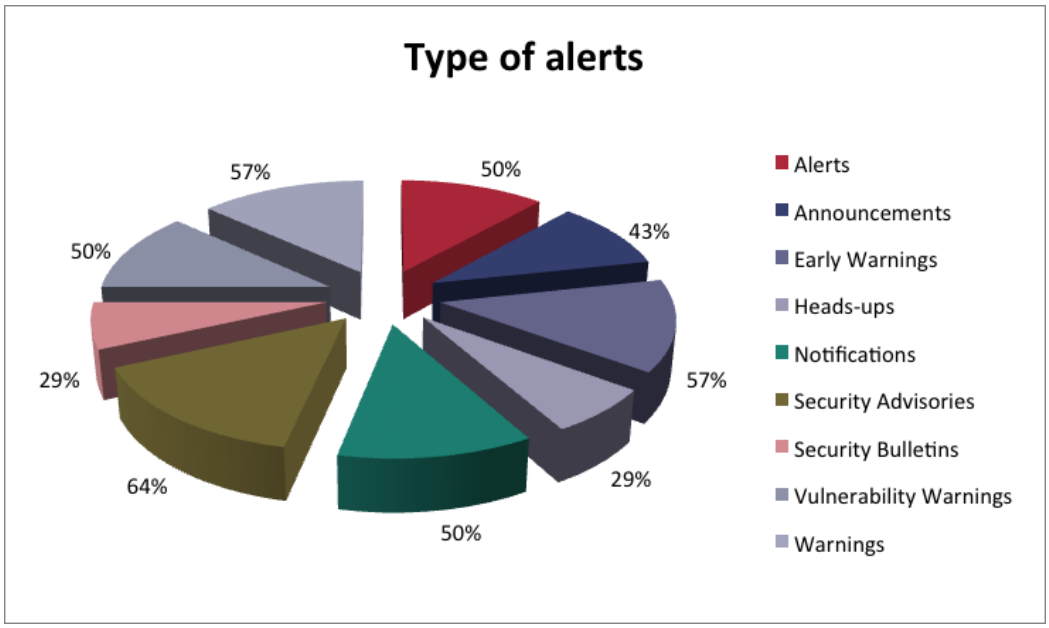

Figure 13: Types of alerts issuing by teams participating in the survey

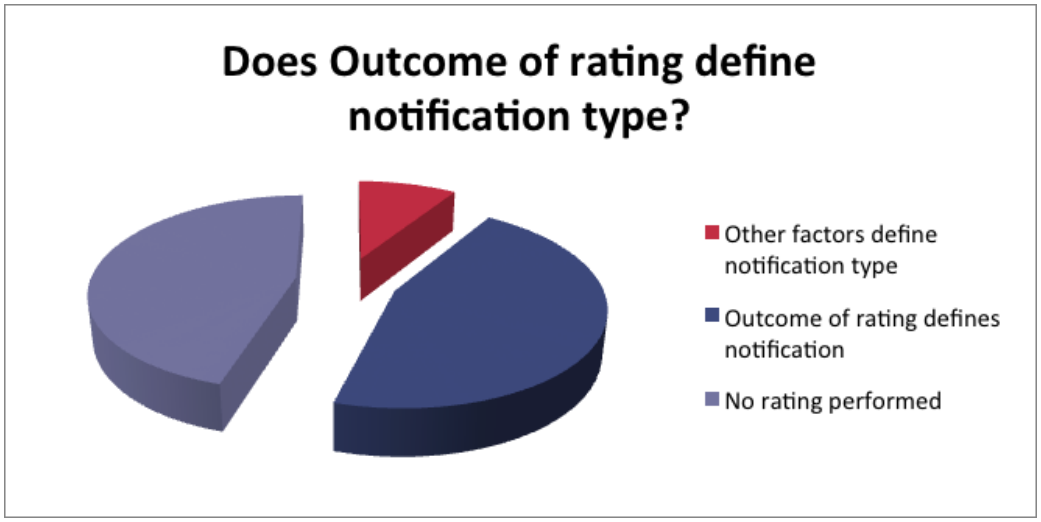

Figure 14: Outcome rating defines notification type

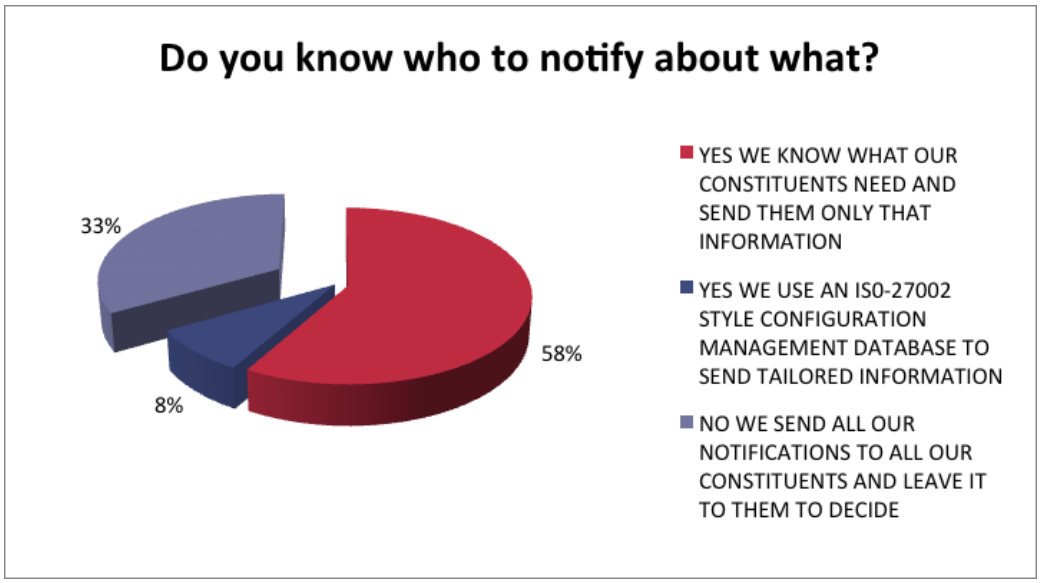

Figure 15: Do you know who to notify about what?

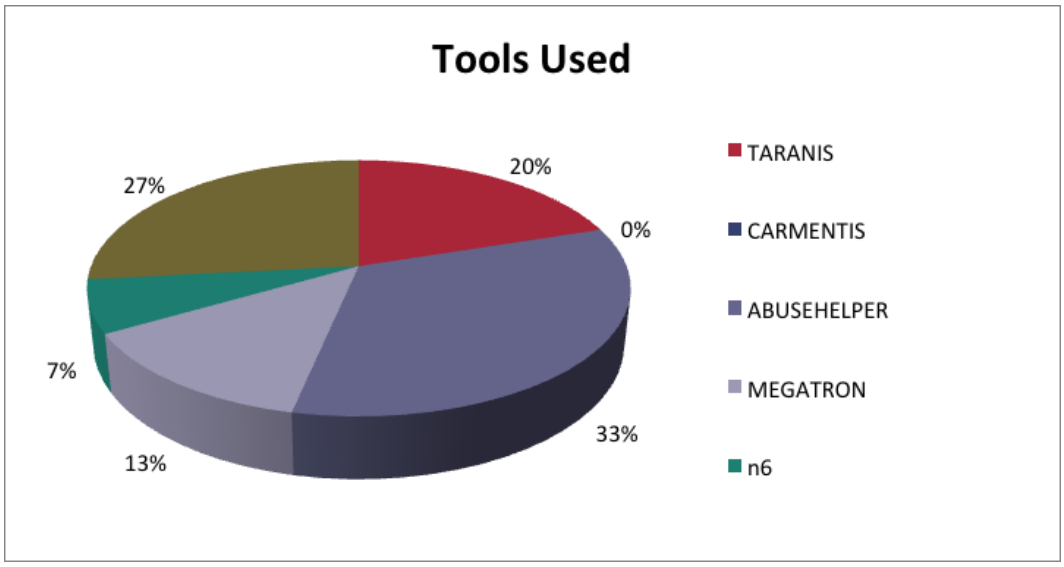

Figure 16: Tools used by teams participating in the survey

Annex B : Public security news feeds sources (Status 11/2013)

http://asert.arbornetworks.com/feed/
http://blog.acrossecurity.com/feeds/posts/default?alt=rss
http://blog.fortify.com/blog/feed/
http://blog.fortinet.com/feed/
http://blog.icann.org/feed/
http://blog.intego.com/feed/atom
http://blog.mozilla.com/security/feed/
http://blogs.mcafee.com/mcafee-labs/feed
http://feedproxy.google.com/Abusech
http://feedproxy.google.com/arstechnica/security
http://feedproxy.google.com/beasecurityadvisories
http://feedproxy.google.com/dvlabsblog
http://feedproxy.google.com/DvlabsNews
http://feedproxy.google.com/DvlabsPublishedAdvisories
http://feedproxy.google.com/DvlabsUpcomingAppearances
http://feedproxy.google.com/isc2Blog
http://feedproxy.google.com/MxLogicThreatblog
http://feedproxy.google.com/PrevxResearchBlog
http://feedproxy.google.com/schneier/excerpts
http://feedproxy.google.com/SkypeSecurity
http://feedproxy.google.com/Vrt
http://feedproxy.google.com/WatchfireApplicationSecurityInsider
http://feedproxy.google.com/wired27b
http://feedproxy.google.com/ZDI-Press
http://feedproxy.google.com/ZDI-Published-Advisories
http://feedproxy.google.com/zdnet/security
http://feedproxy.google.com/FE_research?format=xml
http://feedproxy.google.com/integrigysecurityblog?format=xml
http://feeds.ca.com/CaSecurityAdvisorNewlyDiscoveredVulnerabilities
http://feeds.ca.com/CaSecurityAdvisorVulnerabilityAlerts
http://feeds.ca.com/casecurityresponseblog/
http://feeds.ca.com/CaUnicenterPatchManagementAlerts
http://feeds.ca.com/CS_CASecurityAdvisorResearchBlog
http://feeds.feedburner.com/AttackAndDefenseLabs
http://feeds.feedburner.com/CsirtFoundry
http://feeds.feedburner.com/dsecrg_news
http://feeds.feedburner.com/dsecrg_pub
http://feeds.feedburner.com/dsecrg_vuln
http://feeds.feedburner.com/ForresterSRM
http://feeds.feedburner.com/GoogleChromeReleases
http://feeds.feedburner.com/Rapid7SecurityAlerts
http://feeds.feedburner.com/SansInstituteAtRiskPart2
http://feeds.feedburner.com/SansInstituteRRLast25

http://feeds.feedburner.com/SansInstituteWebcasts
http://feeds.feedburner.com/SANSPenTesting
http://feeds.feedburner.com/SCMagazineNews
http://feeds.feedburner.com/SecuritymarqitNieuws
http://feeds.feedburner.com/securityweek
http://feeds.feedburner.com/Snort
http://feeds.feedburner.com/SpiderlabsAnterior
http://feeds.feedburner.com/verizonbusiness/
http://feeds.feedburner.com/VxHeavens
http://feeds.reuters.com/reuters/technologyNews
http://feeds.sophos.com/en/rss2_0-sophos-advisories.xml
http://feeds.sophos.com/en/rss2_0-sophos-graham-cluley.xml
http://feeds.sophos.com/en/rss2_0-sophos-security-news.xml
http://feeds.sophos.com/en/rss2_0-sophos-sophoslabs-blog.xml
http://feeds.trendmicro.com/TrendMicroSecurityAdvisories
http://feeds2.feedburner.com/infoworldfeed
http://feeds2.feedburner.com/zeltser
http://feeds2.feedburner.com/zscaler/research
http://isc.sans.org/rssfeed_full.xml
http://pandalabs.pandasecurity.com/feed/rss/
http://php-security.org/feed/index.html
http://rss.feedsportal.com/c/32143/f/414040/index.rss
http://rss.feedsportal.com/c/32569/f/491736/index.rss
http://www.accuvant.com/blog/feed
http://www.acunetix.com/blog/feed/
http://www.barracudalabs.com/wordpress/index.php/feed/
http://www.coresecurity.com/content/advisories-feed
http://www.eeye.com/feeds?rss=Zero-Day-Tracker
http://www.eweek.com/rss-feeds-45.xml
http://www.eweekeurope.co.uk/category/news/news-security/feed
http://www.exploit-db.com/feed/
http://www.gcn.com/rss-feeds/security.aspx
http://www.honeyblog.org/feeds/index.rss2
http://www.honeynet.org/feed/blogfeed
http://www.kb.cert.org/vulfeed/
http://www.krebsonsecurity.com/feed/
http://www.microsoft.com/technet/security/advisory/RssFeed.aspx?securityadvisory
http://www.microsoft.com/technet/security/bulletin/RssFeed.aspx?snscomprehensive
http://www.norman.com/feeds/latest_blogs.rss/en
http://www.norman.com/feeds/security_articles.rss/en
http://www.novell.com/newsfeeds/rss/securityPatches.xml
http://www.offensive-security.com/feed
http://www.rsa.com/blog/rssfeed.aspx
http://www.ruby-lang.org/en/feeds/news.rss

http://www.symantec.com/connect/item-feeds/blog/691/feed
http://www.terena.org/feeds/news.rss
http://www.timesonline.co.uk/tol/feeds/rss/tech.xml
http://www.virusbtn.com/library/feeds/news.rdf
https://community.rapid7.com/blogs/feeds/posts
https://hermes.opensuse.org/feeds/62042.rdf
https://www.trustwave.com/feeds/advisories/

Annex C : Relevant ENISA documents cross-reference

'Proactive detection of network security incidents, CERT survey analysis'

http://www.enisa.europa.eu/activities/cert/support/proactive-detection/survey-analysis

'Good Practice Guide for Incident Management'

http://www.enisa.europa.eu/activities/cert/support/incident-management/files/good-practice-guide-for-incident-management

'Baseline Capabilities of n/g CERTs – Updated Recommendations 2012'

http://www.enisa.europa.eu/activities/cert/support/files/updated-recommendations-2012

'EISAS Large-Scale Pilot – Collaborative Awareness Raising for EU Citizens & SMEs'

http://www.enisa.europa.eu/activities/cert/other-work/eisas_folder/eisas-large-scale-pilot

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

Athens Office
1 Vass. Sofias & Meg. Alexandrou
Marousi 151 24, Athens, Greece

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