

Handbook for Computer Security Incident Response Teams (CSIRTs)

Moira J. West-Brown
Don Stikvoort
Klaus-Peter Kossakowski
Georgia Killcrece
Robin Ruefle
Mark Zajicek

First release: December 1998
2nd Edition: April 2003

HANDBOOK
CMU/SEI-2003-HB-002



CarnegieMellon
Software Engineering Institute

Pittsburgh, PA 15213-3890

Handbook for Computer Security Incident Response Teams (CSIRTs)

CMU/SEI-2003-HB-002

Moira J. West-Brown
Don Stikvoort
Klaus-Peter Kossakowski
Georgia Killcrece
Robin Ruefle
Mark Zajicek

First release: December 1998
2nd Edition: April 2003

Networked Systems Survivability Program

Unlimited distribution subject to the copyright.

The original version of this handbook was provided with funding from the following organizations:

U.S. National Science Foundation (NSF);
SURFnet bv;
SURFnet ExpertiseCentrum bv;
M&I/STELVIO bv;
German Federal Ministry of Education, Science, Research and Technology (Bundesministerium fuer Bildung, Wissenschaft, Forschung und Technologie);
Verein zur Foerderung eines Deutschen Forschungsnetzes e.V. (DFN-Verein).

Funding for the revised edition of this handbook was provided by the Software Engineering Institute.

This report was prepared for the

SEI Joint Program Office
HQ ESC/DIB
5 Eglin Street
Hanscom AFB, MA 01731-2116

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

FOR THE COMMANDER



Christos Scondras
Chief of Programs, XPK

This work is sponsored by the U.S. Department of Defense. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2003 by Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number F19628-00-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (<http://www.sei.cmu.edu/publications/pubweb.html>).

Table of Contents

Preface to the Second Edition	ix
Preface to the First Edition	xi
Acknowledgements	xiii
Abstract.....	xv
1 Introduction	1
1.1 Scope of the Document	4
1.2 Intended Audience.....	5
1.3 Use of This Document	6
1.4 Document Structure.....	7
2 Basic Issues.....	9
2.1 CSIRT Framework.....	9
2.1.1 Mission Statement	10
2.1.2 Constituency	11
2.1.3 Place in Organization.....	17
2.1.4 Relationship to Other Teams	19
2.2 Service and Quality Framework.....	21
2.3 CSIRT Services.....	23
2.3.1 Service Categories.....	24
2.3.2 Service Descriptions	25
2.3.3 Selection of Services	34
2.4 Information Flow.....	35
2.5 Policies.....	38
2.5.1 Attributes	39
2.5.2 Content	40
2.5.3 Validation.....	41
2.5.4 Implementation, Maintenance, and Enforcement	42
2.6 Quality Assurance	42
2.6.1 Definition of a Quality System	43
2.6.2 Checks: Measurement of Quality Parameters	45

2.6.3	Balances: Procedures to Assure Quality.....	47
2.6.4	Constituents' View of Quality	48
2.7	Adapting to Specific Needs	48
2.7.1	The Need for Flexibility	50
2.7.2	Legal Issues	51
2.7.3	Institutional Regulations	58
3	Incident Handling Service	61
3.1	Service Description	61
3.1.1	Objective	62
3.1.2	Definition	63
3.1.3	Function Descriptions	64
3.1.4	Availability	65
3.1.5	Quality Assurance	65
3.1.6	Interactions and Information Disclosure.....	66
3.1.7	Interfaces with Other Services.....	66
3.1.8	Priority.....	66
3.2	Service Functions Overview.....	67
3.3	Triage Function	69
3.3.1	Use of Tracking Numbers	70
3.3.2	Use of Standard Reporting Forms	73
3.3.3	Preregistration of Contact Information	75
3.4	Handling Function	76
3.4.1	Incident Life Cycle	77
3.4.2	Incident Analysis	79
3.4.3	Tracking of Incident Information.....	91
3.5	Announcement Function	92
3.5.1	Announcement Types.....	93
3.5.2	A Priori Considerations.....	95
3.5.3	Announcement Life Cycle.....	97
3.6	Feedback Function	100
3.7	Interactions	102
3.7.1	Points of Contact.....	103
3.7.2	Authentication.....	106
3.7.3	Secure Communication	109
3.7.4	Special Considerations.....	110
3.8	Information Handling.....	119
3.8.1	Information Collection.....	119
3.8.2	Information Verification.....	120
3.8.3	Information Categorization.....	121
3.8.4	Information Storage.....	122
3.8.5	Information Sanitizing and Disposal.....	123

3.8.6	Prioritization Criteria.....	124
3.8.7	Escalation Criteria.....	128
3.8.8	Information Disclosure	132
4	Team Operations	137
4.2	Fundamental Policies	141
4.2.1	Code of Conduct.....	141
4.2.2	Information Categorization Policy.....	143
4.2.3	Information Disclosure Policy	144
4.2.4	Media Policy	148
4.2.5	Security Policy	149
4.2.6	Human Error Policy.....	150
4.3	Continuity Assurance.....	151
4.3.1	Continuity Threats.....	152
4.3.2	Workflow Management	155
4.3.3	Out-Of-Hours Coverage.....	157
4.3.4	Off-Site Coverage	159
4.4	Security Management.....	159
4.5	Staff Issues.....	166
4.5.1	CSIRT Staff.....	167
4.5.2	Hiring Staff	169
4.5.3	Arrival and Exit Procedures.....	171
4.5.4	Training Staff	172
4.5.5	Retaining Staff	174
4.5.6	Extension of Staff.....	175
5	Closing Remarks	177
5.1	Closing Remarks from the First Edition.....	177
5.2	Closing Remarks for the Second Edition	178
	Appendix A: About the Authors	181
	Appendix B: Glossary.....	187
	Bibliography.....	193

List of Figures

Figure 1: CSIRT Within an Organization	18
Figure 2: CSIRT Peer Relationships.....	21
Figure 3: Service and Quality Framework as Derived from Mission Statement.....	22
Figure 4: Incident Handling Service Functions	67
Figure 5: CERT/CC Incident Handling Life Cycle	77
Figure 6: CERT/CC Code of Conduct.....	143

List of Tables

Table 1:	Examples of CSIRT Types With Associated Missions and Constituencies	12
Table 2:	Possible Authority Relationships Between a CSIRT and Its Constituency	15
Table 3:	Service Description Attributes.....	23
Table 4:	List of Common CSIRT Services	25
Table 5:	Examples of Possible Information Flow to and from the Incident Handling Service	37
Table 6:	Basic Policy Attributes	40
Table 7:	Policy Content Features	41
Table 8:	Examples of Dynamic Environment Factors and Their Impact on CSIRTs	50
Table 9:	Examples of Liability Issues Arising From Omission	56
Table 10:	Examples of Liability Issues Arising From the Content of Signed Contracts.....	56
Table 11:	Examples of Liability Issues Arising From Information Disclosure	57
Table 12:	Range of Possible Incident Handling Service Objectives Based on Differing Team Types.....	62
Table 13:	Possible Instantiations of Handling Function Attributes.....	76
Table 14:	Analysis Depth Factors	82
Table 15:	Notable Characteristics of Log Files	83
Table 16:	Incident Tracking Information	91
Table 17:	Possible Inter-Team Support Types.....	113
Table 18:	Considerations for Information Sharing	113

Preface to the Second Edition

We have often been asked whether an updated version of the CSIRT Handbook would ever be released. Periodically we have reviewed the document and found that most of the material and guidance provided are still current, relevant, and helpful to new and existing teams. Some of the examples included and organizations discussed were dated, but otherwise the concepts and recommendations covered are still valid for today's work.

In the summer of 2002, the CERT[®] CSIRT Development Team began collaboration with the Trusted Introducer for European Computer Security Incident Response Teams (CSIRTs) service to create a standard set of service descriptions for CSIRT functions. As we finished that document¹ it became apparent that we should, indeed, update the CSIRT Handbook to include this new list of services. As we began to revise the document we felt it was also time to update any out-of-date examples and also any out-of-date terminology. We also included, where appropriate, references to new discussion topics, resources, or CSIRT operational activities that we believe are relevant to the information discussed in this handbook.

In the long run though, we elected to minimize the changes to the original as much as possible. These are the main changes that have been made:

1. Many of the examples provided in the handbook have been updated. We have kept a number of the previous examples, as they are still true conceptually and the guidance available still proves to be useful today. More recent examples have been added that we hope will be more applicable for today's readers.
2. The new CSIRT service definitions have been incorporated throughout the handbook.
3. The handbook has been aligned with other new documents that we have produced or are in the process of producing, specifically the new *Organizational Models for CSIRTs* handbook. This document is a companion piece to the CSIRT Handbook. It provides detailed information on the types of organizational structures and corresponding services that may be implemented to provide a CSIRT capability. We have timed the release of this updated version of the CSIRT Handbook with the publication of the *Organizational Models for CSIRTs* handbook.

As stated in the original CSIRT Handbook preface: We can learn from the experiences of others— and we do this every day. So if you have comments on this version, if you want to

¹ The list of services is available at <http://www.cert.org/csirts/services.html>.

share your opinions, or if you have other suggested additions, please contact us. We regularly attend FIRST Conferences and other CSIRT events, and we can be contacted in person or reached as a group by sending email to the following address: csirt-handbook@cert.org.

Preface to the First Edition

The number of computer security incident response teams (CSIRTs) continues to grow as organizations respond to the need to be better prepared to address and prevent computer security incidents. Just as computer science has struggled to be recognized as a scientific field in its own right, computer security has struggled to be recognized as an essential component of computer science. Similarly, the need for CSIRTs should be recognized within the security arena. As new teams have attempted to form, they have faced the hurdles of having to justify the need for their existence and gaining support and understanding of the problems that they are trying to address. If they have managed to overcome those hurdles, then they have had an additional challenge to face: the lack of documented information on how to effectively form and operate a CSIRT and gain recognition for it. So the need for a handbook of this type is long overdue.

The idea to write this handbook resulted from an email discussion between the original authors (West-Brown, Stikvoort, and Kossakowski) in the summer of 1996. At that time the authors were each working on similar projects in their own organizations: helping other CSIRTs form and develop corresponding policies and procedures. The authors saw a growing demand from newly forming teams for help and assistance in their formation and operation and realized that there were insufficient experts available to fulfill this growing demand. Because the task of forming and operating a CSIRT is fraught with pitfalls that can result in the demise of a team, it was clear that to ensure an infrastructure of competent and respected CSIRTs, supporting information and guidance would be imperative for success.

As with many projects of this type, the handbook development has taken longer than was originally anticipated; it has been something that we've tried to work on when we had spare time. Given that the field in which we work is so dynamic and demanding and experts are in short supply, that spare time has generally been carved from late nights and weekends. We had the luxury of spending most of a week in October 1996 together devoted to scoping the handbook, which resulted in a 22-page structured outline of the issues. With that foundation in place, we returned to our own organizations and began the slow process of writing the content of the various sections and continued document development.

We hope that you will find this resulting first edition a useful reference document in the formation, management, and operation of your CSIRT. We have based material in this handbook on our experiences in forming and operating our own organization's CSIRTs and through assisting other CSIRTs in their formation and operation.

Acknowledgements

We have many people to thank for contributions that made the original handbook possible. First our thanks go to the CERT[®] Coordination Center (CERT/CC), Verein zur Foerderung eines Deutschen Forschungsnetzes e.V. (DFN-Verein), M&I/STELVIO, U.S. National Science Foundation (NSF), SURFnet ExpertiseCentrum bv, and SURFnet bv. These organizations supported this effort through funding, allowed us to spend time and resources on this project, and gave us the opportunity to gain experience and flourish in this field. Special thanks go to our colleagues at CERT/CC, CERT-NL, and DFN-CERT who were busy handling incidents and addressing computer security emergencies at times when we were working on deadlines for this project.

Our thanks also go to the organizations that have sought our help in forming and operating their CSIRTs. Addressing their probing questions and having them share their differing needs and situations with us has enabled us to obtain a more rounded view of the field and has broadened the scope of our experience.

We sought technical review of the first draft of this document from a variety of individuals. We selected a cross section of reviewers ranging from those we knew were interested in forming a CSIRT and were new to the field of computer security, to those who have considerable operational experience from a technical or management perspective. Knowing how busy such people are, we selected 15 reviewers in the hope that maybe 8 would have the opportunity to read the draft and provide feedback in the short time available. To our amazement, 14 of the reviewers provided us with feedback of some sort. The 15th reviewer sent his apologies explaining that he was unavailable due to illness. Our thanks go to all the reviewers who found time to comment on the first draft of this handbook (affiliations shown are from the time of the first edition):

David Finch (MOREnet)
Eduardo Garcia (Price-Waterhouse)
John Horton (DANTE)
Erik Huizer (SURFnet ExpertiseCentrum bv)
Larry J. Hughes, Jr. (NorthWestNet)
Georgia Killcrece (CERT/CC)
Kathleen Kimball (Pennsylvania State University)
Wolfgang Ley (DFN-CERT)

[®] CERT is registered in the U.S. Patent and Trademark Office.

Hannes P. Lubich (SWITCH-CERT)
Jorgen Bo Madsen (NORDUnet CERT)
Ken McNulty (SEI)
Maj. Byron Thatcher (AFIWC/AFCERT)
Wietse Venema (IBM)
Mark Zajicek (CERT/CC)

We would particularly like to acknowledge the efforts of Larry J. Hughes, Jr., Georgia Killcrece, Wolfgang Ley, Hannes P. Lubich, and Jorgen Bo Madsen who all went above and beyond the call of duty by providing very extensive and detailed comments.

The first draft of this document was an interesting work written by three authors with differing native tongues. The task of taking that draft and generating a document ready for technical review was placed on the shoulders of Bill McSteen (a technical writer/editor at the SEI). Bill did an excellent job of smoothing over the bumps in the flow of the document so that the reviewers were able to focus on technical and structural comments. Bill's continued effort was as essential to provide the final version of this document.

Our thanks for this second edition go to Pamela Curtis, also a technical writer/editor at the SEI, without whose expert help this task would have been much more difficult to accomplish. Her technical and editorial assistance helped tremendously and enabled the authors to focus on the content of the material.

We would also like to extend heartfelt thanks to our Technical Manager, Dr. Barbara Laswell, who encouraged us to revise the Handbook and who generously provided the resources and time for us to devote to this revised text.

Finally, we would like to thank our families. Without their support, understanding, and encouragement, none of us would have been able to complete our portions of this work nor would we have found it an enjoyable exercise.

Abstract

This document provides guidance on forming and operating a computer security incident response team (CSIRT). In particular, it helps an organization to define and document the nature and scope of a computer security incident handling service, which is the core service of a CSIRT. The document explains the functions that make up the service; how those functions interrelate; and the tools, procedures, and roles necessary to implement the service. This document also describes how CSIRTs interact with other organizations and how to handle sensitive information. In addition, operational and technical issues are covered, such as equipment, security, and staffing considerations.

This document is intended to provide a valuable resource to both newly forming teams and existing teams whose services, policies, and procedures are not clearly defined or documented. The primary audience for this document is managers who are responsible for the creation or operation of a CSIRT or an incident handling service. It can also be used as a reference for all CSIRT staff, higher level managers, and others who interact with a CSIRT.

1 Introduction

The evolution of the Internet has been widely chronicled. Resulting from a research project that established communications among a handful of geographically distributed systems, the Internet now covers the globe as a vast collection of networks made up of millions of systems.

The Internet has become one of the most powerful and widely available communications mediums on earth, and our reliance on it increases daily. Governments, corporations, banks, and schools conduct their day-to-day business over the Internet. With such widespread use, the data that resides on and flows across the network varies from banking and securities transactions to medical records, proprietary data, and personal correspondence.

The Internet is easy and cheap to access, but the systems attached to it lack a corresponding ease of administration. As a result, many Internet systems are not securely configured. Additionally the underlying network protocols that support Internet communication are insecure, and few applications make use of the limited security protections that are currently available.

The combination of the data available on the network and the difficulties involved in protecting the data securely make Internet systems vulnerable attack targets. It is not uncommon to see articles in the media referring to Internet intruder activities.

But, exploitation of security problems on the Internet is not a new phenomenon. In 1988 the “Internet Worm” incident occurred and resulted in a large percentage of the systems on the network at that time being compromised and temporarily placed out of service. Shortly after the incident, a meeting was held to identify how to improve response to computer security incidents on the Internet. The recommendations resulting from the meeting included a call for a single point of contact to be established for Internet security problems that would act as a trusted clearinghouse for security information. In response to the recommendations, the CERT[®] Coordination Center (also known as the CERT/CC and originally named the Computer Emergency Response Team) was formed to provide response to computer security

[®] CERT is registered in the U.S. Patent and Trademark Office.

incidents on the Internet [CERT/CC 1997b]. The CERT/CC was one of the first organizations of this type—a computer security incident response team (CSIRT²).

A CSIRT can most easily be described by analogy with a fire department. In the same way that a fire department has an emergency number that you can call if you have or suspect a fire, similarly a CSIRT has a number and an email address that you can contact for help if you have or suspect a computer security incident. A CSIRT service doesn't necessarily provide response by showing up on your doorstep (although some do offer that service); they usually conduct their interactions by telephone or via email.

Another similarity between fire departments and CSIRTs is that responding to emergencies is only part of the service provided. Just as important is trying to prevent emergencies from occurring in the first place. So just as a fire department offers fire safety education to raise awareness and encourage best practices, CSIRTs produce technical documents and undertake education and training programs for the same purpose. In the area of improvement, a fire department will influence laws to ensure improved safety codes and fire-resistant products. Similarly CSIRTs participate in forums to improve baseline security standards.

When the Internet Worm incident occurred, the size of the network was estimated at 60,000 hosts; a decade later there were more than 36 million hosts on the Internet and a corresponding increase in intruder activity. The January 2003 Internet Domain Survey [ISC 2003] shows 171.6 million hosts advertised in the Domain Name Service. Clearly a single CSIRT is unable to effectively serve such a vast constituency. In particular a single CSIRT wouldn't be able to address the individual needs of the diverse communities that make up the Internet due to time zone, language, cultural, and organizational issues. Correspondingly, a number of organizations have foreseen the need to be better prepared to respond to intruder activity affecting their community [West-Brown 1995]. This has resulted in a surge of interest in the formation of CSIRTs.

Hundreds of CSIRTs around the world have formed since 1988, and they, like newly forming CSIRTs today, face many challenges as they strive to become operational. Newly forming teams commonly seek guidance and assistance in determining the scope and range of their services and in forming their operational policies and procedures [Pethia 1990a, Pethia 1990b]. When this CSIRT Handbook was originally published in 1998, there were not as many resources available to help new teams establish appropriate and reliable services. Today

² When the first edition of this handbook was published, the term “incident response” was used to describe the core service of a CSIRT (hence the convention for the “IR” letters in the CSIRT acronym). As our understanding of such teams has matured over time, incident response has become one component of a much broader “incident handling” service that encompasses more than just response to an event. We have, therefore, adopted the convention of “incident handling” services throughout this handbook (as well as in our other publications [CERT/CC 2002a]). However, we still continue to use the acronym “CSIRT,” since it is a generic description for a team and is a term that has been widely adopted by the community.

there are many more articles, books, tutorials,³ other documents, and Web resources available.⁴ There are also various organizations, such as the Forum of Incident Response and Security Teams (FIRST) and the TERENA-sponsored TF-CSIRT (a task force for the collaboration of incident response teams in Europe), who promote collaboration among teams and provide resources for helping new and existing teams [TERENA 1995]. The good news is that today's newly forming CSIRTs need not fend for themselves (learning only from their own experiences or making costly mistakes); they can now leverage the experiences of many others to help them develop and implement more effective teams.^{5,6,7}

Although more information is now available to help organizations build their CSIRT capability, unfortunately there is still little available in the area of operational policies and procedures, e.g., generic CSIRT policies, procedures, and templates that can be adapted or revised for use by newly forming teams. Either existing teams have nothing documented to share or they are unable to share their documentation due to its sensitive nature. Seeking expert advice is also difficult because there is still a shortage of experts in the field. Existing experts are highly sought after, have little time to make available, and can be expensive to engage.

Once operational, the need for well-defined services, policies, and procedures does not diminish. Existing CSIRTs lacking clearly defined services commonly suffer from recurring operational problems. For example, they rely on their existing staff to pass on their operational experience to new staff. All too frequently, the consistency, reliability, and levels of service exhibited by such CSIRTs fluctuate dramatically due to the varied perceptions of each of the team members. As a consequence, the constituency served by these CSIRTs may have a false impression of the services offered, which jeopardizes rapport between a CSIRT and its constituency that is essential to the success of the team. Clearly defined and documented services will help the team and, more importantly, will provide guidance for the team's constituency, enabling them to understand the services offered by the CSIRT and how those services should be accessed.

³ The CERT/CC, for example, offers a suite of courses for both CSIRT managers and incident handlers. See http://www.cert.org/nav/index_gold.html.

⁴ See the CERT CSIRT Development Team's List of Resources for more information: <http://www.cert.org/csirts/resources.html>.

⁵ See also the Workshops on Computer Security Incident Handling, Forum of Incident Response and Security Teams, 1989-2002.

⁶ Sparks, Sandy; Fithen, Katherine; Swanson, Marianne; & Zechman, Pat. "Establishing an Incident Response Team." Tutorial, 9th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, Bristol, U.K., June 1997.

⁷ Stikvoort, Don & Kossakowski, Klaus-Peter. "Incident Response Teams: the European Perspective." 8th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, San Jose, Calif., July 1996.

1.1 Scope of the Document

This document provides guidance on the generic issues to consider when forming and operating a CSIRT. Relating back to our fire department analogy, providing an effective service is a complex operation. It can only be a success if it is based on appropriate policies and procedures and if it addresses a range of both reactive and proactive issues. A fire department can be a volunteer or directly funded operation. The service provided is based on available resources and funding. CSIRTs are under the same cost-cutting demands as other organizations. So they must constantly make the tradeoff between the range and levels of service that they would like to provide and what they can afford to provide. This includes identifying the CSIRT services, policies, and procedures appropriate for a given situation. It also means identifying those operational issues that must be addressed to implement an efficient incident handling capability.

In particular, this document helps an organization to define and document the nature and scope of a computer security incident handling service. To this end, we discuss the functions that make up the service; how those functions interrelate; and the tools, procedures, and roles necessary to implement the service. We also focus on incident analysis. Just as a fire department may investigate a fire to understand how it came about (e.g., act of nature, arson, or an electrical design fault), a CSIRT tries to understand how an incident occurred. While a fire department's analysis will include sifting through ashes, a CSIRT's will include looking at system logs and any files left behind by an intruder.

A fire department needs to coordinate with other fire departments who it may call (or be called) on for reinforcements in times of peak demand or to address a crisis. It must interact with other emergency services to respond appropriately and provide law enforcement with the information that it legally requires. This document will discuss how CSIRTs interact with other organizations, such as the sites that report security problems to it, other CSIRTs, law enforcement, and the media. A fire department must handle information, some of which is sensitive as it may pertain to the perpetrator of a crime. Similarly a CSIRT must handle information appropriately. Most CSIRTs offer client confidentiality in the same way that many crisis lines do, shielding the reporters and victims from public disclosure. This topic is critical to the survival of a CSIRT, because if it cannot be trusted to handle information appropriately, nobody will report to it, rendering the CSIRT almost useless. Consequently, information handling is an essential issue of discussion in this handbook.

Some CSIRTs have dedicated staff while others pull together part-time, volunteer staff and trusted security experts to address a given security crisis.⁸ A CSIRT's staff is its interface with the world, and the image that its staff members project through the way they conduct

⁸ More information about different staffing models will be available in a companion document, *Organizational Models for CSIRTs*, to be published in 2003 on our Web site at <http://www.cert.org/csirts/>.

themselves, and the quality of service they provide, are paramount to the CSIRT's success. Finding appropriately qualified staff can be difficult, since such staff are in great demand. However, all too often people responsible for hiring CSIRT staff unknowingly look for the wrong set of skills and qualities in potential employees. Consequently we discuss staffing and hiring issues and steps that you can take to ensure that CSIRT staff provide a consistent, warm, and professional interface for your team.

A CSIRT may provide a range of other services in addition to the incident handling service, such as vulnerability handling and/or the provision of intrusion detection services. Although we have included a high-level description of these other services, the specific procedures and policies are beyond the scope of this document.

The material in this document is presented in a form that is suitably generic to enable the reader to apply it to any type of CSIRT environment, from a fee-for-service team, to an in-house team for a given organization, or an international coordination center.

1.2 Intended Audience

While many new CSIRTs have formed and become operational, the increase in the number of these teams has not kept pace with Internet growth and intruder activity. Many more organizations are recognizing the need for a CSIRT to address their specific needs. Anticipating this need, we have targeted this handbook at those individuals who will be most heavily involved in the establishment of CSIRTs.

The primary audience for this document consists of managers responsible for at least one of the following:

- the creation of a CSIRT
- the operation of a CSIRT
- the creation of an incident handling service
- the operation of an incident handling service

As well as being a useful reference for higher management levels and all CSIRT staff, this document can also be of use to other individuals who interact with a CSIRT and would benefit from an awareness of the issues that affect a CSIRT, such as

- members of the CSIRT constituency
- law enforcement
- media relations

We recognize that some organizations may choose to “outsource” their CSIRT services⁹ to other managed service providers, although we have not focused specifically on that audience in this handbook. We do believe, however, that the information contained in this handbook can be used by such providers and adapted to fit their approaches for providing fee-based services to an organization or enterprise. To that end, the handbook can be a valuable resource document for these service providers in identifying the same types of issues that other CSIRTs face in providing services to their constituency.

1.3 Use of This Document

This document is intended to provide a valuable resource to both newly forming teams and existing teams whose services, policies, and procedures are not clearly defined or documented. Ideally this document should be used at the early stages when an organization has obtained management support and funding to form a CSIRT, prior to the team becoming operational. However, the material can still be a very useful reference document for operational teams.

This material can be used by a newly forming team as the basis for understanding the issues involved in establishing a CSIRT. The information can then be used to assist the development of detailed domain- or organization-specific service definitions, policies, procedures, and operational issues. As a result of applying the material provided in this document, an organization should be on a fast track to a documented, reliable, effective, and responsible incident handling service.

In addition, an existing team can use this document to ensure they have covered the main issues and options that are appropriate for their organization when developing their incident handling service.

Where applicable, the authors identify approaches that have proved successful, as well as pitfalls to avoid. In addition, various alternatives are described that might suit a particular situation or be applicable for a given type of team, such as an international response team, a national response team, an Internet service provider (ISP) team serving its customers, or a team for a single organizational entity such as a university or corporation. It is important, however, to note that this material is provided for reference and guidance. It is not intended to dictate the range or content of services, policies, and procedures that any given team should implement. These must be determined based on the individual needs of the CSIRT and the constituency it serves. Hence, we encourage you to use the material provided in this

⁹ The Networked Systems Survivability Program at the Software Engineering Institute with funding support from the General Services Administration Federal Computer Incident Response Center (GSA FedCIRC) developed the report “Outsourcing Managed Security Services.” This document is available from the CERT Web site at <http://www.cert.org/security-improvement/modules/omss/>.

document to understand the issues appropriate for your team’s unique environment and to choose approaches that meet your team’s particular goals, needs, and requirements.

1.4 Document Structure

The rest of this document is organized as follows. Chapter 2 presents the basic framework of the CSIRT model and describes basic issues that need to be considered and addressed by every CSIRT. It also introduces general CSIRT terminology and concepts including the importance of a clearly defined constituency, generation and implementation of policies, and the impact of organizational and legal issues on a CSIRT. It introduces a range of services that a CSIRT might provide and discusses how those services interact with the incident handling service. This sets the context for the main focus for this document—the incident handling service, which is described in detail in Chapter 3. Chapter 3 describes the construction of an incident handling service and its functional components. Additionally, it discusses the range and nature of interactions that are associated with an incident handling service and how information (mostly of a sensitive nature) is handled. For completeness, Chapter 4, “Team Operations,” addresses practical operational and technical issues that every CSIRT must consider. These issues, such as equipment, security, and staffing considerations, are not all exclusive to an incident handling service, but they are critical to its success. The document concludes with some closing remarks followed by information about the authors, a bibliography of CSIRT-related materials, and a glossary of abbreviations and terms.

2 Basic Issues

A CSIRT may offer a range of services. However, it must at least provide an implementation of the incident handling service discussed later in this chapter and covered in depth in Chapter 3. Without providing at least a component of the incident handling service, the team cannot be called a CSIRT. Consider the analogy with a fire department. A fire department may provide a range of services (fire prevention, awareness, training), and it may undertake fire safety inspections. But at the core is the emergency response component. By providing the emergency fire department, it stays up-to-date and in touch with reality, and it gains community trust, respect, and credibility. Similarly, in an attempt to reduce the effect of incidents through early detection and reporting or to prevent incidents, a team can be proactive through awareness, training, and other services; but without the incident handling service, the team is not a CSIRT.

This chapter presents the basic framework of the CSIRT model and describes the issues that affect every CSIRT. These issues need to be considered and addressed for all CSIRTs regardless of their size, nature, or scope. We begin by describing the CSIRT framework in terms of what it sets out to do (mission), for whom (constituency), what its roots look like (place in organization), and who its peers are (relationship to others). Next, we examine a framework derived directly from the mission statement: the service and quality framework, featuring CSIRT services, quality assurance, policies as major components, and information flow as an essential boundary condition. In the last section, we review the issues faced when adapting a CSIRT to the specific needs of its environment, where the legal issues are a particularly important component.

2.1 CSIRT Framework

In the search for a quick fix to establishing guidelines under which a new team will operate, many people go in search of existing CSIRT guidelines with the hope that they can simply be adopted for use in their environment. However, they soon realize that no single set of service definitions, policies, and procedures could be appropriate for any two CSIRTs. Moreover, teams with rigid guidelines in place find themselves struggling to adapt to the dynamic world of computer security incidents and attacks.

It is important to understand the inherent structure and needs of the environment in which the CSIRT will operate, and the posture that the CSIRT will take in relation to risk management

within that environment. With that understanding, the reader will be better positioned to apply this material to best suit that structure and set of requirements. Ultimately, of course, each team must define its own set of criteria and operating guidelines that supports its environment and constituency.

To obtain that goal in a structured fashion, it is best to start with and to recognize a basic framework for a CSIRT. That framework consists of the questions “what to do,” “for whom,” “in what local setting,” and “in cooperation with whom.” We capture this set of questions for the framework as identifying the

- mission statement: high-level goals, objectives, and priorities
- constituency: constituency type and relationship with the constituency
- place in organization: position within organizational structure and particularly within risk management
- relationship to others: setting of (inter)national CSIRT cooperation and coordination and other interactions

2.1.1 Mission Statement

Many CSIRTs in existence today either lack a clear understanding of their goals and objectives or have failed to effectively communicate that information to the parties with whom they interact. As a result, they needlessly expend effort and resources (often in crisis situations) in an attempt to

- understand if they are using the correct priorities to ensure they respond to the most important activity
- correct any inappropriate expectations of those they interact with
- understand how and whether it is appropriate for them to react to a given situation
- revise their policies and procedures to meet the needs of the situation
- determine if the range and nature of the services they offer should be modified

Until a CSIRT defines, documents, adheres to, and widely distributes a concise and clear mission statement, the situation is not likely to improve. Given the importance of the statement, it should be non-ambiguous and consist of at most three or four sentences specifying the mission with which the CSIRT is charged. The statement will help provide a basic understanding of what the team is trying to achieve; and more importantly, it will provide a focus for the overall goals and objectives of the CSIRT.

In addition, the mission statement of a CSIRT must have the backing and support of senior management in the parent organization (e.g., the corporate security officer, head of

information technology, board of directors, or equivalent). Without such backing the CSIRT will struggle to obtain recognition and resources.

A mission statement is imperative to enable the CSIRT to establish a service and quality framework, including the nature and range of services provided, the definition of its policies and procedures, and the quality of service. Together with the definition of the constituency, this service and quality framework (detailed in Section 2.2) drives and bounds all CSIRT activities. Clearly, if the team is housed within a larger organization or is funded from an external body, the CSIRT mission statement must complement the missions of those organizations.

Many CSIRTs additionally supply a purpose statement that supplements the mission and explains the reason(s) that resulted in the team being established. Armed with this information, the CSIRT should be in a good position to define its goals and appropriate services to support its mission. The public availability of these statements will facilitate the understanding of the CSIRT (its role, purpose, and the framework within which it operates) by other parties who will inevitably interact with the CSIRT during the course of its operation.

2.1.2 Constituency

During the course of its operation, every CSIRT will interact with a wide range of entities. The most important of these is the specific community that the CSIRT was established to serve: its constituency. A CSIRT constituency can be unbounded (the CSIRT will provide service to anyone requesting it), or it can be bound by some constraints. Most commonly, CSIRTs have bounded constituencies that tend to be a reflection of the CSIRT funding source.¹⁰ The most common constraints that are used to bound a constituency include national, geographical, political (e.g., government departments), technical (e.g., use of a specific operating system), organizational (e.g., within a given corporation or company), network service provider (e.g., connection to a specific network), or contractual (e.g., the customers of a fee-for-service team).

Table 1 provides examples for how different types of CSIRTs may fulfill differing missions and serve differing constituencies.

¹⁰ Kossakowski, Klaus-Peter. "The Funding Process: A Challenging Task." 6th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, Boston, Mass., July 1994.

Table 1: Examples of CSIRT Types With Associated Missions and Constituencies

CSIRT Type	Nature of Mission	Type of Constituency Served
International Coordination Center	Obtain a knowledge base with a global perspective of computer security threats through coordination with other CSIRTs. Building a “web of trust” among CSIRTs.	Other CSIRTs around the world
Corporation	Improve the security of the corporation’s information infrastructure and minimize threat of damage resulting from intrusions.	System and network administrators and system users within the corporation
Technical	Improve the security of a given IT product.	Users of the product

An essential task for the CSIRT is to define its constituency and its relationship to that constituency, and then go on to promote the CSIRT to the constituents and gain trust by “doing the job right.” The next few sections focus on some of these issues as they relate to aspects of the constituency and relationships across those boundaries.

2.1.2.1 Constituency Definition

A constituency might be defined in the form of a statement and may be supported by a list of domain names.

Example: The constituency of the Pennsylvania State University response team can be defined simply as “Pennsylvania State University” and as the domain “*.psu.edu”.

On the other hand, it can be difficult (or even impossible) for a different type of team to define its constituency in terms of domain names because its constituency may be very large and dynamic (changing as customers come and go).

Example: The constituency of AusCERT can be defined simply as “Subscribers of the AusCERT service,” but not everyone within the domain “*.au” will be a subscriber to AusCERT services. However, although AusCERT services are tailored towards their subscribing customers, in practice other external CSIRTs have used AusCERT as a point of contact for facilitating communication between Australian sites and other sites involved in computer security incidents.

Even when a constituency seems to be easy to define in the form of a single domain, there can be complications. In an academic environment (such as a university), student or faculty clubs, commercial spin-offs, or systems owned by research organizations might coexist on the university network. Such systems may or may not use the university domain name and may or may not fall under the CSIRT for the university.

Example: The CERT/CC is part of Carnegie Mellon University and is housed at the Software Engineering Institute; however, the CERT/CC manages its computing infrastructure separately from the university (and the Software Engineering Institute). Further, the CERT/CC is not the CSIRT for Carnegie Mellon University.

Depending on the range of services offered by a CSIRT and the nature of those services, a CSIRT may have the need to define more than one constituency. These multiple constituencies might intersect, be sub- or supersets, or be totally separate from other constituencies served by that CSIRT. For instance, a technical CSIRT might provide general security information on its specialized products to an unbounded constituency (e.g., the Internet) via a publicly available Web site, but may provide an enhanced range of services to only a subset of that constituency, such as the registered users of its products.

Even in the case of a CSIRT with a very bounded constituency, the CSIRT will most likely still have to deal with information associated with (or coming from) parties that do not belong to its constituency. For instance, a CSIRT providing an incident handling service to a bounded constituency will undoubtedly wish to accept incident reports that directly affect its constituency from parties outside of its constituency, appropriately handle that information, and ensure that it reaches the appropriate points of contact and is coordinated within its constituency. Many CSIRTs act as a coordination point between their constituency and other external parties (such as other CSIRTs, system administrators, vendors, law enforcement, legal counsel, and the media). These interactions can vary from simply relaying requests to complete sharing of data and full cooperation [Pethia 1990c]. It is important that a CSIRT decides, documents, and states how such interactions will be handled (see Section 3.7, “Interactions,” for more details on this topic).

In some cases CSIRTs specifically choose not to advertise their constituency. For instance, a network service provider CSIRT may consider its customer list to be proprietary information and so will not disclose the information. Similarly, a CSIRT that provides fee-for-service may have contractual agreements with its customers that prevent the CSIRT from disclosing its constituency. In such cases CSIRTs revert to describing their constituency in very generic terms such as “the customers of this organization.” This makes it hard or impossible for such CSIRTs to provide incident response coordination services to their constituency, as other external sites and teams do not know if a given constituent falls within a CSIRT’s constituency and so cannot report the activity directly to the appropriate team. What often occurs in these situations is that the customers are contacted directly by other sites or CSIRTs involved in an incident, and then as necessary the customers may (or may not) choose to seek support from their own CSIRT. It does point out how web-like trust relationships can evolve.

2.1.2.2 Overlapping Constituencies

Not all CSIRTs have unique constituencies. It is not uncommon for two or more CSIRTs to offer any given service to overlapping constituencies. However, experience has shown that

such situations of overlapping constituency will result in confusion between the CSIRTs and their constituencies unless all parties involved have a clear understanding of their responsibilities. There have been cases when CSIRTs with overlapping constituencies have not coordinated with each other appropriately, resulting in duplication of effort and antagonism between all concerned. Similarly, there have been situations when the constituents have not known from which CSIRT to seek support or assistance and as a result made duplicate or inappropriate reports.

Example: Consider a commercial company that has a contractual agreement with a fee-for-service CSIRT, and as a result falls into the constituency of the fee-for-service CSIRT. Additionally, as the result of being located in a given country, the company falls into the constituency of that country's national response team.

Example: German federal government organizations are connected to the German DFN-network for provision of communication and Internet access [Kossakowski 1994]. Many of these organizations fall into the constituencies of two teams:

- DFN-CERT, as a result of the organizations being connected to the DFN-network, and
- CERT-BUND, a team set up by and inside the German Information Security Agency (BSI [Bundesamt für Sicherheit in der Informationstechnik]) to address the specific needs of German federal government sites.

If any incidents affecting German federal government sites are reported, both teams will coordinate their activities as appropriate. While CERT-BUND will provide incident response support towards federal government sites, DFN-CERT might provide technical support in terms of analysis of vulnerabilities together with CERT-BUND. Both teams will discuss and coordinate any response affecting both teams.

Example: The CERT/CC has received (and on occasion, continues to receive) calls from individuals who are members of another CSIRT's constituency. In one case a system administrator at a university in the United Kingdom called the CERT/CC in the U.S. for assistance with an incident. The time in the United Kingdom was 9:00 a.m.; in the U.S. the local time was 3:00 a.m. A CERT/CC staff member was paged and provided immediate assistance to the site. The administrator in the U.K. (who was new to his job) was not aware of the services offered by JANET-CERT. Once he was informed of the existence of JANET-CERT and how it could offer him more appropriate service in terms of local needs and time zone, he contacted JANET-CERT directly. JANET-CERT followed up by providing the necessary support and advice, in fact being able to provide details related to the legal situation that the U.S.-based CERT/CC was not able to provide.

2.1.2.3 Relationship to Constituency

The nature of the relationship between a CSIRT and its constituency will directly affect the nature of the services that the CSIRT offers. As described in Table 2, those relationships fall into three general categories when considered in terms of the authority the CSIRT has over its constituents.

Table 2: Possible Authority Relationships Between a CSIRT and Its Constituency

Level of Authority	CSIRT/Constituency Relationship
Full	The members of the CSIRT have the authority to undertake any necessary actions or decisions on behalf of their constituency.
Shared	The members of the CSIRT provide direct support to their constituents and share in the decision-making process (i.e., have influence in constituency decisions, but are unable to dictate to them).
None	The members of the CSIRT have no authority over their constituency and can act only as advocates or advisors.

A fourth authority relationship, indirect authority, is possible but not common. In such a relationship, the CSIRT can exert pressure on its constituency to enforce sanctions if needed. The influence that a major network service provider (NSP) CSIRT may have over an Internet service provider (ISP) that it provides service to, or the influence that the ISP may have over its customers, are good examples of indirect authority.

Regardless of the authority relationship, some form of incident handling, or vulnerability analysis and response, or training services can be offered. However, services such as incident tracing and intrusion detection (listed in Table 4, “List of CSIRT Services”) may not be possible if the CSIRT has no authority over the constituency. In such cases some form of these services may be possible with contractual agreements in place to support them. But such agreements change the authority relationship to some extent.

Example: Take the situation where a CERT Advisory is released that announces an available patch for a security vulnerability in a widely used network daemon, exploitation of which results in a system compromise. Consider how CSIRTs with differing authority over their constituencies may react to such an announcement:

Full Authority

The CSIRT could *require* all constituents to disconnect from the network until they have installed the patch. Moreover, the CSIRT may manually intervene to disconnect those constituents that do not comply.

Shared Authority

The CSIRT could *advise* and *influence* constituents to disconnect from the network until

the patch has been installed. Additionally, it might *assist* the constituency by helping with coordination and response to the advice.

No Authority

The CSIRT can *advise* the constituency and *propagate information* to the constituency. In addition, the team can *try to motivate* the constituency to install the patch. However, the CSIRT cannot force the constituency to install the patch.

2.1.2.4 Promoting the CSIRT to the Constituency

Once the constituency has been defined, it will be important (regardless of the range and nature of the CSIRT services) to publicly advertise the constituency definition and the CSIRT services to ensure that both the constituency and other parties understand what interactions they might expect with the CSIRT. Particularly if a CSIRT intends to serve as a single point of contact for its constituency for computer security incident reports, it must ensure that it advertises its constituency to ensure that all concerned know to report incidents directly to the CSIRT rather than to an individual constituent. Similarly, constituents need to know which CSIRTs are offering them service and how to report to the appropriate CSIRT.

A team's constituency can be viewed in several ways:

- declared constituency: the constituency that the team claims or wishes to represent
- contractual constituency: the subset of the team's declared constituency who have a contractual agreement to report to the team (regardless of whether they make reports to the team or not)
- reporting constituency: the subset of the team's declared constituency that recognizes the team as representing it and as a result makes reports to it
- others: parties who fall outside the declared constituency of the team and require its services or make reports to it anyway. These might include those who do not know if they have a team that they can report to.

A CSIRT's goal should be to promote itself and its services as widely as possible to ensure that its declared constituency is aware of the team, ensure that other teams know of the CSIRT and the constituency it serves, and to gain broader recognition of the team in general. If the team does not effectively communicate its role and services it cannot expect to increase the size of its reporting constituency or its recognition within the broader CSIRT community.

A CSIRT should promote itself through as many communication channels as possible, including the use of

- constituency email lists and newsgroups
- CSIRT or organizational information/Web server

- presentation, workshop, and tutorial materials
- general awareness materials and newsletters (both regular and “flash”)
- the media (who can reach those portions of the constituency or management levels that do not tend to use email, Web, or other online communication methods)

2.1.2.5 Gaining Constituency Trust

Regardless of the CSIRT’s (authority) relationship with its constituency, it must do more than simply define and publicize the constituency that it claims to serve. It cannot operate effectively without gaining and maintaining the constituency’s trust and respect. Even if a CSIRT has total authority over its constituency, it does not mean that the constituency’s trust and respect can be assumed in such a relationship. This trust must be earned and nurtured. As the team gains the trust and respect of its declared constituency, more of the declared constituency will begin to recognize and support the team, resulting in the growth of the team’s reporting constituency. Experience indicates that it takes about a year from the time that a team commences operations and announces its declared constituency before a stable reporting constituency is established.

Regardless of the constituency defined by a CSIRT, it is rare for any team to achieve 100% recognition by its constituency. It is useful to keep this in mind when trying to predict the impact that a team can have over its declared constituency. No matter how hard a team may try to reach out to its constituency and offer help or influence, it is unlikely that all of the constituency will respond.

2.1.3 Place in Organization

In the basic framework for a CSIRT, one needs not only to state what the team aims to do (mission statement) and for whom (constituency), but also to properly define the “roots” of the CSIRT: its place in its parent organization. This is not just a matter of administrative definition—were it only that, this section would not be necessary.

The place that a CSIRT holds in its parent organization is tightly coupled to its stated mission—and to a lesser degree, its constituency. This is best demonstrated by considering the extreme example of a CSIRT with a very highly visible supportive mission for a Fortune 500 constituency. If placed under the system administration department of its parent organization (a clear mismatch in responsibility), it is destined to fail. To help avoid such pitfalls, relevant aspects of a CSIRT’s position within its parent organization are discussed in this section.

A CSIRT may constitute the entire security team for an organization or may be totally distinct from an organization's security team.¹¹ Alternatively, although an organization may not have a distinct CSIRT, this role may in fact be served implicitly by the organization's security team. Regardless of the implementation, provision of the incident handling service is the key issue. For the purposes of this document, we will consider a CSIRT in its most common and simplistic form, as part (from a small to total overlap) of a larger security team housed within a parent organization, as shown in Figure 1.

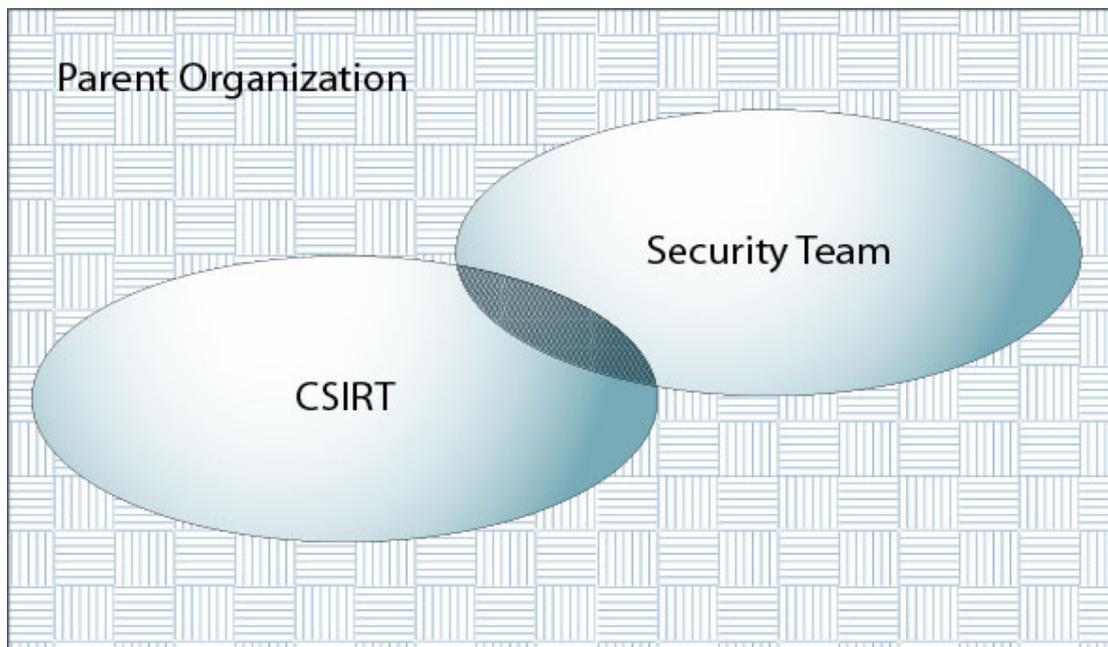


Figure 1: CSIRT Within an Organization

In a corporate environment, a CSIRT must be well embedded within the organization's business structure and commonly resides within, or has some overlap, with the organization's IT security department.

It is also possible for multiple incident handling capabilities to exist within a single parent organization. Such situations arise in vendor organizations and network service providers that may have two separate teams: one to handle incidents involving the company's own network, and another providing services to customers. Vendor organizations may also provide additional services such as those related to addressing security flaws in their products. Multiple capabilities might also arise in a single organization that does not provide services

¹¹ A security team is defined as system, network, and security administrators usually located within an IT department whose job functions involve internal and external security defenses. For example, they handle security issues and technologies such as firewalls, anti-virus filters, secure remote access, and intrusion detection. The term "security team" can refer to individuals who perform these functions or to a group of individuals who work as a team. These individuals might be located in a centralized site, but more often are distributed across the enterprise.

to external parties. For example, a company might choose to handle incidents related to or caused by malicious software through one team and use a different team for network attacks and intrusions.

Before a CSIRT can begin to establish its operational guidelines, it is important to determine the role that the CSIRT plays in overall risk management in the context of its organizational environment and constituency. This role will vary depending on the nature of the parent organization and the nature of the constituency that the team serves. Whatever the resulting role, it is imperative that it is supported by management and understood by all parties involved.

The parts of the organization that host the computing, networking, and communications equipment (on which data resides) clearly carry the technical risk. The business risk also needs to be considered, and that risk may be carried by many different parts of the organization. However, it is important to understand where the responsibility for managing risk resides, and how each part of the organization involved in this area interacts and coordinates their responsibilities.

In a commercial organizational setting, different groups in the same organization may have the responsibility for different aspects of risk management.

Examples: The network operations team responsible for network security issues; the system administrators responsible for host security issues; the physical security team responsible for access to buildings and facilities; the CSIRT responsible for coordination of response to any computer security incident reports; and corporate security responsible for setting company-wide policies and procedures including all other security-related teams and personnel.

Regardless of their specific role in risk management, each group needs to understand how its responsibilities inter-relate to the other organizational components and how to coordinate with other groups to ensure that it does not operate in isolation (or contradict any other group's operations). This includes providing a clear description of each group's duties, interaction/escalation points, and shared responsibilities.

Similarly, an organization may call upon the services of an external CSIRT. If so, the responsibilities and operations of the external CSIRT must be included and equally well defined in the organization's risk management framework.

2.1.4 Relationship to Other Teams

The realm of CSIRTs is the Internet, and therefore the world. There are many constituencies around the world, and a growing number are served by a CSIRT. In this regard, at some level

these CSIRTs have to inter-operate in order to get their job done. This cooperation and coordination effort is at the very heart of the CSIRT framework: just stating the mission, defining constituency, and determining the CSIRT's place within organization are not sufficient without also covering the coordination issue.

Within the context of CSIRTs as they exist today, there is some hierarchical structure that can be observed between the different types of teams. There are teams providing service to clearly marked constituencies and other teams who serve a coordination role across groups (commonly national or international) of CSIRTs. However, this structure is not a true hierarchy, and in most cases the structure is both informal and voluntary. This informal structure is seen as a benefit, as it allows teams the flexibility to share information quickly and effectively with other CSIRTs that they trust and to be more cautious with other teams with which there has not been as much opportunity to determine trustworthiness.

Some formal hierarchies do exist, such as within the U.S. military. For example, the U.S. Army, Air Force, and Navy (ACERT/CC,¹² AFCERT, and NAVCIRT, respectively), serve their own constituencies; while the U.S. Department of Defense DOD-CERT coordinates across all the U.S. military teams.

Note that, for some types of activity, many teams choose to interact directly with other peer teams and not interact at all with a coordinating CSIRT. This commonly happens when the teams involved see no need to bring a coordinating CSIRT into the loop to address a specific problem. However, coordinating CSIRTs usually request that they be informed of all activity in order for them to obtain an overall view of the level of activity in their domain and alert other teams to look for additional or related activity.

As depicted in Figure 2, there are various types of possible peer relationships between CSIRTs. A team may be considered as a coordinating CSIRT if it plays a coordination role among other CSIRTs. The example in Figure 2 depicts both CSIRTs A and B as coordinating CSIRTs. In addition to coordinating among CSIRTs C and D, CSIRT B has another constituency component that is not covered by C or D and is served directly by B. On the other hand, CSIRT A has a constituency that is solely made up of other CSIRTs (B, E, F, and G). However, the CSIRTs in A's constituency do not fall into a rigid hierarchy because there is communication between CSIRTs E and F that can occur independent of their communications with CSIRT A.

The relationships discussed in this section can be used to depict any CSIRT regardless of its setting or purpose. For instance, CSIRTs such as international coordination centers (e.g., CERT/CC), national response teams (e.g., DK-CERT, JPCERT/CC), fee-for-service response

¹² The Army CERT, for example, plays a coordinating role itself across geographically dispersed regional Army teams (each called an "RCERT").

teams (e.g., dCERT, IBM/MSS), teams for commercial organizations (e.g., Motorola’s MCERT, Boeing’s BCERT), network service provider teams (e.g., UNI-CERT, BT-CERT/CC), and universities (e.g., Pennsylvania State University’s PSU-CERT, Stanford University’s SUNSeT) can all be represented using this approach.

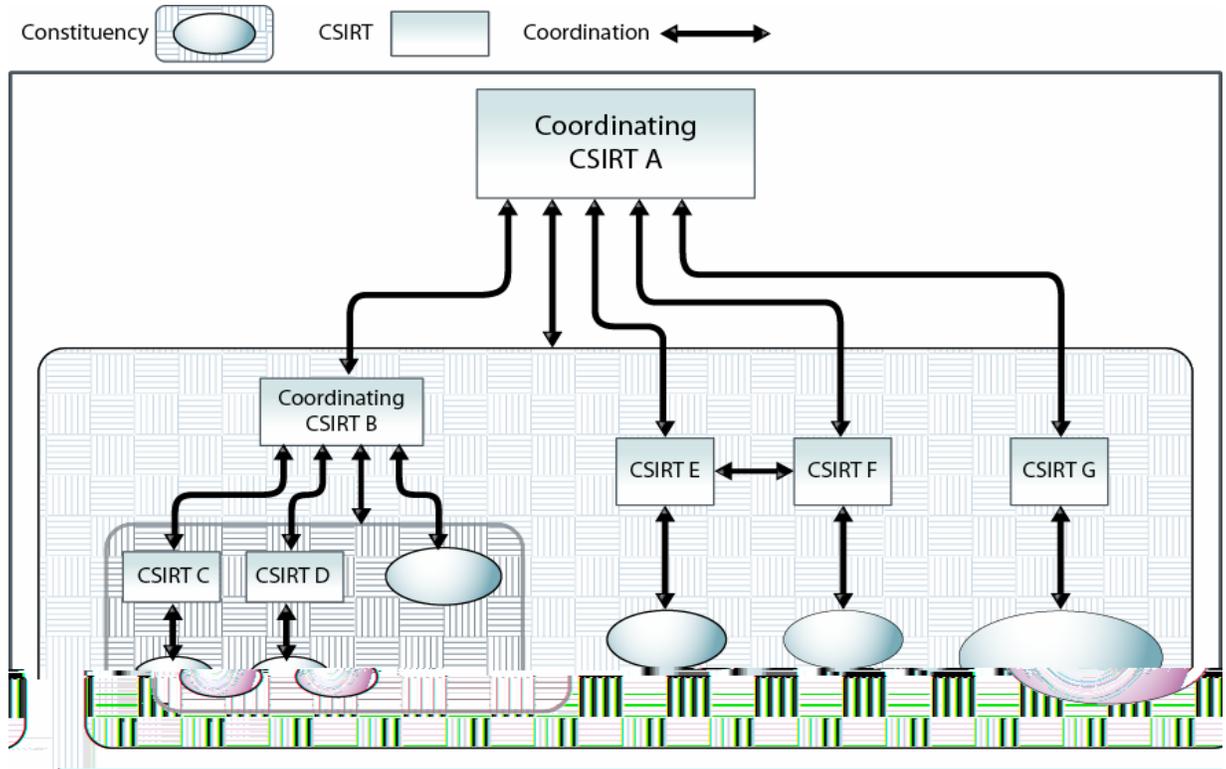


Figure 2: CSIRT Peer Relationships

2.2 Service and Quality Framework

The mission statement of a CSIRT essentially has three derivatives—services, policies, and quality—each of which needs to embody the scope and purpose of the mission statement. The services offered by a team are the methods used to carry out the team’s mission. Services are usually provided to the team’s constituency. Policies are the governing principles under which the team operates. Quality is the desired standard at which all activities will be undertaken. The information flowing within a CSIRT permeates all of the mission statement derivatives. Governed by services, policies, and quality, procedures specify how activities are enacted. This framework is depicted in Figure 3.

Following this framework, the three derivatives of the mission statement (services, policies, and quality) will be discussed in more detail. Although information flow might naturally be a fourth topic to discuss, it will not be covered in this section. For the purposes of this handbook, we will mainly concentrate on the flow of information related to any external

party. The internal flow of information (e.g., between team members, between different services provided) is nevertheless important for the team to define and improve.

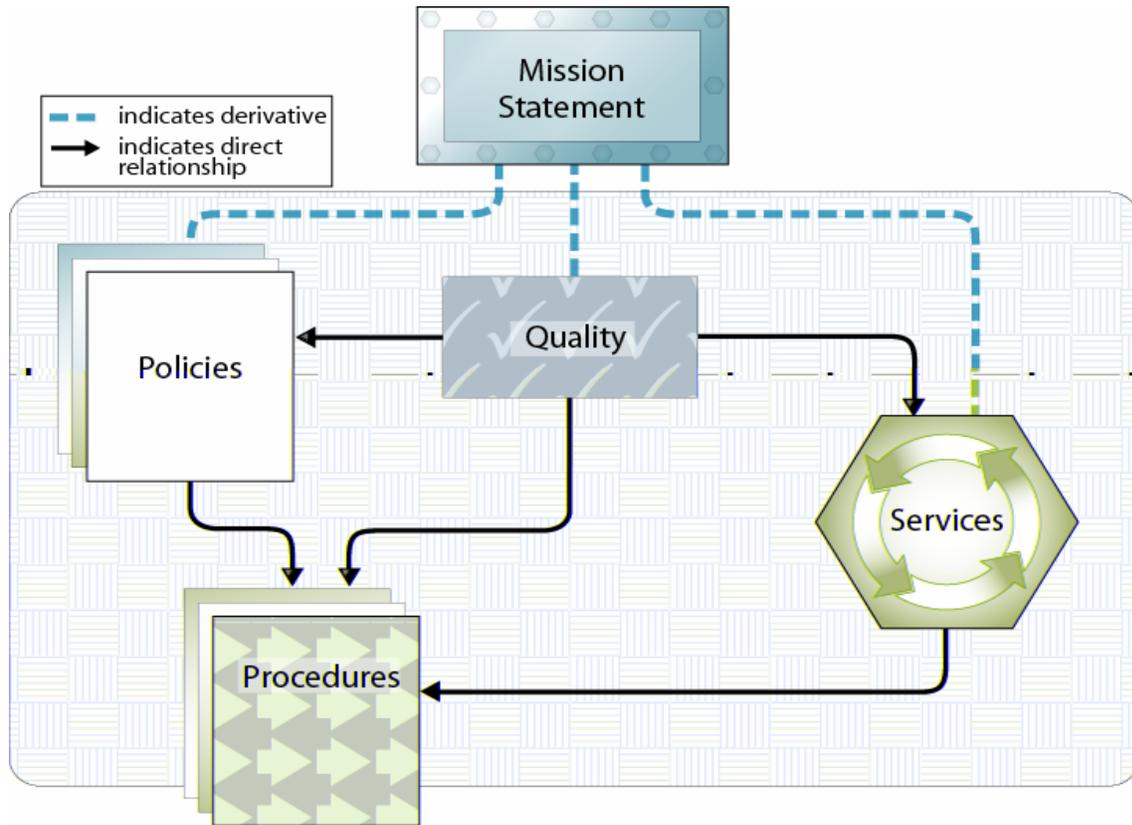


Figure 3: Service and Quality Framework as Derived from Mission Statement

We consider information flow of basic interest only where it pertains to external communication, and as a result, we view the flow of information (e.g., information flow inside the team) as clearly not a basic CSIRT issue. The topic of internal information flow will be discussed in relation to services in Section 2.4, “Information Flow,” and wherever relevant in the subsequent treatment of CSIRT issues in Chapters 3 and 4.

A CSIRT can expect to offer a range of different services to its constituency that directly reflects the inherent promise of the CSIRT mission statement. The incident handling service, which is the focus of this document, will be described in detail in Chapter 3. To provide the necessary context for the discussion of the service, however, this section introduces issues that are generic to all CSIRT services and provides a brief discussion of other common services that a number of CSIRTs offer.

For each service provided, the CSIRT should provide its constituency with service descriptions (or formal service level agreements) in as much detail as possible. In particular,

any service provided by the CSIRT should include an explanation of the attributes and descriptions as outlined in Table 3.

Table 3: Service Description Attributes

Attribute	Description
Objective	Purpose and nature of the service.
Definition	Description of scope and depth of service.
Function Descriptions	Descriptions of individual functions within the service.
Availability	The conditions under which the service is available: to whom, when, and how.
Quality Assurance	Quality assurance parameters applicable for the service. Includes both setting and limiting of constituency expectations.
Interactions and Information Disclosure	The interactions between the CSIRT and parties affected by the service, such as the constituency, other teams, and the media. Includes setting information requirements for parties accessing the service, and defining the strategy with regard to the disclosure of information (both restricted and public).
Interfaces with Other Services	Define and specify the information flow exchange points between this service and other CSIRT services it interacts with.
Priority	The relative priorities of functions within the service, and of the service versus other CSIRT services.

These descriptions are helpful to the team when defining, implementing, and operating the service. Similarly they provide information that should be made available (in some form) to the constituency to both advertise and set the appropriate expectations for the service. Since the nature of the field is one of constant change, reprioritization, and technical advancement, a CSIRT will need to frequently reassess the nature and levels of service it provides to keep pace with the changing environment and the resources available to it. Likewise, the constituency must be informed of any noticeable changes.

2.3 CSIRT Services

For a team to be considered a CSIRT, it must provide one or more of the incident handling services: incident analysis, incident response on site, incident response support, or incident response coordination. As we mentioned earlier in this handbook, the incident handling service includes incident analysis with at least one of the other incident handling services: incident response resolution, incident response support, or incident response coordination (see below for detailed explanations of the differences). In practice, we see that CSIRTs commonly offer other services in addition to the basic incident handling service¹³, depending on the needs of its constituency. These additional services might be provided by the CSIRT alone or in cooperation with other organizational units (such as the IT or security department).

¹³ Such other services might include, for example, distributing advisories, alerts and warnings, vulnerability handling, other proactive announcements, and/or training or awareness building within their constituency.

In addition to the mandatory incident handling service, Table 4 lists some of the most common services that CSIRTs provide and the form that those services might take. Within this set of additional services, a few (such as announcements or vulnerability analysis and response) are even more likely to be offered, as they are closely associated with incident handling activities.

Although this description focuses on services provided by CSIRTs, many of these same services can also be provided by system, network, and security administrators who perform ad hoc incident handling as part of their normal administrative work. Sometimes such ad hoc teams are referred to as a “security teams or other security-related groups.”

2.3.1 Service Categories

There are many services that a CSIRT can choose to offer. The services that each CSIRT provides should be based on the mission, purpose, and constituency of the team.

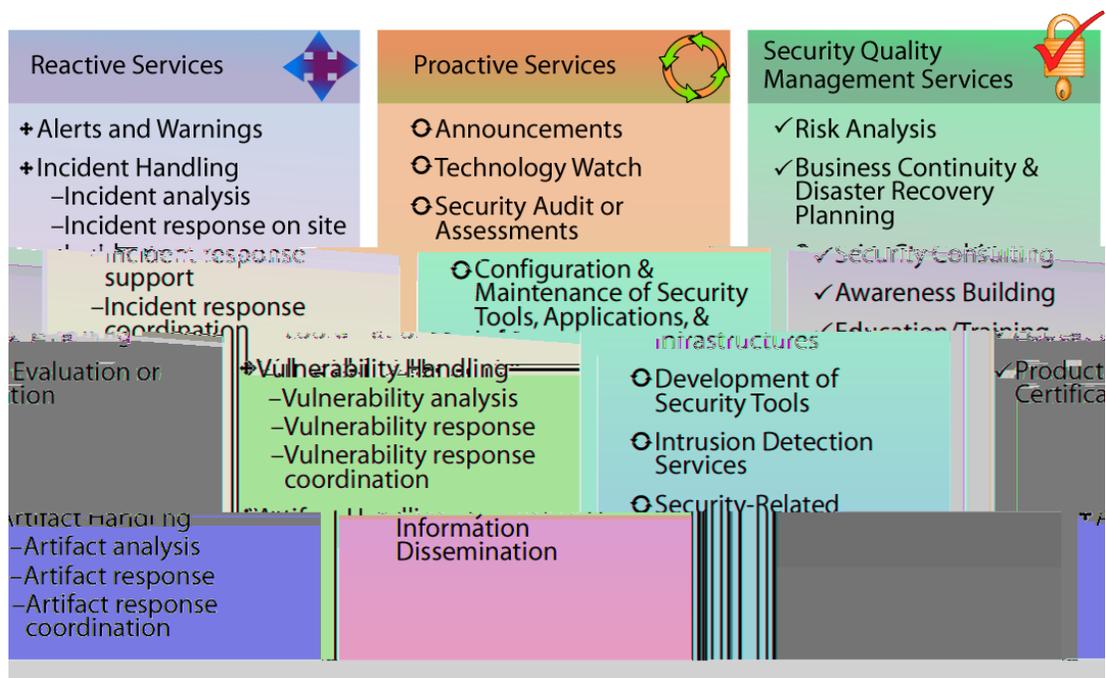
CSIRT services can be grouped into three categories:

- **Reactive services.** These services are triggered by an event or request, such as a report of a compromised host, wide-spreading malicious code, software vulnerability, or something that was identified by an intrusion detection or logging system. Reactive services are the core component of CSIRT work.
- **Proactive services.** These services provide assistance and information to help prepare, protect, and secure constituent systems in anticipation of attacks, problems, or events. Performance of these services will directly reduce the number of incidents in the future.
- **Security quality management services.** These services augment existing and well-established services that are independent of incident handling and traditionally performed by other areas of an organization such as the IT, audit, or training departments. If the CSIRT performs or assists with these services, the CSIRT’s point of view and expertise can provide insight to help improve the overall security of the organization and identify risks, threats, and system weaknesses. These services are generally proactive but contribute indirectly to reducing the number of incidents.

The services corresponding to each category are listed in Table 4 and described in detail below.

It should be noted that some services have both a reactive and proactive side. For example, vulnerability handling can be done in response to the discovery of a software vulnerability that is being actively exploited. But it can also be done proactively by reviewing and testing code to determine where vulnerabilities exist, so the problems can be fixed before they are widely known or exploited.

Table 4: List of Common CSIRT Services



2.3.2 Service Descriptions

2.3.2.1 Reactive Services

Reactive services are designed to respond to requests for assistance, reports of incidents from the CSIRT constituency, and any threats or attacks against CSIRT systems. Some services may be initiated by third-party notification or by viewing monitoring or intrusion detection system (IDS) logs and alerts.

Alerts and Warnings

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 CSIRT or may be redistributed from vendors, other CSIRTs or security experts, or other parts of the constituency.

Incident Handling

Incident handling involves receiving, triaging,¹⁴ and responding to requests and reports, and analyzing incidents and events. Particular response activities can include

- taking action to protect systems and networks affected or threatened by intruder activity
- providing solutions and mitigation strategies from relevant advisories or alerts
- looking for intruder activity on other parts of the network
- filtering network traffic
- rebuilding systems
- patching or repairing systems
- developing other response or workaround strategies

Since incident handling activities are implemented in various ways by different types of CSIRTs, this service is further categorized based on the type of activities performed and the type of assistance given as follows:

Incident analysis. There are many levels of incident analysis and many sub-services. Essentially, incident analysis is an examination of all available information and supporting evidence or artifacts related to an incident or event. The purpose of the analysis is to identify the scope of the incident, the extent of damage caused by the incident, the nature of the incident, and available response strategies or workarounds. The CSIRT may use the results of vulnerability and artifact analysis (described below) to understand and provide the most complete and up-to-date analysis of what has happened on a specific system. The CSIRT correlates activity across incidents to determine any interrelations, trends, patterns, or intruder signatures. Two sub-services that may be done as part of incident analysis, depending on the mission, goals, and processes of the CSIRT, are

- **Forensic evidence collection:** the collection, preservation, documentation, and analysis of evidence from a compromised computer system to determine changes to the system and to assist in the reconstruction of events leading to the compromise. This gathering of information and evidence must be done in a way that documents a provable chain of custody that is admissible in a court of law under the rules of evidence. Tasks involved in forensic evidence collection include (but are not limited to) making a bit-image copy of the affected system's hard drive; checking for changes to the system such as new programs, files, services, and users; looking at running processes and open ports; and checking for Trojan horse programs and toolkits. CSIRT staff performing this function may also have to be prepared to act as expert witnesses in court proceedings.

¹⁴ Triaging refers to the sorting, categorizing, and prioritizing of incoming incident reports or other CSIRT requests. It can be compared to triage in a hospital, where patients who need to be seen immediately are separated from those who can wait for assistance.

- **Tracking or tracing:** the tracing of the origins of an intruder or identifying systems to which the intruder had access. This activity might involve tracking or tracing how the intruder entered the affected systems and related networks, which systems were used to gain access, where the attack originated, and what other systems and networks were used as part of the attack. It might also involve trying to determine the identity of the intruder. This work might be done alone but usually involves working with law enforcement personnel, Internet service providers, or other involved organizations.

Incident response¹⁵ on site. The CSIRT provides direct, on-site assistance to help constituents recover from an incident. The CSIRT itself physically analyzes the affected systems and conducts the repair and recovery of the systems, instead of only providing incident response support by telephone or email (see below). This service involves all actions taken on a local level that are necessary if an incident is suspected or occurs. If the CSIRT is not located at the affected site, team members would travel to the site and perform the response. In other cases a local team may already be on site, providing incident response as part of its routine work. This is especially true if incident handling is provided as part of the normal job function of system, network, or security administrators in lieu of an established CSIRT.

Incident response support. The CSIRT assists and guides the victim(s) of the attack in recovering from an incident via phone, email, fax, or documentation. This can involve technical assistance in the interpretation of data collected, providing contact information, or relaying guidance on mitigation and recovery strategies. It does not involve direct, on-site incident response actions as described above. The CSIRT instead provides guidance remotely so site personnel can perform the recovery themselves.

Incident response coordination. The CSIRT coordinates the response effort among parties involved in the incident. This usually includes the victim of the attack, other sites involved in the attack, and any sites requiring assistance in the analysis of the attack. It may also include the parties that provide IT support to the victim, such as Internet service providers, other CSIRTs, and system and network administrators at the site. The coordination work may involve collecting contact information, notifying sites of their potential involvement (as victim or source of an attack), collecting statistics about the number of sites involved, and facilitating information exchange and analysis. Part of the coordination work may involve notification and collaboration with an organization's legal counsel, human resources or public relations departments. It would also include coordination with law enforcement. This service does not involve direct, on-site incident response.

¹⁵ Note that "incident response" is used here to describe one type of CSIRT service. When used in team names such as "Incident Response Team," the term typically has the broader meaning of incident handling.

Vulnerability Handling

Vulnerability handling involves receiving information and reports about hardware and software vulnerabilities;¹⁶ analyzing the nature, mechanics, and effects of the vulnerabilities; and developing response strategies for detecting and repairing the vulnerabilities. Since vulnerability handling activities are implemented in various ways by different types of CSIRTs, this service is further categorized based on the type of activities performed and the type of assistance given as follows:

Vulnerability analysis. The CSIRT performs technical analysis and examination of vulnerabilities in hardware or software. This includes the verification of suspected vulnerabilities and the technical examination of the hardware or software vulnerability to determine where it is located and how it can be exploited. The analysis may include reviewing source code, using a debugger to determine where the vulnerability occurs, or trying to reproduce the problem on a test system.

Vulnerability response. This service involves determining the appropriate response to mitigate or repair a vulnerability. This may involve developing or researching patches, fixes, and workarounds. It also involves notifying others of the mitigation strategy, possibly by creating and distributing advisories or alerts.¹⁷ This service can include performing the response by installing patches, fixes, or workarounds.

Vulnerability response coordination. The CSIRT notifies the various parts of the enterprise or constituency about the vulnerability and shares information about how to fix or mitigate the vulnerability. The CSIRT verifies that the vulnerability response strategy has been successfully implemented. This service can involve communicating with vendors, other CSIRTs, technical experts, constituent members, and the individuals or groups who initially discovered or reported the vulnerability. Activities include facilitating the analysis of a vulnerability or vulnerability report; coordinating the release schedules of corresponding documents, patches, or workarounds; and synthesizing technical analysis done by different parties. This service can also include maintaining a public or private archive or knowledge base of vulnerability information and corresponding response strategies.

Artifact Handling

An artifact is any file or object found on a system that might be involved in probing or attacking systems and networks or that is being used to defeat security measures. Artifacts can include but are not limited to computer viruses, Trojan horse programs, worms, exploit scripts, and toolkits.

¹⁶ A vulnerability is the existence of a flaw or weakness in hardware or software that can be exploited resulting in a violation of an implicit or explicit security policy.

¹⁷ Other CSIRTs might further redistribute these original advisories or alerts as part of their services.

Artifact handling involves receiving information about and copies of artifacts that are used in intruder attacks, reconnaissance, and other unauthorized or disruptive activities. Once received, the artifact is reviewed. This includes analyzing the nature, mechanics, version, and use of the artifacts; and developing (or suggesting) response strategies for detecting, removing, and defending against these artifacts. Since artifact handling activities are implemented in various ways by different types of CSIRTs, this service is further categorized based on the type of activities performed and the type of assistance given as follows:

Artifact analysis. The CSIRT performs a technical examination and analysis of any artifact found on a system. The analysis done might include identifying the file type and structure of the artifact, comparing a new artifact against existing artifacts or other versions of the same artifact to see similarities and differences, or reverse engineering or disassembling code to determine the purpose and function of the artifact.

Artifact response. This service involves determining the appropriate actions to detect and remove artifacts from a system, as well as actions to prevent artifacts from being installed. This may involve creating signatures that can be added to antivirus software or IDS.

Artifact response coordination. This service involves sharing and synthesizing analysis results and response strategies pertaining to an artifact with other researchers, CSIRTs, vendors, and other security experts. Activities include notifying others and synthesizing technical analysis from a variety of sources. Activities can also include maintaining a public or constituent archive of known artifacts and their impact and corresponding response strategies.

2.3.2.2 Proactive Services

Proactive services are designed to improve the infrastructure and security processes of the constituency before any incident or event occurs or is detected. The main goals are to avoid incidents and to reduce their impact and scope when they do occur.

Announcements

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.

Technology Watch

The CSIRT monitors and observes new technical developments, intruder activities, and related trends to help identify future threats. Topics reviewed can be expanded to include

legal and legislative rulings, social or political threats, and emerging technologies. This service involves reading security mailing lists, security Web sites, and current news and journal articles in the fields of science, technology, politics, and government to extract information relevant to the security of the constituent systems and networks. This can include communicating with other parties that are authorities in these fields to ensure that the best and most accurate information or interpretation is obtained. The outcome of this service might be some type of announcement, guidelines, or recommendations focused at more medium- to long-term security issues.

Security Audits or Assessments

This service provides a detailed review and analysis of an organization's security infrastructure, based on the requirements defined by the organization or by other industry standards that apply.¹⁸ It can also involve a review of the organizational security practices. There are many different types of audits or assessments that can be provided, including

- infrastructure review—manually reviewing the hardware and software configurations, routers, firewalls, servers, and desktop devices to ensure that they match the organizational or industry best practice security policies and standard configurations
- best practice review—interviewing employees and system and network administrators to determine if their security practices match the defined organizational security policy or some specific industry standards
- scanning—using vulnerability or virus scanners to determine which systems and networks are vulnerable
- penetration testing—testing the security of a site by purposefully attacking its systems and networks

Obtaining upper management approval is required before conducting such audits or assessments. Some of these approaches may be prohibited by organizational policy. Providing this service can include developing a common set of practices against which the tests or assessments are conducted, along with developing a required skill set or certification requirements for staff that perform the testing, assessments, audits, or reviews. This service could also be outsourced to a third party contractor or managed security service provider with the appropriate expertise in conducting audits and assessments.

¹⁸ Industry standards and methodologies might include Operationally Critical Threat, Asset, and Vulnerability EvaluationSM (OCTAVESM), CCTA Risk Analysis and Management Method (CRAMM), Information Security Forum's Fundamental Information Risk Management (FIRM), Commonly Accepted Security Practices and Regulations (CASPR), Control Objectives for Information and (Related) Technology (COBIT), Methode d' Evaluation de la Vulnerabilite Residuelle des Systemes d'Informa (MELISA), ISO 13335, ISO 17799, or ISO 15408.

Configuration and Maintenance of Security Tools, Applications, Infrastructures, and Services

This service identifies or provides appropriate guidance on how to securely configure and maintain tools, applications, and the general computing infrastructure used by the CSIRT constituency or the CSIRT itself. Besides providing guidance, the CSIRT may perform configuration updates and maintenance of security tools and services, such as IDS, network scanning or monitoring systems, filters, wrappers, firewalls, virtual private networks (VPN), or authentication mechanisms. The CSIRT may even provide these services as part of their main function. The CSIRT may also configure and maintain servers, desktops, laptops, personal digital assistants (PDAs), and other wireless devices according to security guidelines. This service includes escalating to management any issues or problems with configurations or the use of tools and applications that the CSIRT believes might leave a system vulnerable to attack.

Development of Security Tools

This service includes the development of any new, constituent-specific tools that are required or desired by the constituency or by the CSIRT itself. This can include, for example, developing security patches for customized software used by the constituency or secured software distributions that can be used to rebuild compromised hosts. It can also include developing tools or scripts that extend the functionality of existing security tools, such as a new plug-in for a vulnerability or network scanner, scripts that facilitate the use of encryption technology, or automated patch distribution mechanisms.

Intrusion Detection Services

CSIRTs that perform this service review existing IDS logs, analyze and initiate a response for any events that meet their defined threshold, or forward any alerts according to a pre-defined service level agreement or escalation strategy. Intrusion detection and analysis of the associated security logs can be a daunting task—not only in determining where to locate the sensors in the environment, but collecting and then analyzing the large amounts of data captured. In many cases, specialized tools or expertise is required to synthesize and interpret the information to identify false alarms, attacks, or network events and to implement strategies to eliminate or minimize such events. Some organizations choose to outsource this activity to others who have more expertise in performing these services, such as managed security service providers.

Security-Related Information Dissemination

This service provides constituents with a comprehensive and easy-to-find collection of useful information that aids in improving security [Kossakowski 2000]. Such information might include

- reporting guidelines and contact information for the CSIRT
- archives of alerts, warnings, and other announcements
- documentation about current best practices
- general computer security guidance
- policies, procedures, and checklists
- patch development and distribution information
- vendor links
- current statistics and trends in incident reporting
- other information that can improve overall security practices

This information can be developed and published by the CSIRT or by another part of the organization (IT, human resources, or media relations), and can include information from external resources such as other CSIRTs, vendors, and security experts.

2.3.2.3 Security Quality Management Services

Services that fall into this category are not unique to incident handling or CSIRTs in particular. They are well-known, established services designed to improve the overall security of an organization. By leveraging the experiences gained in providing the reactive and proactive services described above, a CSIRT can bring unique perspectives to these quality management services that might not otherwise be available. These services are designed to incorporate feedback and lessons learned based on knowledge gained by responding to incidents, vulnerabilities, and attacks. Feeding such experiences into the established traditional services (described below) as part of a security quality management process can improve the long-term security efforts in an organization.

Depending on organizational structures and responsibilities, a CSIRT may provide these services or participate as part of a larger organizational team effort.

The following descriptions explain how CSIRT expertise can benefit each of these security quality management services.

Risk Analysis

CSIRTs may be able to add value to risk analysis and assessments. This can improve the organization's ability to assess real threats, provide realistic qualitative and quantitative assessments of the risks to information assets, and evaluate protection and response strategies. CSIRTs performing this service would conduct or assist with information security risk analysis activities for new systems and business processes or evaluate threats and attacks against constituent assets and systems.

Business Continuity and Disaster Recovery Planning

Based on past occurrences and future predictions of emerging incident or security trends, more and more incidents have the potential to result in serious degradation of business operations. Therefore, planning efforts should consider CSIRT experience and recommendations in determining how best to respond to such incidents to ensure the continuity of business operations. CSIRTs performing this service are involved in business continuity and disaster recovery planning for events related to computer security threats and attacks.

Security Consulting

CSIRTs can be used to provide advice and guidance on the best security practices to implement for constituents' business operations. A CSIRT providing this service is involved in preparing recommendations or identifying requirements for purchasing, installing, or securing new systems, network devices, software applications, or enterprise-wide business processes. This service includes providing guidance and assistance in developing organizational or constituency security policies. It can also involve providing testimony or advice to legislative or other government bodies.

Awareness Building

CSIRTs may be able to identify where constituents require more information and guidance to better conform to accepted security practices and organizational security policies. Increasing the general security awareness of the constituent population not only improves their understanding of security issues but also helps them perform their day-to-day operations in a more secure manner. This can reduce the occurrence of successful attacks and increase the probability that constituents will detect and report attacks, thereby decreasing recovery times and eliminating or minimizing losses.

CSIRTs performing this service seek opportunities to increase security awareness through developing articles, posters, newsletters, Web sites, or other informational resources that explain security best practices and provide advice on precautions to take. Activities may also include scheduling meetings and seminars to keep constituents up to date with ongoing security procedures and potential threats to organizational systems.

Education/Training

This service involves providing information to constituents about computer security issues through seminars, workshops, courses, and tutorials. Topics might include incident reporting guidelines [CERT/CC 1998a], appropriate response methods, incident response tools, incident prevention methods, and other information necessary to protect, detect, report, and respond to computer security incidents.

Product Evaluation or Certification

For this service, the CSIRT may conduct product evaluations on tools, applications, or other services to ensure the security of the products and their conformance to acceptable CSIRT or organizational security practices. Tools and applications reviewed can be open source or commercial products. This service can be provided as an evaluation or through a certification program, depending on the standards that are applied by the organization or by the CSIRT.

As a summary of the list of services described in the above section, our experience and discussions with others has shown that whatever services a CSIRT chooses to offer, the parent organization or management must ensure that the team has the necessary resources (people, technical expertise, equipment, and infrastructure) to provide a valued service to their constituents, or the CSIRT will not be successful and their constituents will not report incidents to them.¹⁹

2.3.3 Selection of Services

A CSIRT must take great care in choosing the services it will offer. The set of services provided will establish the resources, skill sets, and partnerships the team will need to function properly. The selection of services should first and foremost support and enable the business goals of the CSIRT's constituency or parent organization. The services provided should be those that the team can realistically and honestly provide based on the team size and range of expertise and skills.

To a large extent, the success of the CSIRT will depend on the overall quality of the services it provides to its constituency. It is better to offer a few services well than a large range of services poorly. As a CSIRT gains the trust and respect of its constituency, it can look to expand its services as staff and funding permit.

When deciding the range and nature of services to provide, care should be taken to ensure that the service selection supports and complements the overall CSIRT mission. In reality

¹⁹ If the CSIRT does not provide the services but outsources the activities to another organization such as a managed security services provider, the same standards for staffing, equipment, and infrastructure still are just as important and must be adhered to, in order to protect the CSIRT, organizational data, and services.

many teams offer a limited set of services but their constituencies insist on adaptations or additional services. If these additional demands are made from influential constituency members and the CSIRT is lacking high-level management support, the tendency is to provide some element of support for these services even if they fall outside of the team's official charter.

Using the attributes and definitions in Table 3, any other additional services can be described in a similar fashion to the incident handling service presented in Chapter 3.

2.4 Information Flow

Whatever services are offered by a CSIRT, it is important to understand which of those services are in some way related to each other and what the interdependencies are. In particular, it is necessary to specify the interfaces between the services and any associated information flow between the services. It is important to identify which services

- rely on information from, or provide information to, another service
- are responsible for providing/requesting the information to/from another service
- have a shared need for a specific function or a specific set of information
- transfer information-dependent responsibilities (e.g., for confidentiality, appropriate use) to another service or externally (other CSIRTs, constituency)

Using this information, it may be possible to optimize the use of resources, to avoid duplication of effort and make efficient use of pre-existing information. For example, all incoming requests could be handled by a centralized help desk that directs (or “triages”) the requests to the appropriate service. For other services handling incoming requests directly, care must be taken to ensure that any information flow linkages with other (related) CSIRT services are appropriately identified and shared to avoid unnecessary duplication of effort.

Example: Consider a CSIRT that handles non-virus related incident reports and a separate department that handles any virus-related activity. Say a constituent contacts the CSIRT to report a compromise on a system that involves both modifications to the system (e.g., additional user accounts added, system changes, defaced Web-pages) and indications of virus-related activity (known viruses or worm programs installed on the compromised machine). At the same time, the constituent also notifies the CSIRT that there have been recent changes in personnel responsible for the affected system. The constituent provides the latest contact information for the new personnel. In this situation, the CSIRT has information that the virus-handling department should be aware of. After following their procedures for verifying information and requesting permission to pass the information on to the other department, the CSIRT passes the relevant information to the virus-handling department. This eliminates the need for the constituent to separately contact the virus-handling department to notify them of activity. It also ensures that both groups

are aware that there is activity that falls under each of their respective groups—the two groups may collaborate or agree that one or the other will take the lead in bringing the incident to resolution.

Care should be taken to ensure that information sharing is handled consistently and appropriately. Different services will have different information handling requirements. Depending on the specific situation, information flow may be restricted due to specific policies (such as an information disclosure policy). Moreover, these differing requirements may even prevent any sharing of information unless either some form of data cleansing can be enforced or appropriate contractual agreements are in place. This issue must be considered before deciding to share information between any given services and reviewed whenever policy and procedure changes occur.

It may be necessary to give different priorities to similar types of requests depending on the source of the request. For example, the incident handling service could obtain simultaneous requests for incident statistics from both the vulnerability handling service (e.g., to assess the frequency with which a given vulnerability is exploited and prioritize further action) and the education/training service (in the process of updating public presentation materials). A higher priority would likely be given to the request from the vulnerability handling service, as it could have a more immediate effect on the overall incident statistics and how that affects the examination and analysis of a vulnerability (and perhaps the constituency, if the vulnerability handling service would be providing guidance to them). It would be a more “reactive” need than the “proactive” needs for updating training, education, and awareness components. This example also raises the issue of information sharing again. The information provided to the vulnerability handling service would most likely include details on the frequency of incidents reported that involve specific methods of exploitation. The information provided to the education/training service, on the other hand, would likely be sanitized of specific details about the information for any public training offering (at least in such a way as to remove details of yet unsolved exploitation methods, sites, etc.).

Some basic examples of possible information-flow relationships between the most commonly provided CSIRT services and the incident handling service are outlined in Table 5. These examples do not attempt to be comprehensive or specify mandatory interactions. They provide a flavor of the type of interactions to be expected. Of course, when considering your own set of CSIRT services it will be important to build a matrix of all possible service interactions, not just those with the incident handling service.

Due to the limited resources available to many teams and the close associations between some of the common services, the distinction between different services may become blurred. When the distinction becomes artificial, it is probably wise to merge the closely related services into one service; the separate parts can then be labeled “functions” within the service according to the terminology of this handbook.

Table 5: Examples of Possible Information Flow to and from the Incident Handling Service

Service Name	Information flow to incident handling	Information flow from incident handling
Announcements	Warnings of current attack scenarios	Statistics or status report New attack profiles to consider or research
Vulnerability Handling	How to protect against exploitation of specific vulnerabilities	Possible existence of new vulnerabilities
Artifact Handling	Information on how to recognize use of specific artifacts Information on artifact impact/threat	Statistics on identification of artifacts in incidents New artifact sample
Education / Training	None ²⁰	Practical examples and motivation Knowledge
Intrusion Detection Services	New incident report	New attack profile to check for
Security Audit or Assessments	Notification of penetration test start and finish schedules	Common attack scenarios
Security Consulting	Information about common pitfalls and the magnitude of threats	Practical examples/experiences
Risk Analysis	Information about common pitfalls and the magnitude of threats	Statistics or scenarios or loss
Technology Watch	Warn of possible future attack scenarios Alert to new tool distribution	Statistics or status report New attack profiles to consider or research
Development of Security Tools	Availability of new tools for constituency use	Need for products Provide view of current practices

The following example highlights the relationship between services and the need to evaluate information flow between services.

Example: Consider the scenario where a CSIRT offers (in addition to an incident handling service) a detailed security assessment service involving assigned team members that attack and test the specified systems. During the tests, administrators of the systems and networks involved are rarely made aware that the security assessment will take place. So, if during the test an insecure host is successfully attacked, the system administrator for this compromised machine may notice the activity, perceive it as a break-in, and report it as such to the CSIRT. If the security assessment service provides the incident handling service with advance notification of the test, the CSIRT team may

²⁰ In the context of the flow and interaction outlined in Table 5, we are not describing the training or mentoring that the incident handlers might need or obtain from other sources (although we recognize that it is needed and is an important component to the knowledge and skills within the team). CSIRT education/training services are generally the “recipients” of output that originates from the other services (e.g., they gain the knowledge from the incident, vulnerability or artifact services, etc., provided by the incident handlers) and integrate information into the training products. These products are in turn packaged and provided to constituents in an appropriate form through training classes, seminars, meetings, or other types of venues.

first verify with the security assessment team members if the activity reported is due to the testing. If not alerted in advance, the incident handling service might begin to expend unnecessary effort to respond to what they consider a legitimate incident report, such as alerting legal counsel or requesting support from other departments. As a result, precious resources of the CSIRT may be needlessly wasted. More importantly, the reputation of the CSIRT may be damaged in the eyes of those outside the team (such as the site management, system administrator, or legal counsel) because it rightly appears that within the CSIRT, one part of the team does not know what another part of the team is doing.

2.5 Policies

Policies are governing principles adopted by organizations or teams. This section will discuss in general terms what policies are and should be, and what properties they should have. But documented policies are not the end of the story. It is important to understand whether they are implementable, enforceable, and function as expected. This section concludes with a discussion of these issues. For more in-depth coverage of global policies (such as information disclosure policy and media policy) that are fundamental requirements for any CSIRT, refer to Section 4.2, “Fundamental Policies.”

The policies of an organization need to be clearly stated and understood by all members of the organization. Without a clear understanding of policy, it will not be possible for the staff to correctly implement and enact their responsibilities.

Where services are essentially defined “for the customer” (e.g., incident response support service or education/training service), the underlying policies for delivering the service are mainly internal guidelines for the team that dictate appropriate behaviors for some specific activity. Example policies in this case could include information categorization, security, media, and code of conduct. The latter two examples may prompt the objection that these are hardly only internal and that they have a lot to do with external communications. True enough, but this external aspect is not something *offered* to the customer; it is not a service in itself, it merely affects the manner and quality in which the service is delivered.

A policy may be service-specific: an incident response support service, for example, may require a specific policy on caller authentication (defined in a procedure for verification of a caller before incident information can be discussed). Caller authentication may not be necessary within another service such as education/training or technology watch. In this section, the emphasis is on the overall policies *encompassing* the services of a CSIRT. However, most of what will be said here will also apply directly to any service-specific policies.

It is important to understand the relationship between policies and procedures, since these are often mingled and mixed. Procedures detail how a team enacts activities within the boundaries of its policies. Procedures can be very beneficial to help make a policy successful, but only on rare occasions can policies exist without corresponding procedures. An extremely simple media policy is “Be very polite to the media and never lie, but only mention generic anonymous information.” However, corresponding procedures help many staff members stay within the policy guidelines, especially in situations of stress. In the following discussion of policies, we will only make reference to procedures where this will add to an understanding of the case.

2.5.1 Attributes

Though it may seem trivial, it is essential to stress that a policy should not be defined as a set of detailed procedures. A policy should outline essential characteristics for a specific topic area (consider the above media policy example again) in such a way that all the necessary information is provided on which detailed procedures can be based to help implement the policy. All policies must be written with comparable levels of abstraction and should undergo legal and appropriate compliance review. Table 6 describes a set of attributes that every policy should have.

Table 6: Basic Policy Attributes

Attribute	Description
Endorsed by Management	Like the mission statement, a policy cannot be enforced unless it is endorsed by senior management.
Clear	Any team member, whether technical, management, or administrative, should be able to easily understand what a given policy is about. Avoid unnecessary jargon, don't be ambiguous, and use very short sentences. If possible (according to your disclosure restrictions), ask someone who is not in security or IT to read your policies. If he or she cannot understand them, rewrite the policies!
Concise	A good policy is a short policy. A long policy is either a bad policy (or uses too many words) or one that may actually include a lot of procedures. Unfortunately, security policies in practice often tend to not be concise, confusingly mixing the management aspect (the policy) with the operational aspect (the procedures), resulting in a mixture that nobody really cares for. Strive to avoid this condition!
Necessary and Sufficient	A policy should include all that is needed to dictate appropriate behavior in some topic area (e.g., security policy), but no more than that—no redundancy, no resiliency. That can be built into the corresponding procedures and quality control.
Usable	Avoid statements that sound nice but are of no use because they are open to interpretation, such as “state-of-the-art security will be provided.” Common sense statements like “treat your customers with respect” could be appropriate inside a policy: they are usable, because people share a common understanding about them.
Implementable	A policy must also be implementable. In the “treat your customers with respect” example, this may mean the addition of a statement essentially saying that regular training must be provided to help the staff understand how to deal with customers.
Enforceable	Policies must be enforceable; otherwise they are of little or no value. Usually when a policy is implementable, it is normally also enforceable unless it contradicts itself. Concrete measures are needed to assess the usage of the policy. Example: An example of a contradictory policy is the security policy that ranks internal information security as priority number 1 but at the same time ensures absolute privacy for its staff; the latter makes it hard or even impossible to enforce security in case of an insider threat.

2.5.2 Content

The content of a policy is mainly a definition of behavior in a certain topic area. Examples include how to behave toward the media, classify incoming information, and deal with the results of human errors. These features are boundary conditions for any policy definition. It is also possible to distinguish some generic features that should appear within the content of policies. These features are listed in Table 7 and, where appropriate, include examples (for consistency these examples will focus on the media policy arena).

Table 7: Policy Content Features

Feature	Description
Mission Link	Describe how the policy is derived from the mission statement.
Identification of Roles	The parties/people involved in (aspects of) the policy should be clearly identified. Example: media, media liaison, and other staff.
Responsibility	Duties and responsibilities of the identified parties should be defined, when appropriate. (Note: you cannot, however, define the duties of the media.)
Interaction	Describe the appropriate interaction between the parties identified within the policy. For example, only talk to the media in person or via telephone, insist on a list of questions to be asked in advance of the interview, insist on reviewing written text before publishing.
Procedures	Essential procedures can be called for, but should not be explained in detail within the policy. For instance, state that a procedure must be in place to verify the identity of a member of the media, or that only authorized staff may talk to media (and only after appropriate training).
Relationships	Identify the relationships between this policy, services and other policies. In the media policy example, a relationship with the security policy is obvious, as well as a relationship with the information intake process of an incident handling service.
Maintenance	Describe responsibilities and guidelines for document maintenance and update (e.g., the request may be received through the triage function).
Glossary of Terms	It is essential to ensure that the CSIRT's definitions of terms are provided; all local organization terms and all acronyms are defined. This will ensure that everyone understands the policy, especially new team members.

2.5.3 Validation

After a policy has been defined it is advisable to check its validity in practice before actually implementing (and possibly enforcing) it. Checking validity means finding out if all the ideas in the policy can actually be translated into real-life behavior.

Example: Only stating “always be nice” is not much help when one is confronted with persistently aggressive people.

The following issues should be taken into account with regard to policy validation:

- If possible, ensure that the people responsible for the policy validation are not the same people who created the policy. This will help prevent any conflicts of interest and ensure an objective evaluation of the policy.
- Pay particular attention to validating the policy attributes and content features detailed in Tables 6 and 7 to ensure that policies are not ambiguous.

- Undertake consistency checking of the policy in relation to other policies, services, and procedures; and also within the policy itself.
Example: Espousing network security yet using the practice of transmitting passwords in clear text.
- Validate implementability and enforceability. Implementing the policy as a pilot and choosing some worst-case scenarios to check on real-life behavior, including enforceability, is a very good way to accomplish this validation of the policy.

2.5.4 Implementation, Maintenance, and Enforcement

Once validation of the policy is completed, feedback should be given to the policy makers so they can make any revisions. Once the revisions are made the policy should be re-tested. Once the validation is complete and no more changes need to be made, the policy can be implemented.

Once that is done, the policy will need to be maintained, i.e., it will be necessary to make regular checks on its behavior in real life. Many of these checks will be equivalent to the validation checks, and some new ones will be added. An example of the latter could be using the media policy to check if the media is indeed informed within a preset number of hours following their request for information. Clearly this real-life behavior could not have been measured previously during a validation phase.

The checks originating from the validation process and the continued validation of the policies are really checks on the behavior of *quality parameters*. Both maintenance and enforcement (what to do if the checks say “something’s wrong!”) are part of the regular quality assurance system, discussed in Section 2.6, “Quality Assurance.” Every policy must have a regular maintainer who keeps track of the quality of service effected through use of the policy and proposes changes to the policy as appropriate. Things change over time, and no policy should be implemented once and used “as is” forever. The excuse “That’s the way it has always been done” is not acceptable in light of the changing nature of the technologies in use today, and the role(s) of the CSIRT and the services it provides to its constituency.

2.6 Quality Assurance

Just defining services, the flow of information between them, and procedures to make things work is clearly not enough to serve a constituency *well*. An associated form of quality assurance is also required. This assurance can range from a statement of the form “we will try,” to fully specified sets of quality parameters backed up by associated enforcement and escalation procedures, along with and liability and penalty clauses.

In the CSIRT arena, standard approaches are rarely used across all teams, although there may be similarities. A few teams attempt to at least prioritize incidents and work on what they regard as high-priority incidents first. One team reported that their most commonly used measure of quality was the “absence of complaints reaching senior management.” However, these were exceptions, as only very few teams undertake quality assurance (QA) either formally or informally. Lack of QA results in inconsistencies in the service provided, services that do not fulfill their purpose, and inappropriate use of staff resources. In this section we suggest a QA approach suitable for the CSIRT environment. Time and experience will tell what other approaches are (more) suitable for this domain.

We will describe the basic quality assurance components and their use in the CSIRT environment. Our QA system consists of three parts: definition of a quality system, checks, and balances.

In the definition, parameters are given that together describe the system’s quality. The checks are there to actually measure these quality parameters. Finally, the balances ensure that the results of these measurements are used to assure quality.

2.6.1 Definition of a Quality System

The first step is to look for the smallest set of quality parameters sufficient for describing the QA level required by the mission statement. When more than one service is offered, several sets of quality parameters may be appropriate, one for each service. And, at an even more granular level, there can be subsets of parameters for functions within services.

A quality system should be defined using a top-down approach, starting with the mission statement and going down to policies and services, functions that comprise those services, and all associated interactions and procedures. The mission statement should be such that one can derive a general sense of the CSIRT’s perceived quality. The mission statement could involve quality perceptions like “timely,” “best effort,” or “flexible.” Clearly, all subsequent quality definitions should be in line with the mission statement.

The set of quality parameters includes those for all policies, services, service functions, and procedures. Each of these elements will have its own subset of unique quality parameters, but some parameters may be common between services or service functions. It is important to bear in mind, however, that quality is dynamic, definable not only within policies, services, and service functions, but also between them, like information flow. Therefore one also needs to take into account the interactions of services when defining quality parameters.

Example: Suppose the mission statement of a team mentions both incident and vulnerability handling services. Obviously it is then practical to define two different sets of quality parameters, one for the incident handling service, the other for the vulnerability

handling service. A typical parameter for the incident handling set would be the maximum time that it takes to react to a constituent's initial incident report. A parameter for the vulnerability handling set could be that one only gives out advice about a vulnerability if a solution is present.

Extending the quality parameters, one can then consider the interaction between the incident and the vulnerability handling services also yielding good examples of quality parameters. For example, such a parameter could be the maximum time that a vulnerability service should take to provide an assessment of a vulnerability to the incident handling staff (or others) when the incident handling service discovers possible evidence of a vulnerability exploitation while analyzing an incident.

To further clarify the diversity and breadth of quality systems, more examples of quality parameters are given below:

- response time for service events (e.g., incident, vulnerability report) and/or priority scheme
- level of information provided for service events (short term)
- time-to-live for service events
- level of information provided on longer term (reporting, summaries, announcements)
- secrecy
- verification

Having identified a suitable set of quality parameters, the quality system definition is completed by assigning values to all *quantities* among the parameters.

Examples:

Parameter	Value
Follow-up time on vulnerability reports	For all non-urgent vulnerabilities, the CSIRT will follow-up with a constituent within two working days of the initial report.
Follow up on high-priority incidents	Every high-priority incident will be acknowledged within two hours. Analysis will start within the first hour of receipt of such a report.
Follow up on low-priority incident reports	Every incident report will be acknowledged within 4 hours. Analysis will start within the first 48 hours after receipt of such a report.

It is important to realize that the quality system is not necessarily a static one, i.e., with all parameters simply defined and assigned specific values. It may well be the case that the state of one parameter dictates the values assigned to other parameters, or that one set of flexible parameters is used.

Example: Consider a crisis situation, when everything looks different from normal. This can be handled with two different approaches:

- a. Suppose a parameter “crisis” exists with possible values “YES” and “NO,” and several other quality parameters are also defined. If “crisis” equals “NO,” all of these parameters are in use and have values assigned. If, however, “crisis” changes to “YES,” a number of the quality parameters are ignored and the remaining ones (like response times) could be assigned more stringent values.
- b. The CSIRT simply uses flexible parameters, such as “95% of all low-priority incidents are handled within five days.” These are communicated to the constituency, who need not be explicitly aware when the CSIRT is in crisis mode.

In both cases the CSIRT should take into consideration that constituents might feel left in the dark if they are not informed about changes in quality parameters for the service levels.

2.6.2 Checks: Measurement of Quality Parameters

It is not sufficient to just define a quality system—it must be checked to validate whether or not it lives up to your expectations. Checking quality parameters (measuring real-life behavior) is thus an essential part of any QA system.

Having defined quality parameters, one also needs to define how to check these parameters and how to measure them. This is by no means a trivial task and dictates some serious *a priori* measures, like establishing a reporting system. Also you need to audit your check system regularly to see if it functions appropriately in real life and if it meets the ultimate demand: to be a good check on quality.

This demand (for a check) explains why quality parameters should be clear-cut and preferably quantifiable: it is hard to measure qualifications like “good” or “bad,” whereas it’s easy to measure parameters such as the average time taken to act upon an initial incident report.

It is useful to note that the frequency used to check on your parameters is really also a quality parameter in itself. Its value must be carefully optimized as well: too few checks clearly endanger QA, whereas checking too often will result in more time being used to live up to expectations, reprioritize etc., instead of actually getting the work done.

2.6.2.1 Reporting and Auditing

To track quality, it is necessary to have a workflow management system (for details of workflow management systems, see Section 4.3.2, “Workflow Management”) to measure parameters (such as response times, problem categories, and priorities) and a reporting system (to measure the use of standard and escalation procedures). Many different levels of reporting exist; a few obvious categories are reporting to operational management, overall management, the constituency, the world. When workflow management software is in place there is a tendency (or desire) to make reports automatically available. Care should be taken that such reports are not generated just because they can be done; they should be generated to provide information that is useful to and needed by the recipient.

Regular auditing (or tracking) of the QA check system itself is also necessary to ensure the quality. The system must be checked for both sloppiness and inadequacy. It is necessary to

- minimize the number of procedures necessary and make them crystal clear
- ensure that the CSIRT staff members understand why the procedures are in place (to enhance the motivation of the staff members)
- avoid the tiny details: it helps staff motivation if they are allowed to think for themselves (and it is nearly impossible to make rules for everything)
- do audits and feed the results into the review cycles

One common mistake is to develop long, complicated rules for evaluating the system, which often results in the need to do very rigid audits to ensure that the rules have been followed. Often these audits become so rigid that they have to be announced in advance, the result being that meeting the audit demands becomes a goal in itself instead of the audit serving the bigger goal of helping to assure quality.

A quality check system can become inadequate even if it was originally designed to be very effective. This happens because quality parameters can change.

Example: You define the initial response time for incident reports from customers as a parameter. This may be perfectly acceptable until you introduce an automated email response service that just “acknowledges” reports in a short message that says: “We have received your report ... Thank you.” While the process might be very fast indeed, it probably is not the quality parameter you set out to initially measure. Therefore you would need to refine your parameters to state, for example: “An initial, automatic response will be sent out by email whenever the customer report is received. A personal response by one of our CSIRT staff will be sent out by email within 24 hours of the initial review.”²¹

²¹ Or some other timeframe (or parameter) you define.

2.6.3 Balances: Procedures to Assure Quality

Doing quality assurance checks on the real-life behavior of quality parameters is not enough. Procedures must be in place to enforce quality when it is at risk. Then, escalation procedures can be defined in the event that standard enforcement procedures fail, or if the quality system itself proves inadequate. Finally, penalty and liability clauses can help enforce quality and at the same time prevent the service provider from becoming excessively vulnerable to potential lawsuits. These procedures and clauses can best be characterized as the “balances” for quality assurance.

In the demanding environment of a CSIRT, where staff stress levels are high and resources are generally fragile and stretched, it is important to ensure that staff members are able to accomplish their work to a high standard of quality without overwhelming them with unnecessary hurdles. As such, there is a need to seek the right balance between procedures, checking, and the ability to get the job done. Correctly written procedures will ensure a buffer for human errors; any procedure not taking (human) error into account is flawed by design (see Section 4.2.6, “Human Error Policy”).

Also it is advisable to give customers (the constituents) some method to enforce quality, though this will normally be an indirect process. Not only does this “sell well,” but the best quality judgment often comes from those who actually use (or suffer) the service. One convenient way of granting constituents influence is by implementing measures such as user groups and/or advisory boards. Admirable though these measures are, the most effective way is probably by implementing penalty clauses: meaning the team has to pay or refund the customer money if it performs below the expected level of service.

Note that it can also be the customer who fails to live up to his part of the deal. If that continues to be the case and is grave enough (e.g., as grave as non-confidentiality), then procedures should also be in place to discontinue or reduce support for such customers.

From the CSIRT staff’s perspective, escalation procedures are usually defined and an integral part of the overall CSIRT processes. However, operational management should be able to swiftly and effectively notify the higher levels of management when quality is truly at risk; waiting for the monthly or quarterly report to have its impact is not sufficient. The routine should include a decision on whether or not to notify customers of the problem and the estimated time to fix it. The decision will depend on the agreed service levels and the direct disturbance caused by the problem. Escalation can also take place when the quality system itself fails and needs to be fixed.

Defining and advertising quality (but not assuring it) will cause the CSIRT to be liable in most countries if service parameters are not met and a constituent claims damage as a result of this failure. However, even in normal cases where a QA system is in place, including

checks and balances, in some countries (notably the U.S.) liability claims are still to be expected. In some cases, adding liability clauses to QA will be useful, especially when penalty clauses are also in place. Such clauses are best handled by legal experts; simply denying responsibility for financial damage is not enough in most countries.

The key point: If you define quality, make sure you assure it. Prioritize your assurance tools: education and awareness building are more effective tools than increasing pressure, especially in the long run. If a workflow management software system is in place, it is possible and advisable to integrate the regular enforcement and escalation procedures into this system. This saves work over time and also creates the possibility of producing reports on the use of these procedures.

Last but not least: Procedures and policies are not made for eternity, and thus must have owners and/or maintainers, and a well-defined life cycle. All too often procedures are created in the project phase—and once the project is over, the change control vanishes, but the procedures are there to stay, out of control, until somebody really stumbles over them.

2.6.4 Constituents' View of Quality

The set(s) of quality parameters for internal use must be complete to ensure that an appropriate quality level is maintained with respect to the mission statement. However, the set of quality parameters communicated to the constituency is generally some subset of those used internally.

From a commercial point of view, it is advisable to communicate a mature (if not the full) set of parameters to the constituency. The message is that the constituency is taken seriously and that you have nothing to hide. On the other hand, from the same commercial point of view and sometimes from a liability point of view as well, it may be wise only to communicate those parameters that are easy for your CSIRT to assure.

A compromise between both extremes is the best option. In any case, avoid communicating quality parameters whose definition is not crystal clear or establishing parameters that are impossible to quantify, however useful these may be to help assure overall quality. Constituents tend to dislike what they cannot grasp.

2.7 Adapting to Specific Needs

In many instances the reason for forming a CSIRT results from a specific need or problem experienced by the organization. Logically, whatever general structure is chosen for the CSIRT, it will be adapted to at least suit the specific need.

Example: A heterogeneous user community that experiences a lot of network-based attacks will establish a coordinating CSIRT facilitating support and coordination within this community. This team will not necessarily concentrate on computer viruses.

Example: The coordinating CSIRT from the example above notices over time, from reports by the users as well as from other peer teams internationally, that computer viruses and worms are becoming a growing concern. The CSIRT identifies and obtains resources to invest in new staff members who specialize in malicious software attacks.

Example: An organization with significant computer virus problems that already has a well-established antivirus team builds a CSIRT that does not respond to computer viruses or other malicious software incidents, but concentrates on network attacks. Agreed interactions and interfaces between the CSIRT and the anti-virus team assures that cooperation takes place whenever useful and necessary.

Every team has its own circumstances to adapt to. The result is that no two CSIRTs are alike in details, only in basic structure.

Example: CSIRTs with full authority (including access to a constituent's systems) but working in an environment with highly sensitive data (military, commercial, health care) must adapt to the extra stringent security measures and extensively screen their personnel. The levels of screening may be different across these CSIRTs, and the services provided may be similar or very different.

Clearly each CSIRT will begin developing and adapting their requirements when defining the mission statement and services. Naturally these changes must also be reflected in the quality assurance system as appropriate. But where adaptation will be most evident is in a CSIRT's policies and procedures—and in the rather practical treatment of team operations (described in Chapter 4).

Fundamental policies will implicitly surface, however, as illustrated in these next examples.

Example: A military or company CSIRT will have a rather restrained media policy covering all issues and topics.

Example: A coordinating CSIRT serving a national research network user community will have a media policy that allows them to describe all technical details but does not allow the CSIRT to reveal the identities of people or organizations involved.

Example: An anti-virus CSIRT will have stringent procedures for how to deal with incoming binaries (such as virus samples), including an isolated test environment and complete backup images to reinstall the test environment back to its initial clean state.

Example: A coordinating CSIRT without any analysis services has stringent procedures as well, but when they receive a virus sample (binary) they will simply isolate the artifact after a mandatory computer virus check. As they have no analysis capabilities in-house, they may not do any further analysis of the binary received.

Two topics remain that deserve attention at this level. The first is the *general* ability of the CSIRT to readily adapt to changing circumstances that every CSIRT must face in order to do a proper job. The second topic is that of law, liability, and regulation. These will be discussed below.

2.7.1 The Need for Flexibility

CSIRTs need to be prepared for the dynamic environment of computer security incidents and attacks. A CSIRT needs to be ready to address any situation that may not be explicitly covered by its existing guidelines or expertise. Some of the factors that make the CSIRT environment so dynamic, coupled with their impact on CSIRTs, are provided in Table 8 as an illustration for the need to be flexible.

Table 8: Examples of Dynamic Environment Factors and Their Impact on CSIRTs

Factor	CSIRT Impact
The rate of incident reports a CSIRT receives cannot be easily predicted.	A CSIRT will experience unexpected and extended peaks in workload or conflicting priorities.
Intruders (or attackers) are constantly devising and implementing new methods of exploitation by devising new attack methods or modifying existing attack methods to open new exploitation possibilities.	The type and complexity of incidents reported to a CSIRT will change over time.
Advances in technology bring new possibilities for exploitation, such as those resulting from Java and ActiveX.	The technical expertise required in a CSIRT will change. CSIRT staff must keep up-to-date with new and emerging technologies.
In some countries laws are just being developed to address what they see as a new problem. Computer crime laws are under review and undergoing active revision in many countries around the world, in an attempt to keep pace with the changing technology and threats posed by intruder activity.	CSIRTs need to be aware of the constantly changing legal framework of the environment in which they operate, and adapt accordingly.
Varying demands will be made on the CSIRT based on the needs, technical expertise, experience, and level of understanding of each of the parties with whom it interacts.	Situations will arise when the resources within an unprepared CSIRT may be insufficient to respond effectively to meet the conflicting demands placed upon it.

Due to factors such as those detailed in Table 8, the types of incidents reported to a CSIRT, priority schemes used, nature of response, and appropriate reporting requirements may very well be expected to change over time. CSIRTs must ensure that they have flexible policies and procedures to enable the team to easily adapt to change—whether the change results from a variation in work load, technical focus, legal issues, or constituency needs.

Although these factors are usually outside the direct control of a CSIRT, some level of advance planning can ensure that the team is prepared for them:

- Be prepared to obtain and use external resources to address a crisis (whether extreme workload or conflicting priorities), or provide a reduced or revised level of service for the duration of the crisis.
- Undertake continuous staff education or professional development in both current and emerging technologies.
- Implement staff training programs.
- Ensure timely access to appropriate information resources.
- Encourage staff attendance at appropriate technical conferences.
- Ensure ongoing cooperation with management, legal counsel, and law enforcement (or others as appropriate).
- Ensure that service definitions, policies, and procedures are not so rigorous that they do not anticipate and allow for change or unexpected circumstances.

Most of these issues are explained in more detail in Section 4.2, “Fundamental Policies.”

CSIRTs should be flexible enough to meet the demands of their dynamic environment when unexpected events arise, but still ensure that such events are handled in a manner consistent with the team’s overall objectives and operating style. Unless the need for flexibility is addressed, the CSIRT guidelines will be too general to provide help and guidance, or too restrictive to accommodate unexpected events.

If changes occur in the CSIRT’s mission and operations that will also affect how they interact with their constituency, such changes must be communicated to the constituency.

Example: A CSIRT changes its operation to charge a fee for the services it previously provided to the constituency under some other funding model. In such a situation, the CSIRT must take appropriate measures to notify the constituency of that fact.

2.7.2 Legal Issues

Since none of the authors are legal experts, we can only offer opinions about what we have experienced or have seen others experience in this subject area. Our approach here is to bring to your attention the issues that you may wish to consider. Readers should check with their own legal counsel to identify the issues that are applicable to their own set of circumstances.

Access to legal advice for CSIRTs is critical; without it, the team can unknowingly take inappropriate or illegal actions that can result in the team’s demise. Small teams who do not

have easy access to legal advice are at a great disadvantage. They should at least seek legal advice prior to beginning service and when making major changes in policy or operating procedures, if at all possible.

Legal issues are a bit like quality assurance: they permeate just about every topic ranging from mission statement to operational procedures. This comparison also yields an interesting difference: Quality assurance is about saying *do-this-and-do-that*, whereas legal issues often revolve around *avoiding* doing or saying the wrong things that may make you, your team, or your organization liable. Of course assuring a stable legal position is not entirely the art of omission; positive action is required as well, such as making sure that possible evidence (for example, log files or other artifacts) is properly dated and authenticated.

Unlike quality assurance, where it is worthwhile to define an overall framework and set up measurements, with legal issues this is less feasible. In fact, legal issues are usually tackled whenever they apply within a given topic or area. This is not a bad approach for CSIRTs, whose core business is incident handling and not the law. The legal issues are boundary conditions and should be handled accordingly, in a thorough but pragmatic fashion. That is not to say, however, that a haphazard approach should be the outcome; an overview should be maintained, possibly by using a fixed set of legal advisors. Seen in this light, the term “legal issues management” is preferable to the commonly used phrase “legal advice.”

Institutional issues are comparable with legal issues; only in this case the national or international laws are replaced by the “laws” or regulations that govern the institution of which the CSIRT is a part. Clearly these regulations must also be adhered to. The biggest difference is liability, which will be virtually absent in the institutional case—unless breaking the institutional rules means making the institution liable!

In the remainder of this section, we will discuss management of legal issues from the CSIRT perspective and then focus on the important topics of liability and the main cause of liability—disclosure of information.

2.7.2.1 Legal Issues Management

Management of legal issues involving CSIRT teams means exercising a coherent view of the legal issues that the team faces. Legal advice should be given by a fixed set of people (mainly legal experts) who are experienced in this area and understand technical terminology and issues that form the basis of daily CSIRT work. This set of people (usually only a few or even one) should cooperate to ensure a joint coherent view. It is important that legal advisors are enlisted for the long haul (years instead of months) because the amount of domain-specific knowledge needed by your advisors should not be underestimated. This is especially true if you have only one advisor; it will take months to get a replacement up to speed. A very practical solution can be to use the legal advisors of your parent organization, but only when

these people are experienced enough to guide you through your specific problems. Continuity must be assured here as well. If the legal staff does not fit this need, you might be better off hiring or retaining a lawyer that better fits your specific requirements, if this is feasible.

The kind of experience that your legal advisor needs can be derived from the following topic areas. These provide examples of the kind of things that the legal advisor will have to look into and give advice on:

Contract Analysis

All contracts should be checked for legal validity, especially those with customers. This not only includes finding statements that are legally meaningless, non-binding or just plain wrong, but also identifying omissions that can be legally harmful to the CSIRT.

Service Definition and Quality Assurance

The service is what you sell (guarantee, promise, whatever applies) to your constituency. Clearly how you define your service and its quality assurance is what you will be held accountable for by your constituents, especially when things go wrong. So whatever it says, it should be legally sound.

Policies and Procedures

Policies and procedures should be checked for legal pitfalls, especially as policies and procedures often include statements that involve strong positive action such as sanctions. Such actions always inherit the danger of being opposed to some other laws. The following examples help to clarify situations in which advance legal advice on a CSIRT's policies and procedures would prove beneficial:

Example: Your policies may say that you are going to fire somebody if he violates your disclosure policy. This may very well cause a conflict with local or institutional laws: in some countries it's trivial to fire an employee, in other countries it's very hard.

Example: Suppose you have stated in your procedures that you will only exchange sensitive data with your constituents in an encrypted way. Suppose your constituent is in trouble and wants you to fax the data to them. If you refuse, even for the best of reasons, although you may comply with your own procedures, it is very doubtful that you are meeting your service goals for that constituent. It would be best if you knew in advance whether the encryption was a legal requirement or simply a preferred practice.

Example: Another instance of the above example would be the constituent who does not want to support encrypted communication at all and does not have the necessary tools available, yet wants to exchange sensitive information.

Waivers and Disclaimers

Disclaimers are found in many places: service descriptions, policies, Web sites, outgoing email, etc. All disclaimers should be checked for legal validity, or at least they should have a legal purpose. If this is lacking, the disclaimer should be removed. On the other hand, disclaimers may be added that have proven their validity in case law. A mythical example of an added disclaimer due to case law is the wonderful story about a little dog being warmed inside a microwave oven after having come home soaking wet. The dog died, and the oven manufacturer was found liable in court. Because of the case, the manufacturer added some appropriate phrases to the oven manuals. Or so the myth goes.

Example: You often read in contracts, on signs in a coatroom, or other places that such-and-such is in no way accountable for something going wrong (e.g., if your coat is stolen). This seems an easy escape but rarely is: often lawyers laugh at such phrases and say that it's up to the judge to decide. However, on the other hand, these escape clauses are not entirely useless; for if they are not there, the case may be even worse from lack of due care.

The CSIRT might require its customers to sign waivers that limit the liability of the CSIRT in some way (e.g., "best effort," "due diligence," or "industry standards"). Legal advisors may be able to suggest areas in which the CSIRT might most appropriately make use of such waivers. The same review and caveats that apply to disclaimers should be applied to the creation of waivers.

Non-Disclosure Agreements

CSIRT staff may be required to sign non-disclosure agreements (NDAs) both when starting and leaving employment with the CSIRT. If so, the same will certainly apply to part-time staff and visitors who share the details of the CSIRT work. This may also apply to the cleaning staff, guards, and others. Just drawing up a non-disclosure agreement and having people sign it may be legally ineffective. The signers must understand what the NDA encompasses. Before implementing an NDA, it should be reviewed by a legal advisor to ensure it is appropriately worded and matches any organizational policies. Without approval and review from the legal realm, the NDAs might end up being only a psychological safeguard and not valid before a court of law.

Proactive Measures

Suppose a law enforcement agency legally requests information from a CSIRT. Is the CSIRT prepared for that event and for what may happen afterwards? Suppose the CSIRT is summoned for a liability case. Is it prepared for that? Being prepared for such cases presupposes two things:

- doing your job the way that you said you would do it (in your service specifications) and demonstrating "due care." What "due care" means also depends on your local laws and should be discussed with your legal advisors.

- documenting and timestamping all significant events in your workflow and the workflow of incidents occurring, within reasonable boundaries

Example: If your CSIRT only saves logs for a specific time and has stated so publicly, and there is no law against this, nobody can complain if logs are not available any more once this time has passed. On the other hand, consider the opposite case: imagine that after the specific time an audit finds data that should have been deleted much earlier. This then might be the foundation of a liability case of its own.

The second point (documenting and timestamping) is where the proactive measures come in, and the legal advisors should advise or provide insight for approaches that support the needs of the CSIRT. Essentially the task is to identify the minimum level at which the CSIRT events (especially the incidents) should be documented, and also to identify the right way of doing this. The “minimum” is meant as that which is required by law, and that which may be required (or come in very handy) in obvious court cases. The “right way of doing it” means that the evidence (the documents, logs, archives, etc.) should be gathered so that it will receive high marks for completeness (within the set purpose), logic, and reliability when the material is legally requested or is investigated in a court case. This is less trivial than it sounds. An example will help clarify this point.

Example: In a Dutch case (State vs. Ronald O., 1993-5) where an alleged intruder was on trial, the evidence put forward by the prosecution included a set of logs. The logs still had original page numbers on them, but several pages were missing; they had been discarded a long time before by the party from whom the logs came because they contained no relevant data. Since pages were missing, the defense pleaded that evidence was being withheld. The judge dismissed the defense’s plea. However, a better way of handling possible evidence (the log files) would have prevented this issue from arising.

Some people advise keeping all data since archives are cheaper than lawyers. Others tell you to dispose of sensitive information as soon as possible so that it cannot be produced even if requested. The appropriate answer for each team will depend on the legal jurisdiction that they fall under as well as the team’s mission. If data is to be kept for possible legal use, consider the media that is used to store the information. Media such as CD-ROMs and microfiche/film, once generated, are not easily forged and can be produced at relatively low costs. Whatever the approach taken by your CSIRT, adequate staff training must be provided in this area (such as how to respond if law enforcement wants to seize CSIRT equipment).

2.7.2.2 Liability

A liability issue is anything that you say, do, or write; or that you do not say, do, or write; and for which people may want to sue you, with a reasonable chance of success in court. In countries such as the U.S., this is a reason for grave concern, given the number of liability cases and the huge penalties often resulting, which can easily ruin entire firms. In many other

countries, liability is not really an issue unless you have made a big mess of your operations, resulting in damage to other parties, such as your constituents.

The matter of liability is so dependent on local law and so legal in character that your legal advisors must be consulted on the subject. Proactive action is needed to prevent liabilities. The kind of action needed may vary depending on the context. The context can range from liabilities arising from the content of signed contracts (e.g., unable to provide service in line with your defined service definition by lack of availability of the service) that a CSIRT has with its constituents to those relating to information disclosure or omission. The examples supplied in Tables 9-11 illustrate different issues arising from these various contexts.

Table 9: Examples of Liability Issues Arising From Omission

Liability Context: Omission	
Issue	Example
Lack of information disclosure	You receive log-files that indicate an intruder's activities, and you fail to follow up on the lead. If this fact is uncovered, you may be liable for failing to act on the information.
Forgetting about side effects	You deal with a "new" vulnerability in a specific incident but neglect to notify the vendor and/or other teams of this vulnerability. Then a month later the Internet comes to a standstill due to exploitation of the same vulnerability.
Non-recognition of legal reporting or archiving obligations	In many countries you are obliged to report to or generate archives for law enforcement regarding any case that may involve a serious crime such as (intended) murder. This can also apply to crimes such as penetration of classified government systems.

Table 10: Examples of Liability Issues Arising From the Content of Signed Contracts

Liability Context: Content of Signed Contracts	
Issue	Example
Inadequate service definition	Your service is not available during public holidays or only on a limited basis; and this is not stated properly inside your contract, or you did not define what you mean by "holidays." There may be the possibility for your constituent to sue you if they experience an intrusion and seek your help during that time, but your service is not available.
Defined service level parameter is not met.	You promise your constituents online support that (for whatever reason) was not available to a constituent in an emergency situation.
Defined quality parameter is not met.	You do not live up to your promised response time when a constituent calls for emergency help during off-hours. If your constituent loses money in such a situation, they may well try to get back some of it through you and will not settle for an excuse not related to work.

Table 11: Examples of Liability Issues Arising From Information Disclosure

Liability Context: Information Disclosure	
Issue	Example
References to individuals or organizations	You give the impression that a party is involved in an ongoing attack. This may damage the reputation and business of the party involved.
Revealing identities	Liability exposure here depends on who is requesting the information. You may be liable if you reveal the identity (without prior consent) of victim sites to other victims, law enforcement, or the media. But you may not be liable if you are required to report the same information to an internal audit.
Distributing false information	You distribute information about a serious bug in operating system XYZ, and this turns out to be false information. The vendor of XYZ may not be pleased.
	You inform truthfully about a problem but advise a fix that does not work. If this is not obvious and damage results from it, you may be liable.
Incorrect advice (i.e., incomplete, outdated, or just wrong)	You advise a constituent to modify their firewall to solve some problems, but your fix silently opens up the LAN to other security problems.
	You present your constituent with information that is seriously outdated when better information is already available at sources open to the CSIRT; the team member just did not catch up, but your constituent may suffer from this.

How to limit your liability is again asking for an obvious answer: Do your job right and document it. Much about what to do has already been said. The following, however, offers a more structured approach to fighting liability and its results:

- Use standard contracts with legally “safe” phrases.
- Remove all statements from your service definitions, quality-of-service levels, and policies that may be untrue, difficult (or impossible) to meet, or are legally unclear.
- Make legally sound disclaimers.
- Define your workflow, policies, and procedures; and install appropriate documentation, enforcement, and control processes such that it is possible at all times to prove that due care is taken during your operations.
- Insure your service if the risks exceed the cost.
- Consider using waivers to limit or prevent the CSIRT from being liable for certain obligations or damage inflicted on a customer or other CSIRT.

2.7.2.3 Disclosure of Information

Information disclosure has the biggest potential of generating liability for a CSIRT. Disclosing information is not just about writing reports and advisories; giving advice on the telephone is also disclosing information. Apart from these “predictable” disclosures, there are also unpredictable disclosures:

- legal court orders
- information leaks from the CSIRT (whether from trusted experts or current or former employees)
- information gained through intrusion (physically or through the network)

Several examples of disclosure of information leading to liability have already been illustrated in the previous section. It cannot be emphasized enough that these cases of liability can be grave indeed, possibly involving huge claims. Some additional interesting examples of the possible impact of information disclosure, whether predictable or not, will help this understanding:

Example: If sensitive information about one of your constituents leaks out or is given out without thought, this may seriously endanger the security of your constituent's site, their reputation, or their business.

Example: If a site is under an ongoing investigation, and a related alert from another site is given or leaked to the suspect site, this may warn the suspect and hinder or even ruin the investigation. Often the CSIRT will not know about the ongoing investigation, a situation that cannot be reasonably anticipated or controlled. The CSIRT could limit its exposure to such a situation with an appropriate waiver.

Preventing information disclosure from creating liabilities is mainly a matter of controlling workflow and procedures such that due care is demonstrable at all times (as has been stated previously). Clearly the information disclosure policy must be of a restricted type. In other words, the policy should say that information should only be handed out on a need-to-know basis.

In most cases the CSIRT defines the terms under which information is disclosed. However, the CSIRT may have mandatory reporting requirements placed on it by organizational, local, or international relationships (with law enforcement, interest groups such as the Forum of Incident Response and Security Teams [FIRST], and others). The requirements and their consequences must be clearly understood, because they may affect information disclosure by the CSIRT and expose the CSIRT to liabilities. The most common example is that the CSIRT must comply to a demand for a report from internal auditors, whereas complying with a request from external auditors may or may not be mandatory, depending on the jurisdiction under which the CSIRT operates.

2.7.3 Institutional Regulations

Apart from local (and international) laws, your CSIRT will also have to live by the local regulations of its parent organization. If these regulations are seen as laws, then most of what has been suggested above also holds true for this case. The liability aspect for the team itself

may be minimal or absent (making the risks involved in breaking local regulations relatively small). However, it may be that breaking these regulations makes the parent organization liable. Then the case is the same as above, only with the added complexity of having to deal with the parent organization as well. If the risks are high, it is worthwhile creating a legal isolation for the CSIRT, such as a separate corporate body. This separation may make the risks easier to control. However, it may also pose other problems for the CSIRT when trying to interact with other organizational units within the parent organization.

Examples of institutional regulations are

- U.S. Department of Energy (DoE) regulations (e.g., CIAC, the CSIRT for DoE, is subject to those)
- company regulations (such as those in financial institutions or large corporations)
- military regulations
- international auditing standards
- other federal or national regulations

3 Incident Handling Service

In the previous chapter we discussed an overview of the basic issues that are of concern for each CSIRT. We now go on to discuss the mandatory issues related to the incident handling service in detail. In this chapter we will describe the fundamental components of an incident handling service and the procedures that need to be in place to support them.

Another insight into the structure of this chapter is to note that any description of the service must have at least two dimensions:

- **specification—the logical dimension**
A description of the purpose and structure of the service and its functions (Sections 3.1-3.2)
- **implementation—the technical dimension**
The actual set of tools, procedures, and roles necessary to implement the specified functions in a specified manner (Sections 3.3-3.8)

We conclude this section with a discussion of two general characteristics of the incident handling service (or, for that matter, for any CSIRT service): interactions (Section 3.7 “Interactions”) and information handling (Section 3.8 “Information Handling”).

3.1 Service Description

The services offered by a CSIRT should be clearly defined. Each definition needs to be understood and available to the CSIRT and the parties with whom it interacts; these definitions might be provided at different levels of abstraction. As discussed in Section 2.2, “Service and Quality Framework,” it is important that each service provided by a CSIRT is detailed in a corresponding service description. In this section, we will discuss the issues to consider when creating an incident handling service description.

The issues below are ordered logically to facilitate use as a template for filling out a CSIRT’s service description. However, when considering a description that is to be made available to others (e.g., to the CSIRT’s constituency), we encourage teams to refer to the results of the IETF working group, “Guidelines and Recommendations for Incident Processing” (GRIP) [RFC 2350]. In addition, example descriptions of several service levels from a technical perspective (independent from funding issues) can be found within the final report of the TERENA Task Force, “CERTs in Europe” [TERENA 1995].

3.1.1 Objective

To facilitate the development of its policies and procedures, the CSIRT should have a clear definition of its objectives. Continuing with the top-down approach, for example, the objectives for the incident handling service will be derived from the CSIRT mission statement, which in turn was derived from the mission of the security team, the parent organization, or other sponsoring entity. In accordance with the CSIRT's stated objectives, the range and extent of functions appropriate to fulfill those objectives can be defined. Table 12 shows some possible service objectives based on different types of teams with differing missions (this is by no means a complete list of service objectives).

Table 12: Range of Possible Incident Handling Service Objectives Based on Differing Team Types

CSIRT Type	Nature of Mission	Possible Service Objectives
International Coordination Center	Obtain a knowledge base with a global perspective of computer security threats through coordination with other CSIRTs.	Provide technical support in response to computer security incidents through coordination with other CSIRTs around the world. Through incident handling activities, seek and document technical details of current or potential intruder threats. Create and disclose information on detection, prevention, and recovery from intruder threats.
National Team	Maintain a national point of contact for computer security threats and reduce the number of security incidents perpetrated from or targeted at systems in that country.	Provide technical support in response to computer security incidents in the national language and time zones. Provide technical information to detect, prevent, and recover from vulnerabilities. Act as a liaison to national law enforcement agencies.
Network Service Provider Team	Provide a secure environment for the connectivity of their customer base. Provide an effective response to their customers for computer security incidents.	Provide technical support in response to computer security incidents. Ensure the security of the network infrastructure. Act as a liaison to national teams and/or others.
IT Vendor	Improve the security of its products.	Provide technical support in response to vulnerabilities. Coordinate with CSIRTs to analyze the basic source of incidents. Create and disclose public alerts about new patches and best current practice.
Corporate Team	Improve the security of the corporation's information infrastructure and minimize threat of damage resulting from attacks and intrusions.	Provide a center of excellence for incident handling support to system and network administrators and system users in the corporation. Provide on-site technical support for incidents impacting company systems to isolate and recover from intruder threats and attacks.

3.1.2 Definition

Before you can describe how your incident handling service can be implemented to achieve its purpose, it is important to understand the scope and depth of service that you need to provide with the resources available. A good place to start is to identify the issues that will constrain the level of service that you can provide. The service provided will be constrained not only by the stated objectives of the service, but also by the resources (physical, financial, and expertise) available to the team and the team's scope of authority in relation to its constituency. There are many different types of incident handling service in existence today. The following examples indicate how different services constrained by different limiting factors can still provide important roles and achieve useful purposes.

Example: The most common limiting factor is one of funding, which affects both staffing and the physical resources available to run the service. However, many security teams that exist today provide a minimal incident handling service consisting of simple instantiations of the triage, handling, and feedback functions (see Section 3.2, "Service Functions Overview") all combined as a single "service."

Example: CSIRTs on the national, organizational, and service provider level with limited funding concentrate on the coordinating of activities across their constituencies instead of providing detailed or on-site support. These teams play the role of a trusted broker by providing a central point of contact to and from their constituency and communicating direct incident information to the parties affected by an incident.

Example: At the other extreme, a CSIRT might have funding for several staff members but be unable to attract, obtain, or train staff with the necessary in-depth technical expertise. In such a situation, a team might be unable to provide comprehensive incident handling service with all functions in place independently. The lack of in-depth technical expertise prevents the team from providing an in-depth handling function, i.e., not being able to fully grasp what specific incidents are technically about. In this case, the team must rely on information generated by other more technically adept teams to use and disclose within their own constituency. The limited scope of such a team, e.g., relying on the service outputs of other teams, will degrade the service function to just the relaying of information.

With an understanding of the available resources, limiting factors, mechanisms that you may be able to leverage from within your existing organizational structure, and the purpose that you are trying to achieve, it should be possible to define the incident handling service. To do so, bound the level of service that you are able to provide and then impose that level of service across the range of functions necessary to provide the service.

It might be appropriate to produce two resulting service descriptions based on the same set of criteria and definitions. One description, for external consumption, would provide information such as to whom the service is available, how they would request the service, or

the CSIRT support a customer should expect. The other description, for internal consumption, would include the external description (the who, how, and what) plus more specific and detailed guidance for the internal implementation of the service and how the administration of the services is handled in the CSIRT. This latter description could, for example, include how the information is tracked and recorded, who is responsible for that function, specified approaches for prioritizing requests, and the determination of what exactly is provided to the customer within the boundaries of the service descriptions (e.g., there may be circumstances when either more or less support is provided, depending on current incident status, management direction, or customer funding). From this standpoint, the external description should really be developed as an outgrowth or a subset of the internal description. Depending on the type of constituency, the whole text might be rewritten for external consumption to make it more understandable to people who are not experts in the CSIRT field, or to provide additional background information for when service level support might change from expected behaviors.

3.1.3 Function Descriptions

The incident handling service generally encompasses reporting, analysis, and support (see the list of services in Table 4). The service can be further described to encompass four main functions: triage, handling, announcement, and feedback. A more detailed description of these functions begins with Section 3.3, “Triage Function,” and continues through Section 3.6, “Feedback Function.” The triage function is like an expert secretary; assessing incoming information and passing it on to the right desk (that is, function). The other functions are self-explanatory and need no further introduction at this stage.

For each of these (or additional) functions, clear descriptions should be documented for use within the CSIRT. These descriptions will assist in the generation of associated procedures. Aspects of the individual descriptions will be used to constitute other elements of the overall service description that will be made available to the parties that may access the service. However, various other implementation details that might be important for internal team use may simply confuse external parties, so it is not normally helpful to publish them outside of the team.

The function definitions should at least contain the following information:

- objective of the function
- implementation details and pointers to associated procedures

Examples: Is the function only triggered by internal action (i.e., from another function or service within the CSIRT) or can it be triggered externally (i.e., by a constituent or other party)? How is it triggered or accessed? What forms are used (e.g., email, telephone, reporting)? What data is required or desired to flow to or from the function by those accessing it? What is the life cycle of events?

- priority criteria used within the function
- level(s) of service provided
- expectations setting and quality assurance criteria used

3.1.4 Availability

Defining the availability of a service is not just a matter of answering the question “Who can contact who when?” but also “under what conditions?”

- **Who may access the service?**
Are particular aspects of the service restricted to the declared constituency (such as announcements or technical support for incidents) and other aspects available to a broader audience (such as accepting incident reports that affect the declared constituency from anyone)?
- **Times during which the service is available**
Are different levels of service available at different times? For instance, the feedback function might be available only during stated business hours, whereas the handling function might be accessible during business hours, or on a 24x7 (24 hours a day, 7 days a week) basis for all or some incident types, or to some particular subset of the constituency.
- **Conditions under which the service will be provided**
For example, are incident reports accepted only through completion of mandatory information requested using the team’s reporting forms?

3.1.5 Quality Assurance

Users of the service should be provided with information that sets their appropriate expectations for use of the service. Differing expectations might be set with other parties. For instance, a team is likely to offer greater quality expectations to their funding body and to their declared constituency than to other parties. It should be made clear exactly what is provided by and what is excluded from the service. It is also reasonable to give some indication of the time frame for a response that a user of the service can typically expect. Additionally the CSIRT should indicate what the users of the service can expect from the CSIRT in terms of handling different types of information provided to it. The expectations set should be in harmony with the priority criteria in place for the service.

3.1.6 Interactions and Information Disclosure

The users of the service need to understand what interactions take place between the CSIRT and other parties affected by the service and how information (disclosure) is handled. For instance, what can a user of the service expect to happen to any artifacts or log files that they supply to the team during an incident? Will these be shared with other teams, vendors, or experts, and if so under what conditions will that transfer of information take place and how will the information be sanitized and protected? These issues are discussed in more detail in Section 3.7, “Interactions” and Section 3.8, “Information Handling,” specifically Section 3.8.8, “Information Disclosure.”

3.1.7 Interfaces with Other Services

Points and criteria for information flow in the CSIRT between the incident handling service and other services with which it interacts depends on what other services the team provides. For example, triage is common to many services, and often a single triage function is provided for multiple services.

3.1.8 Priority

It is important to not only prioritize events within each function of the service, but also to understand the relative priorities between the functions that constitute the service and the relative priority of the incident handling service and other services offered by the CSIRT. The relative priorities assigned will reflect the overall goals and objectives of the team and the services offered. If resources are limited, the handling function most commonly takes precedence over feedback and announcements. This will also be true if a team is facing a dramatic incident rate increase without correspondingly employing additional members of staff. Regardless of how the situation arises, the concentration on the handling function will leave few resources for other activities, which will be apparent to the constituency.

Triage, however, is a prerequisite for the handling function to operate effectively. So limited triage might take place at a reduced level for all feedback and announcements to keep the constituency informed. Until the team can revert to its usual operating state, detailed triage effort must be focused on the handling function, and the current operating situation can be explained to other requesters to keep them informed.

Issues of prioritization are discussed in more detail in Section 3.8.6, “Prioritization Criteria.”

3.2 Service Functions Overview

As stated above, the incident handling service usually includes other activities that support the delivery of the service, consisting of the triage, R 7(li.9(ng)12, R 7(an)10.9(touncem)18)-2.ntice, ed trngt

Triage Function

The triage function provides a single point of contact and the focal point for accepting, collecting, sorting, ordering, and passing on incoming information for the service (see Section 3.3, “Triage Function”). In addition, this triage function can be the channel through which all external (outgoing) information is passed. It supports different input channels suitable to the needs of the team and constituency. An initial priority and possibly an associated tracking number are assigned to any apparent new event. As part of the triage function, additional actions (such as archiving, translation, or media conversions) can be undertaken to make it easier for subsequent incident handling activities.

Handling Function

The handling function provides support and guidance related to suspected or confirmed computer security incidents, threats, and attacks. This function can include a number of different activities. A review of the report (e.g., an incident report) is done to determine what occurred and/or the specific type of activity involved. Analysis of the report may include reviewing supporting evidence or materials (such as log files) to identify who is involved (or needs to be contacted) or the assistance that is requested/provided. The team will need to identify the appropriate responses (aligned with the CSIRT’s mission/goals/services), and the actual notification or follow up with the reporter(s) or constituent(s). This handling is described in much more depth in Section 3.4, “Handling Function.”

Announcement Function

The announcement function generates information tailored for the constituency in various formats to disclose details of ongoing threats, steps that can be taken to protect against those threats, or sanitized trend information on the scope and nature of recent attacks reported to the team (see Section 3.5, “Announcement Function”). For the purpose of this document, the scope of this function will be limited to its direct applicability with the incident handling service. However, in a CSIRT providing a broader range of services, publishing announcements can be considered as a service in its own right and would likely offer a much broader range of information derived from other services such as vulnerability or artifact analysis.

Feedback Function

The feedback function provides support for giving feedback on issues not directly related to specific incidents (see Section 3.6, “Feedback Function”). Feedback can be provided upon explicit request (e.g., by the media) or unsolicited, on a regular basis (in annual reports, for example), or case-driven (e.g., proactively informing the media). This function will provide at least a minimum set of support for frequently asked questions and might be seen as an interface for media requests or input to the team at large.

3.3 Triage Function

The goal of this function is to ensure that all information destined for the incident handling service is channeled through a single focal point regardless of the method by which it arrives (e.g., by email, fax, telephone, or postal service) for appropriate redistribution and handling within the service. This goal is commonly achieved by advertising the triage function as the single point of contact for the whole incident handling service. If a team wants to limit the ability of constituents and others to bypass the triage function, direct contact information for individual team members (such as telephone numbers or email addresses) should never be given out.

Because this is a common requirement across many CSIRT services, teams usually advertise a single point of contact for the whole CSIRT; and, regardless of the service required, a single triage function is provided for all the services that the CSIRT offers.

Example: Within DFN-CERT, the person undertaking the triage function is called the CERT Hotliner. This person is responsible for reading all email to the response team's alias, opening all postal mail, reviewing incoming faxes, and answering all telephone calls. The DFN-CERT hotline and the personal telephone lines for all other team members are forwarded to this person's telephone to ensure that all incident-related calls are dealt with centrally.

To stimulate the reporting and the collection of all relevant information, the constituency must be provided with easy to use and efficient mechanisms for reporting:

- a clearly defined point of contact
- specific details on the availability of the defined point of contact
- simple but defined procedures to follow
- clear guidelines on the kind of events to report
- supporting documentation (e.g., reporting forms and references to other available documentation) for reporter use

Once the information is received by triage, an acknowledgment of receipt will be sent, then the information will be sorted, prioritized, tracked, and passed on to other functions within the service. Additionally the triage function must decrypt encrypted messages and check digital signatures, preserve this information for later use, and allow for actually reading the content. To undertake this task, it is necessary for the triage function to have access to the data repository used by each of the other functions of the incident handling service.

Based on the information content and the data in the repository regarding existing service events, an initial sorting will take place to identify which function of the incident handling service should handle the information. The next step is to determine if the information is

directly related to any current or past event. If it is directly related to some existing or previously tracked event, it will be tagged as part of that event. Otherwise it will be tracked as a new event of a given type and tagged appropriately. In addition to being sorted and tagged, the triage function commonly assigns an initial priority to the information in accordance with the priority scheme in use by the functions within the service. If information enters in the form of hardcopy materials, it is common for the triage function to ensure that this information is entered online or a reference made online to the physical location of the materials.

Tools for entering, accessing, and tracking information and events can greatly facilitate and semi-automate data manipulation and searches. Such tools can support the staff responsible for triage by helping establish the identification of

- new events (incidents, requests, vulnerability reports, other information notices)
- information directly related to currently tracked events
- information directly related to a previously closed event
- events that are being tracked separately, but may have a direct relationship
- information that is considered out of the scope of the incident handling service

If the information contains insufficient detail or is incomplete, it is likely that the triage function will become slow, inaccurate, or incapable of serving its role. In such cases it may be necessary to seek more detailed information from the sender before the information can be appropriately triaged, which delays the process. In addition to direct tool support for the triage function, other steps can be taken to enhance the quality of the information, such as tracking numbers, standard reporting forms, and preregistration of contacts. The next three sections deal with these topics.

3.3.1 Use of Tracking Numbers

If a team uses a tracking number scheme and can encourage or require others to use the assigned numbers in all follow-up correspondence, this will greatly facilitate the triage process. To facilitate automated support, the numbering scheme should provide simple identifiers for human and tool recognition. In a robust tracking system, the tracking numbers are the “tags” that the system uses to automatically sort incoming information and store (correlate) it with other related activity, without human intervention at this initial stage. This streamlines the process and enables the triage function to focus more intensely on correct correlation of untagged information. Tracking numbers can easily be used in the subject line of email messages, documented on fax cover sheets, and specified in voice messages.

Tracking numbers should be used to track events under each function of the incident handling service. Different prefixes might be used for the different services. Since external communications have to be considered, part of the number should identify the team “owning” the number. Feedback, incidents, and announcements should each have their own variety of tracking number.

Example: CERT/CC uses the prefix identifiers CERT# for tracking incidents. VU# is used for tracking vulnerabilities. INFO# is used for identifying other, lower priority information. In addition, other prefixes are used for a variety of internal and external documents.

3.3.1.1 Unique Intra-CSIRT Tracking Numbers

A fundamental requirement for tracking numbers is that they be unique. Commonly, teams allocate numbers from a predefined range of integers as the basis for their numbering scheme. Within a team’s own incident handling service and preferably across all of their CSIRT services (tracking numbers can also be used for other services such as vulnerability handling and artifact analysis), a best practice is to use a unique prefix for each function, and also ensure that the tracking number following the prefix is unique. If the same number is to be used for more than one function, confusion and other difficulties might arise if parties forget to provide the prefix and refer just to the number. Ideally, it should not be possible to have incident number 60 and feedback number 60. The tag number itself should be sufficient to refer to a unique event. If a team plans to reuse numbers, strong controls must be enforced to ensure that there is enough time between closing a particular event and reusing its number. The delay must make it very unlikely that the number can be misconstrued as pertaining to an activity or event previously tracked with that number.

Example: In 1994 the DFN-CERT used numbers between 1 and 65,535. There were no plans to reuse any of these numbers. After four years of operation, approximately 600 numbers were used. Even with the increased rate of assignment of unique tag numbers, there will still be a significant number of years before old numbers will need to be reused or a different set of numbers needs to be adopted.

Example: The CERT/CC also uses a randomly generated set of numbers to track incident and vulnerability reports. No incident or vulnerability will receive the same initial tracking number, unless the reports are related and cross-referenced or subsequently merged into one larger activity. Even then, depending on the nature of the activity, there may be references to yet other tracking numbers for other relationships across and between activities. With the large volume of reports handled by the CERT/CC, however, other supplemental schemes were needed to handle other types of tracking identifiers for tracking and assigning reference numbers to other types of information (such as CERT Summaries, Information items, CERT Incident Notes and CERT Vulnerability Notes, to name a few).

Instead of using a limited integer number space for tracking numbers, other approaches have been adopted that provide an unlimited number of possible identifiers. Such approaches are desirable when the teams involved deal with large constituencies or wish to ensure a scalable approach that will work for several years without the need for procedural changes.

Example: In early 1994 AusCERT initially used an incident numbering scheme of the form YYMMDDHHMM. This was generated from the date and time that AusCERT “opened” the incident. While this addresses the size of the available numbers, it provided some other information about the report that was not wanted—such as how long this incident was known to the team and was first identified (or being tracked) by the CSIRT. Therefore they adopted another scheme.

3.3.1.2 Unique Inter-CSIRT Tracking Numbers

Tracking numbers need to be unique not only within a CSIRT, but also among different CSIRT constituencies. As multiple CSIRTs may be involved in responding to an incident, they will each use their own identifiers to refer to that incident. If this is not the case, there is the possibility that two teams will use the same identifier for different incidents, which could lead to confusion and cause delays in providing appropriate responses or feedback.

Example: Currently both the CERT/CC and the DFN-CERT allocate integer numbers within a given integer range for incidents. To ensure uniqueness, both teams provide a prefix to indicate their own tracking number. For instance: CERT#12345 and DFN-CERT#12345 are two separate and unique tracking numbers that refer to two totally unrelated incidents.²² For today’s teams, it is a generally accepted best practice to record all numbers for all teams involved when exchanging correspondence or other information about CSIRT activity with other teams. It greatly facilitates the identification, tracking, and correlation across such events.

If a team’s tools support recognition of various tracking number formats used by different teams, it will further facilitate the triage function. Teams are encouraged to reference each of the tracking numbers of other involved teams during their communications on related events to allow efficient identification and processing by all.

3.3.1.3 Tracking Numbers are Public Information

Because tracking numbers are used in the team’s external communications, they should be considered as public information and hence should not disclose sensitive information such as the names of hosts or domains involved. Other sensitive information to avoid in a tracking scheme includes information that would indicate the number, nature, or scope of events

²² It is possible that each team, through coincidence, may have its random number generator issue the identical number for related incident reports. In practice this is unlikely to occur; and even if it did, the use of the prefix would clearly identify the relationship of the activity to each CSIRT.

(particularly in the case of incidents) reported. For these reasons, use of some random-number-generating scheme is a better practice.

3.3.1.4 Tracking Number Life Cycle

The life cycle of tracking numbers also needs to be considered. If an identifier is used to track an event, then it is usually the case that the tracking number initially allocated will remain with that event from the point at which the event is identified until the event is handled from the team's perspective and is considered closed. But there are situations that arise that do not fit such a simple model and need consideration, such as:

- **Information is incorrectly triaged:**
Triage may incorrectly identify an event as new when it is in fact directly related to some other event.
- **Information is incorrectly tagged:**
Information may arrive with an incorrect tracking number and as a result be tracked inappropriately.
- **An event is reopened:**
If an event is closed and new information arrives for that event, then the event will be reopened.
- **Events merge:**
New information arrives that directly links two events that were previously tracked separately. This is difficult to archive. All incidents should be appropriately cross-referenced. Whenever incidents appear to be related they should be analyzed in more detail to determine if both incidents should be merged or not.²³

3.3.2 Use of Standard Reporting Forms

The use of standard reporting forms will facilitate the provision of complete and appropriate information being supplied to the team by parties reporting to it. This also facilitates the timely identification of new reports with associated activity and routing of information to the right function. It also improves completeness and comprehensibility of initial communications, which makes further processing easier. For most services, useful forms can be designed and implemented for use by the constituency or others (e.g., vulnerability reporting forms within a vulnerability handling service).

²³ Note: Even if they are merged, the issue of which identifier to must be considered—if many different sites are involved (and the incidents were formerly tracked separately), it can be complex to manage the overall event. In such cases, the original tracking identifier and a new “merged” identifying number may be used; alternatively a completely new number may be assigned once all the individual reports are merged. In either case, the affected parties (sites, constituents, etc.) will need to be notified and requested to use the appropriate tracking identifier. While the CSIRT may ask them to do this, recognize that there is no guarantee that the affected sites/constituent members will do so.

Within the incident handling service, forms may be made available for reporting incidents and for making information requests. To be of use, these forms need to be as clear and concise as possible and readily available for people to use when required. In support of both the triage function (in determining the relationship of the report to currently tracked activities) and the handling function itself, incident reporting forms commonly request the following types of information:²⁴

- contact information for the reporting site and any other parties communicating in response to the incident
- names and network addresses of hosts involved in the incident
- the nature of the activity
- description of the activity and relevant information (such as logs, associated time-zone information, and other artifacts)
- tracking numbers that may have already been assigned (by a local security team or another CSIRT)

Example: During a coordination effort, logs from one attacked machine are submitted to the CSIRT by a reporting site. The logs are of the form:

```
Mar 2 02 10:34:12 myhost tcpd[52345]. connect REFUSED from cumber.some.where
```

Without knowing the corresponding time zone for the above log, the team will be unable to provide the administrator of `cumber.some.where` with enough information to enable them to check their local logs for users that were logged in around this time. This problem is further exacerbated in international environments or countries with multiple time zones, as the possible time frame for the activity broadens.

Sometimes teams have trouble convincing people of the need to make a report in the first place. If some prospective reporters feel that the reporting form is cumbersome and not very effective, they may be more reluctant to report an incident. A team might choose the risk of losing some initial information in preference to not obtaining a report at all and let their constituency call or send information in “free form” (unformatted). The effects of this, however, mean the CSIRT staff will need more resources (time) to extract relevant information from such reports and enter it into the CSIRT tracking system. If forms are provided they must be as clear and concise as possible and must allow for easy reporting. This also applies for the number of forms used by one team. Consider whether it might be possible to use one basic reporting form or template that constituents can use to report a variety of different problems, requests, or other information to the CSIRT. In addition to providing forms and expecting the constituency to use them, the team must raise the awareness of the benefits of form use and must encourage people to report using forms.

²⁴ These are excerpted from the CERT/CC incident reporting form [CERT/CC 1997a]. A more recent incident reporting form (as an online reporting form or in text version) is also available at http://www.cert.org/nav/index_red.html.

3.3.3 Preregistration of Contact Information

In addition to the use of reporting forms, depending on the size and nature of a team's constituency, it may be possible to take some proactive steps to solicit information in advance that will be helpful to the triage function (as well as other functions comprising the incident handling service). This process can also be extended to solicit information in advance from other parties, such as other CSIRTs, law enforcement, etc. Such a registration process can help to prevent the need for standard questions to be handled on a case-by-case basis for every new report/request.²⁵ Useful items to preregister include, for example:

- trusted points of contact and associated contact information (must be routinely verified, at least once a year)
- information disclosure restrictions
- (verified) keys for encrypted and/or signed exchange of information

In some cases it may be useful to preregister other information such as domains or time-zone information for the site. But this will depend on whether or not the hosts covered are located in the same time zone as the registered contacts.

Example: Given the (numerically) small and well-defined scale of its constituency in 1994, AusCERT initially established a constituency registration process as detailed in *Forming an Incident Response Team* [Smith 1994]. This process included establishing trusted points of contact and information disclosure restrictions. AusCERT later changed this process when it became a fee-for-service team, but kept the same underlying concept of obtaining as much contact information as possible from its subscribers.

Example: CERT-NL serves a constituency defined by contract (between a national research network as the Internet service provider and its customer sites) and therefore can support a registration process for site security contacts and cryptographic keys to allow confidential and authentic email communication. Furthermore, because its constituency is of the academic type, CERT-NL is able to uphold a default information disclosure policy.

Example: Because the CERT/CC has such a broad-based constituency, it is not possible to obtain this type of preregistration contact information from every potential reporter. However, the team does have such details for regular reporters or those with whom the team interacts (such as trusted experts, sponsors, vendors, etc.).

²⁵ Kossakowski, Klaus-Peter. "The Role of Site Security Contacts." 7th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, Karlsruhe, Germany, September 1995.

3.4 Handling Function

The goal of this function is to provide response and support for reports received from the constituency (and possibly others). At a minimum, the function should provide some instantiation of the following attributes:

- **Reporting point:** A location for receipt of incident reports pertaining to its constituency
- **Analysis:** Some level of verification of the report and technical understanding of the activity. This will include identifying the appropriate responses that will be provided under the notification attribute.
- **Notification:** Passing appropriate response/recovery information to (at a minimum) constituents and preferably other affected sites and CSIRTs

The definition of the term “response” will vary from team to team based on the team’s definition of an incident and the objectives of the individual team’s incident handling service. In addition, other factors need to be considered, the most important of them being the priority assigned to a specific incident report and the relationship to the sites involved (e.g., if they belong to the constituency of the individual team or some other affected site).

Table 13 lists some possible instantiations of the functions necessary to carry out the handling service.

Table 13: Possible Instantiations of Handling Function Attributes

Attribute	Possible Instantiation
Reporting Point	<ul style="list-style-type: none"> • Deal with incoming reports that affect the constituency and pass them on (as appropriate) to the sites affected within the constituency • Deal with reports from the constituency that affect sites and CSIRTs external to the constituency, and pass them on accordingly • Both of the above
Analysis	<ul style="list-style-type: none"> • Examine log files • Identify affected sites • Point to technical documents or advisories • Provide technical support • Provide workarounds and fixes • Provide on-site assistance
Notification	<ul style="list-style-type: none"> • Point to resources that provide or can help establish appropriate points of contact • Provide a list of appropriate points of contact • Undertake contact of other parties affected in the incident • Undertake contact of other parties affected and law enforcement

Before talking about analysis in more detail, it is helpful to have an overview of the life cycle of an incident from an initial report through analysis to notification and closure. In order to

be able to perform the incident analysis function well, a set of specific information must be tracked. This is also discussed in this section.

3.4.1 Incident Life Cycle

Whatever a team’s definition of an incident may be, it will likely conform to the life cycle described in this section. As mentioned in Section 3.3, “Triage Function,” part of the life cycle of an incident may take place within the triage function, where an incident can be initially categorized, identified as a new event to track or as part of some existing incident already being tracked. The appropriate tracking number is assigned to it (either a new tracking number or the number for an activity already being tracked and to which it belongs). Note that a new incident can also be identified during the handling function as a result of incorrectly triaged information, information provided to the team under an incorrect tracking number, or new information being discovered as a result of more in-depth technical analysis. Figure 5 provides an illustration for the CERT/CC incident handling life-cycle process. This is described in more detail in the following paragraphs.

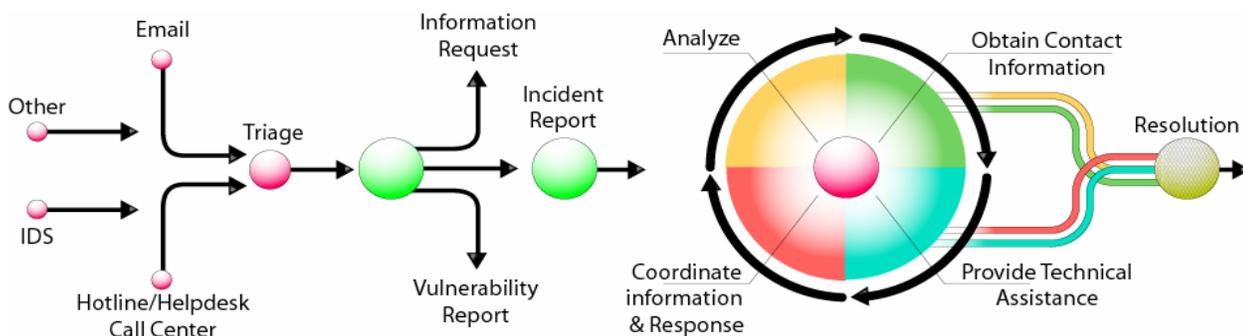


Figure 5: CERT/CC Incident Handling Life Cycle

Once an incident is opened, it may transition through many different states, with all the information relating to the incident (its change of state and associated actions) until no further action is required from the team’s perspective (the “circle” portion of the life-cycle illustration) and the incident is finally closed. It is also important to note that an incident (or event) can cycle through the analysis portion multiple times during the activity’s life cycle.²⁶

The closing of an incident normally occurs when none of the parties involved in the incident are identifying or reporting new information to the CSIRT and the CSIRT has undertaken its actions of appropriately responding to all parties affected by the activity. A team might also close an incident even if new reports are anticipated, but it makes no sense to follow up further (e.g., if there is nothing more the team can do).

²⁶ This CERT/CC Incident Handling Life Cycle graphic is extracted from the CSIRT training courses (see <<http://www.cert.org/training/>>).

Example: A CSIRT continues to receive more external notices of virus reports. The virus report has already been received, analyzed, and handled. Appropriate notification has been made available to the CSIRT's constituency. There is nothing new to be gained by further analysis. The incident is closed.

Example: A company CSIRT may not close an incident until any legal case associated with it is completed.

Example: A coordinating CSIRT serving a large constituency may close an incident if no further technical support is needed by the sites involved in the incident.

As a result, the criteria for closing an incident can vary from team to team.

It is equally important to note that, even if a team closes an incident, a site involved may still consider the incident open if they remain involved in resolving the incident, are preparing to recover their systems, or are involved in a court case against the perpetrator. In the latter case, if the CSIRT is a coordinating team with a broad constituency (say a country-level team), the CSIRT may not be involved in the legal proceedings that are being pursued by an affected site.

During its life cycle, an incident may transition through many different states, such as

- action required: Actions are required by the team in response to the incident.
- waiting: The team is waiting for a response from other parties external to the team.

When a CSIRT decides to close an incident, it should ensure that all of the affected parties are being or have been informed of the closure. This will help to set the appropriate expectation and avoid confusion in cases where someone thinks the incident is still open and wonders why they hear nothing further from the CSIRT. The team can either separately inform all parties involved when they close the incident or inform parties during ongoing incident correspondence. The former is more time consuming and is likely to generate a flurry of trivial email responses, or may result in someone finally providing a response that causes the incident to go back into an "action required" state. The latter encourages correspondents to provide information in a more timely fashion and is a more effective use of the often limited CSIRT resources.

Example: In previous years when there was more direct one-on-one incident response, the CERT/CC incident handlers would tell contacted sites that if no further feedback was provided by the correspondent by a specified date, their thread of the incident would be considered closed by CERT/CC.²⁷

²⁷ More recently, the CERT/CC focuses on a more strategic coordinating role of helping many (one-to-many) so less direct interaction with the end user occurs. Depending on the type of incident

Closed incidents may need to be reopened if new information is made available to the team, such as a report of rekindled activity at one of the sites involved. When the need arises to reopen an incident, the original tracking number should be reused, if possible.²⁸ However, if the activity is not considered to be a continuation of the original incident/report, it is appropriate to generate a new incident for the activity and issue a new tracking number.

Similarly, as mentioned earlier, new information may become available that directly links two or more incidents that previously appeared to be unrelated. In such cases, a team needs to decide if the incidents should be merged into one (if so, identify which tracking number should be used and who should be informed of it) or if they should remain as separate incidents and marked as related. Whatever scheme is adopted, all procedures, tools, and databases that might be affected will need to be capable of supporting such events. Such technical problems can be solved, but the human issues are not so trivial to solve and need to be taken into account. Unfortunately, even after an incident has been renumbered or closed, you may find someone replying to an old message containing an out-of-date tracking number to report completely new activity.

3.4.2 Incident Analysis

During the life cycle of any incident, analysis provides information that plays a major role in the decision-making process and next steps to take in accordance with a team's policies and procedures.

The first instance of incident analysis actually takes place during the triage function, occurring whenever new information comes in; this kind of analysis has been covered already and is not the topic of this section. Here we will focus on the more in-depth technical analysis of log files, malicious code, and incident texture.

Example: Consider an analogy with a hospital emergency unit. The triage function initially decides which incoming patient goes where (e.g., the first "initial" high-level analysis). Then more comprehensive analysis aspects come next, such as blood tests, scans, EKGs, and X-rays. The results of these tests help determine the next actions, such as medicine and/or surgery.

handled, the CERT/CC may keep a report "open" and continue to solicit additional information from correspondents. (It should be noted that the general auto-responder message from the CERT/CC infers that with the heavy incident load, not every incident report will receive individual follow-up.)

²⁸ Unfortunately, some trouble-ticketing systems will not allow ticket numbers to be reopened once closed.

Different types of analysis can be CSIRT services in their own right, separate from the incident handling service. One could, for example, offer an artifact handling service, additional to (or even totally separate from) the incident handling service.²⁹

Artifacts can be found in the remnants of intruder's activities. Searching for and analyzing artifacts, followed by neutralizing them as cost effectively as possible, is a craft of its own. A discussion of such separate services is not the goal of this section. However, since artifacts are often discovered during incident handling and since artifact analysis often is part of incident handling to some extent, reference will be made to artifact topics as appropriate, but not in any great detail.

There are two general classes of incident analysis to consider:

- **Intra-Incident Analysis**

Analysis of the issues concerning a specific incident. The most common types are as follows:

- analysis of any artifacts left by intruder activities (log files, exploits, viruses, Trojan horse programs, toolkits, etc.)
- analysis of the software environment in which the incident took place
- analysis of the web-of-trust within an incident

- **Inter-Incident Analysis**

Analysis of issues concerning relationships across and between incidents, that is, the analysis of the texture of ongoing incidents. This analysis is aimed at finding symmetries between separate incidents that might indicate equivalent or related sources of intruder activity.

Analysis is a very large topic area. We have chosen to cover it in detail in this chapter because a good analysis is critical to the provision of a competent incident handling service. We begin with a discussion of the importance of an overall analysis review (“the bigger picture”) and issues that affect the depth of the analysis undertaken.

3.4.2.1 The Bigger Picture

It is important to retain an overall grasp of all analysis results—the “bigger picture.” The bigger picture is largely concerned with trends (possible types of future attacks, security improvement), statistics (e.g., number of hosts involved, rate of incident reports), and case studies (e.g., understanding the intruder community or the impact on specific systems and applications). Each CSIRT will build its own bigger picture that is most relevant to its constituency.

²⁹ The CERT/CC has chosen to do this. See the CERT/CC “Overview of Incident Trends” presentation at <http://www.cert.org/present/cert-overview-trends/module-1.pdf>.

Obtaining this bigger picture can be difficult, as different staff may be assigned to tackle different types of analysis. Various people within the team will have different pieces of information resulting from their piece of the analysis. To obtain the bigger picture from the information available to the whole team, it is important to institute a process to collate the information. This can be accomplished by team members interacting in regular meetings or ensuring that incident supervisors obtain the information from team members.

Refining the bigger picture is especially useful in identifying lessons learned and so can help to improve response to future incidents. By studying lessons learned and experience as the result of handling incidents over time, the case history information obtained will often help staff to make the right decisions in the future, and sooner. Implementing a knowledge base to assist in this process can be a great help, especially for continuity's sake, since unlike personnel, knowledge bases neither change employers nor have holidays. Case studies can also be an excellent learning tool for new CSIRT staff members.

Example: The following shows how insight into the bigger picture can be provided through access to a good case history. As the result of both intra- and inter-incident analysis, it was noticed that on several occasions incidents had been identified where compromised systems had a combination of a certain weird directory name along with a Trojan-horse program located elsewhere in the file system. The next time this weird directory name was found, it prompted the team to search for the associated Trojan horse *immediately*.

It is useful and advisable to make the (appropriately sanitized) bigger picture available to other teams and possibly law enforcement.³⁰ This can be done in the form of free text news flashes or a common format for disclosure. A common format has not yet been adopted for use in the CSIRT community. Additionally, a team might choose to publish such reports to their constituency through their announcement function to keep the constituency in the loop, raise awareness, and provide insights into new trends and developments (see Section 3.5, "Announcement Function"). Along other lines, there is effort underway to seek more formalized ways to represent data (see for example, the "TERENA's Incident Object Description and Exchange Format Requirements" [RFC 3067]). The goal of this work is to ultimately provide a "common data format for the description, archiving and exchange of information about incidents between CSIRTs...including alert, incident in investigation, archiving, statistics, reporting, etc."

3.4.2.2 Analysis Depth

To what level of detail should the analysis be undertaken and what level of resources should a team expend when analyzing incidents? Analysis depth depends on a range of factors. Some examples are given in Table 14.

³⁰ Carpenter, Jeffrey J., Dunphy, Brian P. "Moving Towards the Exchange of Incident Statistical Data." 10th Annual FIRST Conference on Computer Security Incident Handling and Response, Forum of Incident Response and Security Teams, Monterrey, Mexico, June 1998.

Table 14: Analysis Depth Factors

Analysis Depth Factor	Description
Team's mission and technical capabilities	A team whose mission is to safeguard the security of their constituents will have to go to great lengths to investigate ongoing incidents in a thorough way. The team will need the technical capabilities to do so. If capabilities in a certain area are lacking, it will result in less detailed analysis. In such cases, the analysis for that area could be subcontracted. ³¹
Severity of the incident	When there is sufficient funding and staff resources available, incidents of lower priority might be investigated more often and to greater extent. On the other hand teams with limited funding or staff resources will need to be very selective about the depth of any analysis undertaken and will most likely focus on high priority incidents.
Chance of repetition	If it is likely that the intruder will strike again at another time or place, it is worthwhile spending time analyzing the incident. Investigating the incident will reduce the impact that might result from repetition of the incident by providing relevant information to constituents, other teams, and possibly also law enforcement. The analysis of such incidents may also be of use internally, keeping other team members aware of the bigger picture.
Possibility of identifying new activity	There is little point in analyzing an incident in great detail if the activities exhibited by the intruder and the tools and methods used are commonly known (there will be nothing new for the team to learn from the analysis). However, if it is suspected that the intruder is using a new method of attack or a new variant of an existing method or tool, then in-depth analysis is necessary to understand the activity.
Support from constituents	If a site reports an incident but does not provide the information needed to perform a detailed analysis, this might effectively stop any further analysis.

There is a whole range of actions that a CSIRT can take if it has the time to both analyze events thoroughly and appropriately disclose the results to its constituents and to other teams. A list of possible actions in order of increasing resource demand is

- examine log files
- examine malicious code and software environment
- provide workarounds or fixes
- actively resolve problems
- examine site security, in conjunction with the site's network ("trust") relations

Example: A team might use a commercial vulnerability scanner to actively investigate if there are no obvious holes in the site's host systems, seen from internal (intranet) and external (the Internet) perspectives. Checking the security posture of a site in this way is fairly easy to do. But such activities must have management approval and need to abide by policies and procedures of the team and site to avoid any misunderstandings that imply that the team "broke into" a site. Such activities will consume a great deal of time because the results will need to be analyzed by the team very carefully to avoid any liability for overlooked security weaknesses.

³¹ See *Outsourcing Managed Security Services* on the CERT/CC Web site at <http://www.cert.org/security-improvement/modules/omss/>.

3.4.2.3 Log-File Analysis

Every decent hardware platform and operating system and many programs (especially server type software) provide the facility for alarms and logs. Alarms are triggers designed to draw attention when some predefined (usually undesirable) event takes place, such as a packet flood. Logs are files in which events (both harmful and innocent) are recorded. Alarms ring when a specific log entry meets predefined alarm criteria.

Alarms are mainly of interest to the operators in question, whereas logs generate a wider interest, mainly because of their portability and wealth of detail. Log files can provide information such as

- who logged in when from where
- what kind of login occurred (SSH, telnet, rlogin, X, etc.)
- to what destinations email was sent
- what errors occurred

It is up to the operators of the systems involved to ensure their logs provide the necessary level of detail. Clearly an incident handling service can give advice here, during the course of an incident, but also in a preventive way by telling the constituency about good log-file practice. It is up to the CSIRTs to accept relevant logs, process them, and act on the results.

Changes within the DNS system may take place so that host names or IP addresses are no longer valid or point to different hosts. So if log files are to be of any more than just incidental use, they must display certain characteristics (see Table 15) and must be analyzed as soon as possible. In addition, the full value of the logs may only be realized when reviewed alongside the configuration files (like `/etc/syslog.conf`) of the tool that generated them.

Acceptance, receipt, and processing of log files involves some generic issues for the CSIRT to consider; these are discussed below. Additionally, all material within the premises of the CSIRT must be protected (see Section 3.8.4, “Information Storage”). Of similar importance is to carefully dispose of all material that is no longer needed or in use (see Section 3.8.5, “Information Sanitizing and Disposal”).

Table 15: Notable Characteristics of Log Files

Characteristic	Description
Timestamps	Timestamps must be present in the log for virtually every internal event recorded. Use of time-synchronizing software like NTP (Network Time Protocol) is strongly advised to avoid confusion when comparing logs from different sites (or even different machines from the same site). For the same reason, timestamps should include time-zone information.
Origin of Log	All details about the machine (Internet name, network address, machine type) that produced the log must be collected. It is also important to know what software (including version number) was used to generate the log and any associated configuration files.
Authentication of Log	Without authentication, it is not possible to say if a log file is authentic and wasn't created after the event in question or tampered with before the activity was discovered. (After all, logs are still mainly text files that anyone can produce with their favorite text editor on any computer platform.) Under some laws, it is advisable to let two people date and sign important log files (i.e., on printed versions of the log files), preferably on the same day that the log was produced. Such actions are mainly of interest if legal action is possible. Another alternative approach is to have logs written to another host machine with very controlled access.

Categorization

What category (secret or public) does the log file belong to? There should be a policy on categorization of information to be applied by the triage function, and subsequently the information should be handled appropriate to its category.

In the case of log files, it is often necessary to attach more than one category to one log. Generic information is generally less sensitive than specific information revealing machine names, network addresses, and names of employees.

Example: IP sniffer logs can commonly contain any information, ranging from not sensitive at all to explicit username/password combinations, as well as emails and their attachments.

A broad categorization of the log file must usually be done *before* the log is actually obtained, because the category (based on the apparent information in it) may impose boundary conditions on the way the team receives the log and what they can do with it.

Receiving

The log should be delivered to the team with the necessary level of due care corresponding to the category of the log information. Sensitive information must be transferred in a safe way (e.g., using encrypted channels), whereas non-sensitive information can be transferred in plain text using email. Sending the log files on disk or tape might be appropriate if larger amounts cannot be transferred encrypted over the network.

Verifying

Is the log file genuine? An authentication method should be agreed on with the party sending the report. Digital signatures provide a solution to this problem, with MD5 and RSA being popular algorithms/protocols to implement this. Whereas MD5 (a checksum algorithm) only ensures that the data received equals the data sent, RSA (a public key algorithm) helps in establishing the identity of the sending party and the authenticity of the data received. No method, however, can 100% verify that a log has not previously been tampered with by either an intruder or some other party.

Cleansing

Sensitive but irrelevant information is often best disposed of or sanitized immediately, to eradicate any possibility of disclosure.

Example: Often passwords can be removed from incoming logs immediately. Password information is seldom of much use to a CSIRT, but you don't want it to leak out. You may ask the parties involved changing all passwords, but this often takes much time, and some constituents will not even comply. So it's best to avoid unnecessary risk. The original log must remain unchanged, as it might be of use during a legal investigation.³²

Disclosing Log Extracts

If incident follow-up is undertaken and other constituents and teams are to be informed about the activity and the part of the incident that relates to them, it will usually be necessary to send them information from the log files. As a rule, complete and unabridged log files will not be sent out. Relevant extracts will be produced and sent to the parties involved, containing only those details that are specifically relevant to them.

3.4.2.4 Artifact Analysis

Intruders often leave all sorts of files on the systems that they compromise. These can range from Ethernet sniffer log files, password files, exploit scripts, source code to various programs and other tools. Generically we label these *remnant files*. We address scripts, source code (malicious code), and other programs of a potential malicious intent as *artifacts*. Some of these files may not be at all malicious, but we don't know that until they have been analyzed. The correct default assumption to make is that an artifact script or program is malicious until proven otherwise.

Intruders may have replaced ordinary files by others that differ in content from the original, but not in name. These Trojan-horse programs are popular among intruders. They are programs that seem to do everything the original program was intended to do, but that do it the wrong way—or (even worse) do what the original program was supposed to do and also

³² This is actually a very dated example, but given the still widespread use of cleartext password protocols like FTP and POP3, as well as Basic Authentication for HTTP, this example still applies.

do something else (e.g., updating the intruder on what is happening, sending sensitive information to a hidden log, or forwarding information to other external destinations).

A Trojan-horse version of a telnet daemon may log the username/password combinations that people are typing and send these logs to the intruder by email or store them somewhere on disk for the intruder to fetch. You cannot definitively identify fake programs such as Trojan horses by file attributes such as date or size. Intruders go to great pains to make the Trojan horse identical to the original program with respect to all file system attributes except content, meaning that only a proper cryptographic checksum analysis can detect a difference between files. Keeping off-line, read-only lists of checksums of system files and important programs is therefore a good practice (unless you favor reinstalling your systems from scratch after even a minor intrusion). Programs such as Tripwire³³ and anti-virus programs such as ThunderByte make lists of this type part of their routine and inform you when the checksum on a file has changed.

Whether or not the CSIRT should analyze malicious code as part of their incident handling service (or as a separate service) is an important question to answer. Various CSIRTs have differing views. The following examples are from different extremes.

Example: Some teams do not have the resident expertise to analyze malicious code and must rely on other CSIRTs or experts to provide such analysis.

Example: Commercial teams who have taken on the job of securing a site's network will usually fully investigate the matter, including analysis of malicious code if required.

No matter who performs the task, proper analysis of malicious code should be addressed to some extent. How else is one going to derive an intruder's fingerprint, which may help in analyzing other incidents? How else does one know in what directions to seek further, if not by actually observing what the code tries to do? Just eradicating all artifacts and building the system from scratch is a very expensive solution to an intrusion. And it is often a very naive one, especially if the flaw that enabled the intrusion in the first place has not been removed. Artifact analysis can help by understanding what the artifact does, for "inside the artifacts lurks the intruder."

Assuming the CSIRT takes the responsibility of analyzing malicious code as part of the incident handling service, the following points should be considered.

Where to analyze the artifacts?

Usually the malicious code will not be analyzed on the victim's systems. The constituent will

³³ Tripwire[®] is available commercially from <http://www.tripwire.com/>. A publicly available version is available from <ftp://ftp.cerias.purdue.edu/pub/tools/unix/ids/tripwire/>.

want to return to normal operation as soon as possible, and you don't want to further endanger the constituent's environment. From the intruder's point of view, what is simpler than writing a piece of malicious code that attempts to destroy the information on the hard disk if the code is invoked in the wrong way?

Care should be taken to make a copy of the artifact(s) plus surroundings that, where possible, exactly mirror the original environment. This should not be performed by the constituent just sending some files, but by a member of the CSIRT or some other informed and authorized individual. This means the constituent might grant the team member temporary access to the system involved, or undertake the task themselves using instructions and guidance provided by the CSIRT staff.

Ideally the artifacts are analyzed in an isolated laboratory, isolated in a networking sense. Test computer equipment and a test network environment should be available for artifact analysis. Also, total loss of the test environment's data should be of no consequence. If the test environment must be accessible from the outside for practical reasons, this should only be possible through a very restrictive firewall.

Example: Several response teams tested a flaw in INND (a netnews daemon) in production (e.g., not isolated) environments. This resulted in News "control messages" escaping from their test systems and doing what they are supposed to do inside NNTP,³⁴ i.e., spread all over the world. Unfortunately these control messages exploited the INND flaw in such a way that /etc/passwd files were sent to specified email addresses (in this case the CSIRT teams doing the testing). Thousands of such messages were received. Had the teams used appropriate isolated laboratory environments, this would not have occurred.

When to involve expert groups?

Given the size of the average CSIRT, it is most likely impossible for each team to know everything about every operating system version and network protocol. Therefore it is often advisable and even desirable to share the analysis process with some other group of experts. Because of the sensitivity of the work, and depending on the nature of the CSIRT and its constituents, these experts should be identified in advance and appropriate precautions taken (e.g., screening or non-disclosure agreements) before any information or analysis is shared or undertaken.

Example: CERT-NL has an expert group associated with it. This is a voluntary effort by experts from the CERT-NL constituency. The experts benefit by receiving new information first-hand. CERT-NL benefits by obtaining feedback from the experts.

³⁴ The News protocol.

Example: There are times when more mature teams cooperate when performing analysis (mostly vulnerability analysis or artifact analysis), with one team taking the lead and the others contributing as “experts.”

When to stop?

Criteria should be defined in advance for determining the limit of depth and breadth of the analysis before it is stopped or transferred to another entity (e.g., a separate artifact analysis service). Such bounds can be as trivial as a limit on the amount of staff effort spent, or they can be based on an evolving assessment of the problem’s complexity.

3.4.2.5 Analysis of Software Environment

Just analyzing Trojan-horse programs, Ethernet sniffer logs, and exploit scripts (i.e., artifact analysis) is not enough. Study of the environment in which these remnants were found is of equal importance to solve the puzzle and see the bigger picture. Take for instance an exploit script. The success of such a script is determined by its surroundings: the shell environment, the software present, available privileges, and so forth.

Operational systems are made up of tens of active programs and drivers and hundreds of ready-to-run programs. The file system is usually complex, with rights distributed in a semi-random way to numerous users and groups. An exploit script itself is relatively easy to analyze, since it tells us what it does, although it may contain provisions for random sequences of events or may be written in a way that makes understanding its actions difficult.

An exploit executable is altogether different, as analysis of its activity can only be fully understood when it runs. Exploits that make use of race conditions (unforeseen states of running code with undefined outcome) don’t make things easier, for they may only be reproducible if the test laboratory situation is a very good mirror of the original software environment.

The analysis of the software environment is firmly tied to the artifact analysis, so firmly that one cannot be separated from the other. As a result, most of the content of the previous section on artifact analysis also applies here. Essentially:

- Whoever performs the artifact analysis (constituent, incident handling service, or separate artifact analysis service) should also undertake the analysis of the software environment.
- Obtaining and analyzing the artifacts also means obtaining and mirroring the original environment as genuinely as possible. Preferably one should use the same operating system versions, patch levels, drivers, and configuration files, etc. This requirement illustrates how involved and complex artifact analysis can be. What one would really like to do is perform the analysis in the original surroundings; but as indicated before, this is seldom feasible, since understandably constituents will usually refuse to act as guinea

pigs. The risks for them are too great. Some constituents, however, might be prepared or able to participate in such analyses (for example, in an academic environment where skilled technical staff members are available and easily accessible, isolated test equipment may be available, and there is time and interest in the task).

The analysis of artifacts and the associated software environment may unveil known vulnerabilities in specific software (in a specific environment). The victim can then be helped with appropriate advice [Garfinkel 1996] and in cases of widespread exploitation, a “heads-up” can be sent to constituents and other CSIRTs. On the other hand, the analysis may unveil as yet unknown (or at least unpatched) vulnerabilities. This problem should then be transferred to a vulnerability handling service. That service might exist within the CSIRT or external to the team (e.g., trusted experts, colleagues, or vendor teams). In the ideal case, the vulnerability would then be promptly patched and everyone affected would be appropriately informed.

3.4.2.6 Intra-Incident Web-of-Relations Analysis

Advanced intruders usually weave a whole web of connections over the Internet, using a set of their favorite vulnerabilities to gain access to systems, and then using those compromised systems as “stepping stones” to attack yet other systems. Intruders make the web complex to evade detection and apprehension. If the center of the web can be identified (such as the system compromised as the first one in the chain of intrusions), then it may become possible to locate or identify the perpetrator.

Example: If an intruder uses the telephone system to gain access into the Internet via one of the public online services, the telephone company may help with tracking down the intruder’s telephone number used to make the connection. (Usually this means law enforcement must be involved.)

Example: If an intruder is spinning their web from public terminal rooms (at a university, for instance, or at an Internet cafe), one needs to catch the offender in the act. When they next resume their activity, the location of their machine can usually be tracked through its IP number within an organization. If the Internet cafe by itself does not try to detect attacks, they will not be able to identify such abuse of their infrastructure.

One has to take great care when trying to locate an intruder. During the process, the intruder and investigator may both make use of the same systems. Often the intruder has the highest privileged access on these systems (i.e., UNIX root privileges) and may be alerted to or undermine the investigator’s activities [Stoll 1989, Shimomura 1995].

The analysis of the web-of-relations inside an incident is of great importance to help contain an incident. The more one understands an intruder’s operations and relations, the easier it becomes to counteract their activity and help prevent others from becoming new victims, and finally may even enable capture of the perpetrator.

When performing this analysis, one should trace the relations that appear inside log files, and keep track of the intruder's signature.

Trace log-file relations

Log files or parts of log files associated with the intrusion (e.g., telnet logs of the intruder's activity, sniffer logs) should be carefully examined and every link to other systems should be investigated. Usually this means involving constituents and other CSIRTs on a need-to-know basis, providing them with only the portions of the logs relevant to them. It is, however, advisable to alert your fellow CSIRTs (or at least those teams with whom you have a sound relationship) and give them information about the way that the intruder seems to perpetrate his attack(s). This will help the other teams to be more proactive and to recognize similar activity when it occurs. Ideally, the other teams would be expected to return the favor and inform you if the situation were reversed.

When the search yields new log files with relations, these need to be similarly analyzed. This can be quite a time-consuming job. This is especially true in incidents involving Ethernet sniffer logs. Such incidents have caused several CSIRTs a tremendous amount of effort tracking and notifying all affected sites of possibly compromised systems.

Keep track of the intruder's signature

Throughout the incident, you (or others involved by you) need to make sure the intruder's signature is abstracted and compared with the signature as you know it. The signature is the way the intruders go about their work, the scripts used, the passwords tried, the programs they attempt to break, the vulnerabilities exploited, and the file or subdirectory names they invent or use to store their tools.

Keeping track of this signature will enhance your understanding of the intruder. You may also find instances where the signature seems to be totally different. Your intruder might either be very creative or this might be the signature of another intruder whose path you have crossed by accident. Alternatively the signature could be the result of intruders working together and sharing information. By following one intruder's trail, you might start finding their colleagues' trails too. These trails could show both remarkable similarities and clear differences.

3.4.2.7 Analysis of the Texture of Ongoing Incidents

Not only should all relations *within* every incident be investigated, but also separate incidents should be compared with one another (this adds to understanding the "bigger picture" mentioned earlier in this section). The same two aspects stand out again as main quantities to evaluate during any analysis activities: intruder's signatures and the web-of-relations.

Because the analysis of a seemingly coherent incident may show that another intruder's tracks are confusing the proceedings, analysis of the texture of incidents may show that separately treated incidents may well belong together. The similarities may result from the same intruder with the same signature or from a group of intruders apparently working together with a similar or almost identical signature and web-of-relations; alternatively, there might be completely different, unrelated attacks against the same target. Only through the analysis process will such understanding be obtained.

3.4.3 Tracking Incident Information

During the life cycle of every incident, it is very important to track information pertaining to the incident at varying levels of detail. This facilitates the organization of information and effective response to the incident. This recording of information in a logical, organized way will also provide a historical record of reported activity and all actions taken to assist in the distribution and allocation of incident workload. This record can also provide statistical and trend information that may be used within the handling function or by other functions or services; for example, management or sponsor reports, customer quality assurance, staff performance measurements, etc.

The level of detail recorded may vary from team to team based on their specific requirements, the level of incident handling service provided, and the depth of analysis undertaken. For every incident, a good practice is to capture and track at a minimum the information detailed in Table 16.

Depending on individual policies, teams might store online incident information only for a short period of time while it might still be needed, such as a few weeks after closing an incident or to generate regular statistical information. Usually the information is kept at least a little longer, to allow the possibility of an incident reopening. If the tools support the reopening of incidents, this is a must. In some cases CSIRTs have reported incidents reopening a year (or longer) after the original report! Such cases are mainly due to sites failing to regularly review logs for incident activity, but can also be the result of a law enforcement case finally reaching a trial date. Depending on the nature of the team and services provided, there may be no need to store the whole set of collected incident information indefinitely, although it might be helpful for historical purposes. This topic is discussed in more detail in Section 3.8, "Information Handling."

Table 16: Incident Tracking Information

Information to be Tracked	Description
Local CSIRT's unique incident tracking number	Unique tracking number supplied by this team. This is used to track all information and actions relating to the incident.
Other CSIRTs' incident tracking numbers	Tracking numbers assigned by other teams involved. This facilitates the appropriate coordination of this incident across other teams.
Keywords or categorizations	Information to characterize the incident and help establish relationships between different incidents. This information may change during the incident life cycle as new information becomes available.
Contact information	Names, phone numbers, email addresses and other contact information for all parties involved in the incident. It should include details of preferred encryption methods and associated keys.
Policies	Legal parameters or policies that affect the way the incident might be handled.
Priority	Priority of the incident according to the CSIRT's priority scheme. An incident's priority often changes during its life cycle.
Other materials	Location of other materials associated with this incident, such as log files or hard-copy materials.
Incident history	Chronicle of all email and other correspondence (e.g., details of telephone conversations, faxes) associated with the incident.
Status	Current status of the incident.
Actions	List of past, current, and future actions to be taken in respect to this incident. Each action should be assigned to a specific team member. These actions, as appropriate, might also capture completion dates and other deadlines.
Incident coordinator	A team may choose to assign a staff member to coordinate the response to this incident. This person might not always be available, which raises its own problems, but with one person seeing all information related to the incident a better picture can be built.
Quality assurance parameters	Information that might help to measure the quality of the service. References to service level agreements that might affect the handling of this incident.
Textual description	A free-form description that accommodates other information not covered in any other tracking field.

3.5 Announcement Function

As previously stated in Section 3.2, "Service Functions Overview," the announcement function generates information tailored for the constituency in various formats. The purpose of announcements varies from disclosing details of ongoing threats and, steps that can be taken to protect against those threats, to sanitized trend information on the scope and nature of recent attacks reported to the team. For the purposes of this handbook, the scope of this function is limited to its direct applicability with the incident handling service. For a CSIRT providing a broader range of services, however, publishing announcements can be considered as a service in its own right and would likely offer a much broader range of information derived from other services, such as vulnerability or artifact analysis.

Example: Announcements such as CERT Advisories [CERT/CC 1988] provide information on preventing threats that are generally discovered as the result of incident

reports, vulnerability reporting, and testing based on patches, as well as security updates made available by the vendor community.

Since the formation of the first CSIRT, announcements to teams' constituencies have generally, in some form, been part of a CSIRT's daily business, whether as a distinct service or as part of the incident handling service. As previously mentioned, however, this announcement function is optional, as it is not critical to the provision of a basic incident handling service. The main objectives for announcements are to disclose information to the constituency, to assist them in protecting their systems, or to look for possible signs of attack by providing notification of potential, current, or recent threats; and further to suggest methods to detect, recover from, or prevent threats. When disclosing information related to a specific attack type, care should be taken to ensure that the level of disclosure is sufficient to allow recipients to understand the threat and check for it, but not detailed enough to enable the information to be used to implement the attack. This is the most challenging task of the announcement function.

A list of announcement types and a discussion of the announcement life cycle follows in Sections 3.5.1 through 3.5.3. Other issues that should be considered when making announcements to a constituency are covered more generally in Section 3.8.8, "Information Disclosure."

3.5.1 Announcement Types

Announcements can take on many forms, from those providing short-term information related to a specific type of ongoing activity to general long-term information for improving awareness and system security. Each has its own tradeoffs and benefits.

Heads-up

The heads up process usually takes the form of a short message, issued when detailed information is unavailable. The purpose is to inform the constituency or other parties of something that is likely to be important in the near future. Announcing a heads-up has multiple benefits. First, the CSIRT can proactively warn or inform their constituency to a potential issue or threat. Second, the recipients may already know something about (or have additional information relating to) an issue detailed in the heads-up than the CSIRT. This gives the constituency the opportunity to provide feedback to the team. Third, the recipients may stumble on information related to the content of the heads-up at some later time. They will then be in a better position to recognize the information and its potential importance. There is a caveat, however, that information in such documents is likely (and often expected) to change, so it might be worth including a disclaimer prominently in the text of the "heads up" to clearly identify when the information is unconfirmed or speculative.

Alert

Alerts are short-term notices about critical developments containing time-sensitive information about recent attacks, successful break-ins, or new vulnerabilities. There may already be complete information regarding the subject of an alert, but something may have changed to require the publication of new information. Examples are documents such as the CERT Summary on the “named” problem [CERT/CC 1998b], a variety of CERT Incident Notes and Vulnerability Notes [CERT/CC 1998c, CERT/CC 1998d], and other more recent similar alerts, such as the CERT Current Activity page at http://www.cert.org/current/current_activity.html.

Advisory

Advisories are often one of the most common documents produced by CSIRTs.³⁵ Advisories provide mid-term and long-term information about problems and solutions suitable to raise awareness and help avoid incidents. They typically contain information about new vulnerabilities, but may also contain information about intruder activity. Advisories are often well-researched and include substantial technical detail relating to patches and workarounds [Cormack 2002].³⁶ Advisories are typically directed at a technical audience such as system and network administrators, but sometimes contain additional background information for less technical readers. Examples are CERT Advisories [CERT/CC 1988].

For Your Information

These are documents that contain mid-term and long-term information, similar to advisories, but shorter and less technical to address a wider audience. These might be called briefs, bulletins, or newsletters as well. Such documents typically contain information of a tutorial or instructive nature that can be used by non-technical personnel with an interest in security. This could include management or legal staff and members of the press. An example is CIAC C-Notes (originally called CIAC Notes) [CIAC 1994].

Guideline

A guideline is a sequence of steps that lead someone familiar with the basics of his craft through a process meant to expand that person’s knowledge or even to work direct improvements (in system or network security). They can be lengthy documents aimed at helping technical staff improve their fundamental understanding of security and their day-to-day practices [CERT/CC 2000]. Other examples include the Site Security Handbook [RFC 2196] and CERT Security Improvement Modules [CERT/CC 1997c].

Technical Procedure

Technical procedures are more lengthy than guidelines, with more technical details

³⁵ For example, the CERT Advisory is probably the most common way in which people recognize the work done by the CERT/CC.

³⁶ McMillan, Robert D. “Vulnerability/Advisory Processes.” 8th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, San Jose, Calif., July 1996.

addressing an expert audience and often targeted at a specific problem domain. Examples of this are the CERT Tech Tips such as the “Problems with The FTP PORT Command” [CERT/CC 1998e] or the Security Improvement module “Security Public Web Servers” [Kossakowski 2000].

3.5.2 A Priori Considerations

Defining a set of announcement types is just the first step towards a comprehensive announcement process. Several other factors need to be considered and addressed before such announcements are issued. These factors range from the criteria that trigger an announcement to how it will be distributed, and these are discussed in further detail below.

3.5.2.1 Announcement Triggers

Criteria need to be in place to determine what will trigger the development and distribution of each type of announcement. These criteria could be anything from just another team’s information to identifying a surge in current attacks being reported to the team. Obviously the information required to meet the criteria must be tracked and monitored regularly. Usually the information source is either the CSIRT itself, based on the activities reported to the team, research done by the team, or the source of some other team’s announcement.

3.5.2.2 Categorization Criteria

It is useful to derive criteria to help categorize announcement material, that is, help to pick the right type of announcement for that material. Criteria based on the source of the material are not hard to define; criteria based on content type are much harder.

Examples: Material that is derived from public mailing lists such as Bugtraq may cause a heads-up but certainly not an advisory, unless the content is double-checked and/or verified through other investigation (source-based criteria). Likewise, if a CSIRT does not generally handle virus issues, a new surge in viruses is not likely to cause an advisory or guideline, but it may yield an alert or heads-up (criteria based on content type).

When categorizing the content of material, the target audience for the announcement must also be considered.

Example: A very technical description of a Sendmail exploitation may well trigger a very technically detailed advisory aimed at experienced system administrators, not the general “user” population. A less-detailed technical description of a problem in some popular Web browser might better result in a “for your information” aimed at a much wider audience.

3.5.2.3 Prioritization

Several (more or less subjective) factors will affect the perceived importance of each individual announcement. Care should be taken to preassign objective priorities to each announcement based not only on specific type of announcement but also on its content (i.e., a handful of broad topics such as denial-of-service attacks or viruses). CSIRTs need to use the right vehicle (e.g., announcement type) to get the appropriate information disseminated to the right audience(s).

Example: Originally the CERT/CC had just one method of announcing information publicly to its constituency: the CERT Advisory. This document type was the mechanism by which CERT staff disseminated important information that system and network administrators should “pay attention to.” This proved effective for a while; however, clearly some issues are more important than others. Having just one publication vehicle to disseminate many different types of information can dilute that publication vehicle (e.g., everything cannot be the highest priority; true for incidents and also for announcements).

Over time, the CERT/CC developed other notification schemes, such as the CERT Summary, CERT Incident Notes, CERT Vulnerability Notes, and CERT Current Activity. In this way, the CERT/CC uses the Advisory announcement type for the most important issues that CERT staff want the constituency (e.g., system and network administrators) to immediately review and address as appropriate to their particular situation. Other CERT/CC publication vehicles are referenced in Section 3.5.1.

3.5.2.4 Clearance of Information in Announcements

According to the CSIRT’s policies and procedures governing the disclosure of information, the information intended for use in the announcement must be cleared for disclosure at the appropriate level (whether for public or restricted distribution). Some general clearance rules should be set beforehand to help this process run smoothly in practice.

Examples: An obvious clearance rule for a public announcement would be that it may never contain details about individuals or individual sites. Another such rule is that if the information going out is based on or simply a redistribution of materials provided by other teams, appropriate permission must be obtained from those teams, and the appropriate attributions included.

3.5.2.5 Distribution Channels

Depending on the announcement type, different issues need to be considered when choosing appropriate distribution channels:

- sensitivity of information: Is the distribution channel safe enough?
- audience (constituency) addressed: Is the channel adequate to reach this audience?

- speed: Is the channel fast enough?
- cost: Is the expected result in producing and distributing the announcement worth the money (e.g., time, effort, materials, value)?

Other issues relating to distribution channels include the manner in which the information will be distributed or disseminated. This is discussed in Section 3.5.3.5, “Distribution.” Whatever mechanisms are considered appropriate, they should be set up and tested in advance and then advertised to the constituency.

3.5.3 Announcement Life Cycle

Having decided on appropriate announcement styles and initial criteria, the next step is to define processes and procedures to handle the actual generation of announcements (type, form, style, etc.). In general, the five phases described in this section can be recognized in the life cycle of an announcement, ranging from the first evidence for its need to its ultimate distribution.

3.5.3.1 Initiation

When possible announcement material is identified (e.g., during incident analysis or through observing likely sources such as mailing lists), a determination must be made as to whether the material meets the general criteria referenced in Section 3.5.2, “A Priori Considerations,” or is otherwise important enough to be announced. If so, the type of announcement, the content type, and the intended audience will have to be made explicit, and it is helpful if an internal tracking number is allocated. This facilitates the tracking/recording of the development of the announcement.

The priority or importance of the topic, announcement type, content type, and audience determine the following important parameters:

- style and detail in which the announcement will be written
- information clearance measures to be taken
- distribution channel to be used

In addition, other issues that need consideration or decisions at this point include the proposed time schedule, internal resources needed, responsibility for the task, and other aspects such as collaboration with other parties (to provide content or improve the quality of the announcement, or to synchronize publication of announcements with other collaborators).

3.5.3.2 Internal Prioritization of the Announcement

This phase may be revisited regularly for all announcements currently under development. It is possible for some teams to have a number of concurrent efforts underway for

announcements (in various stages of development): yet unissued announcements are prioritized based on the predefined criteria and other relevant criteria determined at the time the content begins to be developed. Certain types of announcements (by their very nature) might have the highest priority if they are time critical, and can supercede other important announcements already in the queue. Others that simply provide general statistical information may be of the lowest priority and deferred to a later date for announcement. Relative priorities might not be immediately obvious when two comparably important announcements are competing for resources. In such cases it is best to prioritize based on the severity of the issue, e.g., threats addressed and/or the size of the constituencies involved.

3.5.3.3 Development

This phase is composed of the technical description, editing, and overall writing of the announcement. Most teams generate a standard template for each type of announcement that indicates the appropriate layout and content of the material. This “boilerplate” can greatly facilitate the development of the material, and it also helps to maintain a consistent look to the announcement form and so sets the constituency’s expectations for what the content will cover. Drafts of the announcement may then be provided for internal and limited external review to obtain detailed comments from experts not directly responsible for developing the announcement. When providing other parties with this information, any restrictions to use must be made apparent.

Example: An early draft announcement may be sent out to all FIRST teams to seek their review and comment, but is labeled “not for further distribution.” To further protect it against attacks or possible disclosure, drafts and comments should be encrypted and digitally signed when distributed for such review and comment.

3.5.3.4 Final Preparation

Prior to final dissemination, some technical issues still need to be addressed, such as generating cryptographic checksums of the announcement itself or items that it references. However, most of the issues that remain at this stage are non-technical. They are usually concerned with the overall presentation and content (i.e., dates, headers, footers, acknowledgments, and disclaimers). Before releasing the announcement, the team must make sure that all the references it contains are valid (i.e., URLs and patch files are correct and accessible).

Another matter for consideration is whether or not it is appropriate to offer an advance distribution of the announcement to a limited audience, such as fellow CSIRTs or your team’s media contact. This gives such parties the opportunity to prepare for any responses they may receive. For a fellow CSIRT, this might mean preparation of an announcement of their own, e.g., if they need to translate the information to another language or will simply “wrap” your announcement with their own boilerplate. In such cases it is appropriate to retain distribution

restrictions on the announcement until the moment of ultimate disclosure, in case there are any last-moment changes required.

Example: Some CSIRTs send out encrypted final drafts of their advisories to all FIRST teams with the proviso that those teams can use the information to prepare their own announcements, but they are not free to further distribute the information until a final public distribution version has been made available. As a result, any possible conflicts of interest can be minimized. If one of the reviewers disclosed the information prior to the final public distribution date, this would damage the whole process, as the information would be leaked while others believed that it was still confidential.

Team members providing triage, feedback, or other incident handling functions need to be prepared to copy or address any possible responses to an announcement from the constituency, media, or other parties. Advance briefings on all announcements should be provided for these team members.

Finally, every outgoing announcement should be allocated an external tracking number, which usually is serially allocated per announcement type. Then the announcement should also be digitally signed to protect against tampering.

Example: Every CERT Summary distributed by the CERT/CC has a number of the form CS-YYYY-XX (where YYYY is the year in which the summary was issued and XX the number sequentially allocated from 01 for each Summary issued in that year). Authenticity and integrity is provided by a digital signature generated with PGP (or GPG). This is also true for other CERT documents (e.g., CERT Advisories, Incident Notes, and Vulnerabilities Notes).

3.5.3.5 Distribution

This activity is related to the effort involved to distribute or disseminate the final announcement via the distribution mechanisms that the team advertises for announcements of that type. This might include placing the announcement on appropriate information servers such as the team's FTP or Web server, or distributing it via other mechanisms such as mailing lists, automated fax distribution, or news mechanisms.

Example: The CERT/CC issues many of its announcements (such as CERT Advisories and CERT Summaries) to a mailing list known as the cert-advisory mailing list, and to the USENET newsgroup comp.security.announce. The latter is moderated by CERT staff and intended solely for the use of CERT announcements. In addition, the CERT/CC archives all announcements (including those not directly disclosed via the mailing list and newsgroup) on its Web server.³⁷

³⁷ <http://www.cert.org/>

3.6 Feedback Function

Providing support for recovering from and dealing with incidents is the major objective of most CSIRTs. Being effective in this role will lead to other requests and issues being directed at the team that are not necessarily specific to the incident handling activity (or perhaps requests that are not even part of the CSIRT mission or services provided to the constituency). Unfortunately, simply ignoring such requests and issues (even if they are beyond the scope of your CSIRT's roles and responsibilities) can affect the team's reputation and the attitude of the constituency toward the team. Hence, it is in the interest of all CSIRTs to provide some level of appropriate feedback at some minimal level to such requests.

On the other hand, from a knowledge management perspective, the type of incoming requests received by the CSIRT can provide some insight into the current needs of the constituency and other interest in the team. As a result, providing feedback to such requests can help to deliver a better service and at the same time clarify the expectations of the constituency instead of ignoring obvious problems and misconceptions. The CSIRT should strive to answer requests, no matter what the request (even if it is a reply saying "We cannot respond to this request, it is beyond the services we provide," or pointing to other information resources).

Example: If a team does not reply to questions directed at it, the requester may think of the team as unhelpful or unable to help. Other requesters might think the team is arrogant or worse (mismanaged, uninformed, etc.). To avoid this perception, the team should at least provide a statement of the purpose of the team and why no further feedback is provided. Keep in mind that some of these requests can be the result of "investigative" journalism (seeking to elicit information that may not be public), and can come in many shapes and sizes.

Incoming requests commonly fall into one of four categories:

1. **general computer security requests**

Such requests commonly seek information on avoiding incidents through proactive security measures or how to interact with the CSIRT if an incident should occur. As CSIRTs regularly deal with incidents, they have the knowledge to provide this type of information. Therefore it seems natural for people to direct questions of this type to the team. Whenever possible, a team should make use of such opportunities to proactively help the constituency raise awareness, avoid (or limit) incidents, and improve the overall security.

2. **media requests**

These are requests from members of the media who may be seeking input for a story relating to a general security article, announcement, or specific incident. Whenever possible, the CSIRT should be prepared to appropriately handle such requests from the media, while ensuring that the team's information disclosure policy is not violated.

3. **other requests and issues**

There is a whole range of other requests and issues that constituents may submit or that a team might wish to provide feedback on. These include requests for the team to provide a speaker at a conference or a request for permission to make use of copyrighted material available from the team. Handling such requests may help promote awareness of the team and, when feasible and appropriate, should not be ignored. Sending out the CSIRT's annual reports can be placed in this general category, even though that is a matter of disseminating information proactively rather than responding to a request.

4. **out-of-scope requests**

These requests have nothing to do with the services provided by the CSIRT, but even so, as mentioned above, a simple acknowledgment with a reference to some Frequently Asked Questions (FAQ) or policy statement on how to deal with out-of-scope requests is more useful than just ignoring the request.

Example: Typical real-life examples that are obviously out-of-scope are: How do I connect to the Internet? Do you have the postal address for my old friend in Hamburg, Germany? I need a pen pal. Others may also be beyond the scope of the offered services of the CSIRT, such as, Will you perform penetration tests on my network? Which operating system version|application|software tool should I use?

3.6.1.1 Life Cycle

Teams may choose to track each different type of request with different types of tracking numbers. Or they may track all requests with a single type of tracking number and document the different type of request made or the nature of the response given for each request.

Requests have a life cycle similar to those of incidents. However, it is uncommon for these requests to remain open for very long after the initial response from the team, although some may result in further dialogue.

3.6.1.2 FAQ and Other Default Feedback

Responses to requests can be handled individually, but this is often time consuming and the CSIRT may not be able to assign dedicated resources to this work. Most teams choose to develop previously prepared documents such as a team FAQ document, as mentioned earlier. Such a document answers some general questions about the team, gives details of the incident handling service provided by the team, and tells how to access documents that address specific needs tailored for the constituency. Once such documents are available, most requests can be handled by providing pointers to or copies of the appropriate document(s). Even in the case of out-of-scope requests, the team's FAQ might be an appropriate response if it outlines the services provided by the team and states that all other requests are inappropriate. On the other hand, a simple standard reply could be developed that politely indicates to the requester that the CSIRT does not have (and so cannot provide) the information being sought.

Example: A newly formed CSIRT might initially choose to track and respond to requests from their constituency to identify common types of requests received and then develop their FAQ to quickly handle such general (or similar) requests in the future. Over time, such an FAQ can be updated and expanded, placed on the team's Web site, and/or provided in response to future requests. The CSIRT Development Team's FAQ evolved in this way [CERT/CC 2002b].

For media requests, depending on its policies, a team might use an existing organizational media office or interact with the media through a team member or associate experienced in interacting with the media. Once the team has established a policy for where to direct media contact requests, all media requests should be handled according to that policy, and no additional support should be provided to the media through the feedback function. Depending on the team's policy, it might be appropriate to provide the media with publicly available information about the team, such as the team's FAQ.

3.6.1.3 Organization of Feedback Function

If standard responses are available, technical staff without detailed technical knowledge or a direct Web interface can be used to provide feedback to common enquiries within this function. Other alternatives might be to point the requester to other sources, such as online archives of technical guidelines made available by other teams, or to other technical experts.

An internal FAQ (or other guidance) for the team members that describes the procedures related to handling the various types of requests is also very beneficial, ensuring and enabling a consistent reaction to all requests. Such a document should also detail how to prioritize requests. For instance, requests from sponsors might obtain the highest priority, followed by requests from constituents. Or a team might choose to prioritize on the type of request rather than the requester. This internal FAQ can also be a learning tool for new CSIRT staff members.

3.7 Interactions

Throughout the incident life cycle, most of the activities of a CSIRT involve interactions with other parties. Due to the importance and implications of such interactions, great care must be taken to establish contacts with the "right" parties (see Section 3.7.1, "Points of Contact"). For the majority of interactions (i.e., communications) it is equally important to ensure authenticity (Section 3.7.2, "Authentication") and preserve confidentiality (Section 3.7.3, "Secure Communication"). This section concludes with an outline of the items to consider concerning particularly important interactions such as those with constituents, other teams, and law enforcement (Section 3.7.4, "Special Considerations"). The examples below illustrate effects that can have a negative impact on the site and the CSIRT.

Example: Say that an incident is in progress. Someone calls a CSIRT and claims to be an administrator at site A. The CSIRT provides technical details of the incident and appropriate technical solutions. The next morning there is a headline revealing the identity of victim site A, together with a detailed report about the incident. It turned out that a journalist heard rumors about the incident and tricked the team into giving out the information.

Example: Unencrypted email messages between a CSIRT and site A are monitored and copied by a third party during storage on an Internet mail host. Later the email is distributed to Internet newsgroups to a large number of readers.

On the other hand, effective communications can also have a very positive impact, as in the following example:

Example: In February 2002, the CERT/CC published CERT Advisory CA-2002-03 on “Multiple Vulnerabilities in Many Implementations of the Simple Network Management Protocol (SNMP).” Throughout the development and preparation of this advisory, a large number of new vendor contacts had to be established to ensure the “right” contacts were identified and contacted. This advisory was, by far, one of the most complex advisories published to date; involving coordination, secure communications, and interactions with a very large number of vendors (over 100) and other experts. Updates to the vendor information continued to be made almost a year later as new information was obtained from the vendor community.³⁸

3.7.1 Points of Contact

During the course of any incident, contacts are established as necessary. To establish the “right” or appropriate contacts; however, is an art in itself. It is important to pass information on, but more important is finding the person best suited to handle and receive the information and/or the person authorized to take actions or make any necessary decisions. Therefore, establishing and maintaining good contacts must be an ongoing effort of the CSIRT, with the intention of building a web-of-trust to meet the needs of the incident handling process.

For our purposes, contacts can be considered broadly in two categories: incident-related and non-incident related contacts. These are discussed in more detail below.

3.7.1.1 Incident-Related Contacts

These are the contacts that a CSIRT will need when handling a specific incident. They can include contacts within and external to an organization experiencing an incident. Examples of such contacts are

³⁸ The printed version of this advisory from the CERT Web site is over 50 pages long.

- upper management (managers, department/division/bureau heads)
- sponsors
- other departments
- technical (system and network) administrators
- security officer
- legal counsel or legal compliance department
- internal audit department
- risk management group
- network operation center
- network information center

In large organizations there may be a predetermined initial point of contact (POC) that the CSIRT notifies concerning an incident report occurring at that particular site. However, it may be essential for the CSIRT to then be placed in contact with a specific department or appropriate individual(s) who can respond to the activity. Without direct contact to the appropriate technical or management level staff, the CSIRT may waste precious time and resources.

3.7.1.2 Non-Incident-Related Contacts

Non-incident-related contacts can be used to provide background information for (or about) the team, help it to fulfill its service, potentially support the team's operation, or be used for obtaining input from domain experts. The following list provides example categories for some of these non-incident-related contacts that could be considered when generating a contact database. Examples are

- (constituency) site security contacts
- other constituency site contacts (like management, physical security, human resources)
- sites external to constituency
- Internet service providers
- other CSIRTs
- law enforcement, legal counsel
- vendors
- experts
- media

Constituent Site Contacts

As previously mentioned, there is a very good reason for maintaining different kinds of

contacts within one organization: the possible need for escalation. While it is usually acceptable to handle an incident in cooperation with one single department, upper management should be involved when it is obvious that an incident has consequences that require management authority or oversight or an incident that crosses multiple divisions/departments.

3.7.1.3 Finding Contacts

Finding the right contacts for organizations is not always a simple task. For non-critical contacts, one can use publicly available resources, like telephone directories or similar services available on CD-ROM or through a search on the Internet.

Whenever a critical decision must be based on a contact, using the wrong contact may result in leakage or disclosure of critical information to inappropriate parties or (usually worse) to outsiders. It also demonstrates a lack of control and attention to detail within the CSIRT, which is bad for its reputation.

To keep the confidence of the constituency, great care must be taken to identify (vet) and use the correct contacts. If publicly available contact information can be forged, manipulated, or corrupted (a potentially very real threat, with the risk varying from printed media, CDs, to network directory servers), it should be verified before use. Better still and always preferable is to obtain the contact information directly from the source, from the contacts themselves or their management (or designated representatives).

3.7.1.4 Maintaining Contacts

This is a seemingly simple task, but in reality, maintaining contact information can be a more daunting challenge than finding it. Contact information becomes (partially) obsolete when people leave an organization, are promoted, are reassigned to other types of work, or just relocate to another desk/office/building with another telephone number. One can ask contacts (e.g., constituent sites) to pass on information relating to these kinds of changes; however, the reality is that this rarely happens. For non-critical contacts, it is best just to accept some potential for outdated or incorrect information in the database and correct the information when it becomes known. For critical contacts, this is less appropriate, and regular check-ups on a periodic basis (quarterly, semi-annually, annually, when contracts change, etc.) can help address this problem, in addition to asking the contacts to provide changes and updated information.

Example: CERT-NL demands of each of its constituents that management appoints a “site security contact” (SSC) and relates the contact information (and any changes pertaining to it) to CERT-NL. For practical reasons, it is even advised that the constituents create generic email addresses of the form `ssc@some.site.nl` for their site security contacts.

The local administrator is then responsible for maintaining the email address. This makes

the relay of information possible without prior involvement of CERT-NL. Furthermore, CERT-NL advises its constituents to create “security entry points,” with an email address of the form `sep@some.site.nl`. This security entry point is like a local CSIRT intended to handle incidents and other security issues in real-time, separate from the site security contact, who may be on holiday or ill.

Example: Other CSIRTs advocate that at the least the RFC-defined standard email addresses for “postmaster” and “security” for each domain are provided.

3.7.2 Authentication

An important aspect when interacting with others is authenticity. This term usually applies to ensuring that someone is really the person she/he claims to be. By using technical communication facilities, it is inherently more difficult to check the authenticity of a caller or called person. Therefore great care must be taken. Information that must be protected should be revealed only after the caller or called person has been authenticated and the other party is authorized to access the information. As this information might become important later, each contact and its origin must be logged.

To know that a person is the person that she/he claims to be is important, but only half the story. Appropriateness and authority are the other half. In addition to checking for authenticity, it’s essential to determine whether the person is the “correct” or appropriate individual with which to interact in the organization. By “correct,” we mean that the person is authorized to receive, accept, or act on the information.

Examples: During an incident a call is placed to the security manager of organization XYZ. Because the manager is not available, the secretary takes the call. The secretary’s identity might be authenticated; however, it still might not be appropriate to discuss with or disclose to the person details that are intended for the manager. It might be more appropriate to leave a message for the manager to return the call as soon as possible.

Alternatively, a senior manager of XYZ might telephone the CSIRT and demand all kinds of action to be taken with regard to the same incident. If this person was not the team’s registered point of contact for such issues, the CSIRT would need to refer him or her back to the registered point of contact of that organization to make the demands (if appropriate) through the appropriate “chain of command.”

Without such procedures for authentication in place, teams and their constituencies are potentially susceptible to social engineering attacks (discussed below).

3.7.2.1 Social Engineering

Social engineering is a phrase commonly used in the CSIRT community to refer to a situation where someone presents a fake identity to trick another person into doing something that they would not normally do if the real identity were known [Gordon 1995, Greening 1996]. A classic example of social engineering (like the example above) is of someone pretending to be a high-ranking official and telephoning the guard, telling him to open the gates or else. Amazingly enough there is evidence that brute-force psychological attacks similar to this are still successful today. Two examples of this type of attack are relatively well known:

- **unsolicited media calls:**

When a media representative thinks that an incident is going on, (s)he may try to get insider information. By not revealing her/his identity or explicitly pretending to be “just another victim,” a team member might reveal information in an effort to help a victim to recover.

- **intruder calls:**

Social engineering is a well-known technique for intruders. If the intruder thinks the CSIRT may be monitoring their activity (such as an intrusion), they might call the team in an attempt to find out if their activity has already been detected. They might pretend to be a contact from the site in order to elicit information about the activity, much like the media example given above.

Other social engineering approaches can be seen in the well-known ploys to entice people into extracting an attachment or visiting a URL. These may contain nasty side-effects.

Example: With regard to email, unsolicited items are sent to unsuspecting recipients. These may contain a return address, a provocative envelope, or something else that encourages its receiver to open it. Malicious email often contains a return address of someone we know and often has a provocative Subject line. The unwary recipient may not exercise due care in handling that message.³⁹

3.7.2.2 Technical Possibilities and Limitations

Modern telecommunication facilities like ISDN provide the “caller ID” feature. The telephone number of the calling site is signaled to the called site, and if the telephone has a display, this number can be shown to the person receiving the call (however, there are also features to block caller ID).

Depending on the technical communication facilities, support can be available to prove or verify authenticity. One of the most well-known authentication methods in relation to today’s

³⁹ An easy-to-read article on “Home Computer Security” provides additional information about social engineering attacks (as well as other computer security issues) that might be helpful to your constituency. It is available from the CERT web site at <http://www.cert.org/homeusers/HomeComputerSecurity/#3>.

networks are digital signatures, such as those used in the secure mail systems PGP and S/MIME.

Example: To authenticate the origin of all outgoing email, a digital signature produced with PGP authenticates each email message issued from DFN-CERT.⁴⁰ Every recipient can check (verify and authenticate) this signature with PGP. This check relies on the authenticity of the public key the DFN-CERT member used for the digital signature. Therefore, it is the responsibility of the recipient to check the signatures using the published PGP public fingerprint of all the DFN-CERT team members.

Other tools like S/MIME depend on a hierarchical key certification process, where certification authorities (CAs) or trusted third parties (TTPs) check the authenticity of a user and the relationship between the user and their public key. If they are able to verify this information, they will certify the key's authenticity.

It is important to note that digital signatures can also provide a high level of authenticity and protection against disclosure or other attacks by using associated encryption capabilities (e.g., both PGP and S/MIME are capable of this). It is important to understand the limitations of the mechanisms used and to use each mechanism within these limits. When there are inherent problems or tradeoffs, organizational approaches can help provide the necessary security.

Example: CERT-NL uses a new team-key each year. As the team-key is used for day-to-day operation, it is stored on systems that, more or less, have a direct connection to the Internet. There is, however, a CERT-NL master certification key that is kept off-line (is never used on an Internet-connected host) and whose use is controlled by a strict procedure. Every time a new CERT-NL team-key is generated, it is signed using the master certification key. All the keys of the CERT-NL staff members are also signed using the master key. This overall system neatly blends practical demands and security. Constituents must verify that the staff keys are properly signed by the master certification key and can then safely use the staff keys without checking the fingerprint of each staff member individually to verify the key. To bootstrap the process, all constituents must obtain and verify the public master certification key.

3.7.2.3 Databases

Another area where tools are involved is the use of information databases, particularly those containing contact information. As internal databases form an integral and important part of the interaction process and CSIRT communications infrastructure, they should be very carefully protected. If an attacker could manipulate the database, they could compromise the data, and seemingly authenticated data could be entered and the team members would trust it.

⁴⁰ Many other CSIRTs also PGP-sign messages, as well as other documents they produce or disseminate.

The same problem exists when using public information sources. Here the possibilities for manipulation are greater, and hence the invested trust in such information by the CSIRTs is limited.

Example: The DNS system and Whois databases (two widely used directory services on the Internet) are often used for contacting victim sites when no better point-of-contact information is available. As it is possible to masquerade as a DNS server for another system, every public information server must be considered as “not trusted.” Besides questioning the authenticity of the information available, one may well also question the integrity of the data: for example, Whois information is often outdated or contains errors. In the worst case such flaws may lead to passing information on to the wrong person.

Example: In Europe a directory of European CSIRTs is maintained by the Trusted Introducer service.⁴¹ This service maintains related records of European CSIRTs in the Whois database (IRT objects). This allows the lookup of the responsible CSIRT based on IP addresses.

3.7.2.4 Anonymous Information

The final area is how a team should deal with anonymous calls or calls that cannot be authenticated at all. No sensitive or substantial information should be passed to any anonymous callers. But when they provide new information, a team must decide if such information is useable and if (and how) such information should be handled. It may not be possible to verify the information provided, so it should be tagged as such and its anonymous origin must be tagged too. One of the best reasons for considering the use of anonymous information is that it makes no difference, for example, whether a warning comes from an anonymous caller or not—either way, to be safe, you will check whether the warning is valid.

3.7.3 Secure Communication

Authenticating the origin of important data is only part of its safe handling. It is also important to adopt security mechanisms suitable to protect the information during its transmission across networks. This does not only apply to computer and telephone and other telecommunication networks (including remote systems and access), but also information transmitted via more traditional means, such as post or couriers, which are also vulnerable to attack (or loss).

In the same way that cryptographic mechanisms can help to ensure authenticity, they can ensure confidentiality. Efficient encryption mechanisms are available, although for various reasons, specific mechanisms are not universally allowed or are not exportable to other countries (government regulations).

⁴¹ For more information on the Trusted Introducer service, see <http://www.ti.terena.nl/teams/index.html#DIR>.

Wherever cryptographic mechanisms are used, key management is a major issue to address, by means of a policy and operational procedures.

Example: FIRST uses PGP to protect email communications. As it is very difficult to use public key encryption with a large number of parties (in January 2003 FIRST had almost 130 member teams) conventional (symmetric) encryption is used. All FIRST members share the knowledge of the same pass phrase, which is changed regularly. In addition, digital signatures can also be used to provide authenticity, enabling other teams to check the origin of a message. The procedures for use and maintenance of the keys are distributed among FIRST members [FIRST 1998].

In the case of telecommunication networks, additional black boxes can be applied, since confidentiality is not usually a default feature of telecommunication services. Such encrypting devices are available in the open marketplace, although the protection provided might depend on the implementation and other factors like export restrictions, which limit the availability of products all over the globe.

Example: Some teams use Secure Telecommunication Unit (STU III) or SECURE TERMINAL EQUIPMENT (STE) devices, which can protect telecommunications. Such devices are controlled equipment that have special handling/reporting procedures and requirements for their usage and are limited to certain usage communities (the U.S. and Canada, for example).

Example: Other countries have developed encryption technologies for GSM cellular phones or ISDN connections. Usually the application of such technologies will only work with communication partners that share the same technologies as well as having exchanged the cryptographic keys needed.

3.7.4 Special Considerations

The following text will present interaction considerations specific to given environments. The objective is to explain the practical considerations for interactions that have already been introduced. The parties involved in interactions are not described in detail, but the important issues are explained through examples.

When conducting interactions, one of the first issues a team should address in its policies and procedures is the level of service it is willing or able to provide to different parties. This statement might include details like response times or might describe specific forms for exchanging information. By doing this, available resources are considered and devoted to particular tasks and priorities.

As each teams' situation will differ, the examples below, where possible, indicate both beneficial approaches to take and pitfalls to avoid. Although the examples provided include a

wide range of possible partners, others might certainly exist. But we believe we've covered the most important interactions to consider. Any others that you may identify can likely be treated similarly to one of the categories below or will be similar to that for the media (i.e., open, public, and unknown).

3.7.4.1 Constituency Sites

The CSIRT's primary objective is to help its constituency. Most of the issues to consider have already been covered in earlier sections of this handbook. For interactions one additional consideration to be discussed is the need for different kinds of contacts even within a single site. Of course, if the same person at the site fulfills multiple roles, a site may still only have a single contact.

Since the escalation process in dealing with incidents will need decisions (like the decision to report to law enforcement), contacts for each phase of escalation are necessary.

Example: While the technical details of an incident are passed on to an administrator responsible for the daily operation of the network connection, some information must also be directed to management. If, for example, a site reporting a new incident already informed law enforcement, other sites need to know this information to consider their own decision in the light of this fact. On the other hand, this information might not become available to the CSIRT, and can therefore not be passed on to others.

When defining policies and procedures, the CSIRT must prevent a single site or constituent from consuming all of the team's available resources unless the team considers the activity to be of such importance that it should take precedence over all other activities. During periods of limited staff resources (e.g., vacations or conferences), prioritization will become even more important to distribute the activities among the available staff.

Documented and publicly available policies will allow the sites and constituents to understand limitations and restrictions, but even so steps should be taken to alert the constituency to these times. For instance, a holiday message might be distributed that provides information for high priority reporting procedures. This appropriately sets the expectations of the constituency, who will be more patient with the CSIRT than they may otherwise be without such measures.

Depending on the size of the constituency and the service provided, pre-registration may be a possible option. Clearly pre-registration of a constituency is only a possibility if the constituency size is relatively small (in the hundreds) and is fairly stable. It might also be possible if the relationship between the constituency and the CSIRT is on a contractual basis, such as with a commercial fee-for-service team or network service provider, where it is easy to add the pre-registration criteria as a supplement to an existing contract. During the pre-

registration, issues such as information disclosure restrictions, trusted points of contact, and preferred method of secure communications must be addressed.

3.7.4.2 CSIRTs

Incidents that require no external interactions with other parties are rare in today's "unbounded" networked environment; they arise only if an incident is local without any external relations or side effects. Even then, external interactions may become necessary, such as when law enforcement is involved.

Besides direct contacts at constituency sites, the most important cooperation partners for CSIRTs are fellow teams. While handling an incident, direct help and information exchange are most important, and there is potential for teams to provide mutual support. This is particularly true if the teams have been in the CSIRT business for a long time or have particular technical expertise. Other examples of support might be provided in one of the forms described in Table 17.

By exchanging information, cooperating teams usually benefit, making it easier for them to either fulfill their duties or provide a better service. But sharing information in the first place is not as easy as one might think. When considering the issues outlined in Table 18 it becomes clear that the extent to which teams are able or willing to exchange information and to cooperate on confidential issues depends on any existing trusted relationship they may have with each other. The existence of a formal (written) agreement between two teams might make it even easier for the teams to exchange information, assuming a clear understanding of all the issues described above already exists.

If two teams want to initiate a cooperative relationship, they will first need to establish the necessary foundation of trust. Building such trust is not an easy task and will require time. One of the most important steps towards such a relationship is getting to know each other. The teams should exchange visits and try to understand each other's goals, objectives, procedures, and policies as much as possible. This will help the teams to make a realistic assessment of whether a deeper relationship is achievable and beneficial. The teams might want to start by collaborating on a small project with minimal risk rather than starting on a larger, more complex and risky task.

Table 17: Possible Inter-Team Support Types

Support Type	Description
Education/Training	This might range from issues like “Forming a new CSIRT” to technical tutorials to understand the nature of incidents.
Out-of-Hours Coverage	While one CSIRT may only provide service during business hours, another fellow CSIRT may take calls during other hours as part of a collaboration agreement. This is particularly relevant if a team operates under the indirect control of a coordination center.
Technical Expertise	To address technical questions and share this knowledge with other teams.
Cooperative Work	To address problems that are too difficult to solve with the resources of a single team, two or more teams might come together and collaborate to seek the solution to such a problem. This handbook is a good example of this kind of cooperation.
Other Opinions	While working on the solution to a particular problem, the members of the team involved may be too close to the problem to view it objectively. To avoid the negative impact that might arise in these instances, another team might be asked to review and provide an opinion on the proposed solution before it is publicly distributed. Existing CSIRTs have a long history of exchanging draft advisories and often incorporate many suggestions before the final advisory is released.
Point of Contact to Other Teams or Experts	Since a team might need a trusted contact for a specific site or network, they can ask other teams whether they have an established contact or if they know somebody else to ask. This also holds true for contacts to technical experts and vendors.

Table 18: Considerations for Information Sharing

Issue	Description
Confidentiality/Secrecy	Since the information might also be valuable to other parties, its confidentiality must be maintained. This is true for transfer, storage, and actual use. The mere reaction of a team member might be enough to reveal at least some part of the information, for example the existence of a new bug or security hole or the existence of some other incident.
Appropriate Use	While the information belongs to one team, it must be clear to other teams that to obtain access to the information they must adhere to any restrictions that the original team places on the information and conform to what the original team considers as “appropriate use.” Most of the time the official signing of a non-disclosure agreement is necessary to obtain such access. Part of a non-disclosure agreement will list the rights and duties by which appropriate use is established.
Disclosure	Since the information may be distributed to the public at some future point in time, disclosure restrictions should be stated up front. Some teams put time constraints on information. While it is forbidden to disclose the information by any means before the time limit, it is perfectly acceptable to incorporate this information in an advisory to be disclosed after the time limit. Setting a timeline is not easy in an international environment. Differences in time zones mean system administrators in one area of the world can be finishing work, or already at home, while others are just starting their working day.
Proper Acknowledgments	Since the information was collected, analyzed, and made available by other teams, the team using it should consider a fair and proper acknowledgment of the original source.

Just as there are teams that you know from previous interactions, there are also teams that you have heard about but are less familiar with. Since you have no knowledge whether the team is suitably qualified or even genuine, it can be a difficult decision to pass information on to them. If you have some initial knowledge about the team, the decision may be easier to make. One way to obtain such information is to ask other teams that you have a good

relationship with what their experiences may have been with those team(s) that you are less familiar with. It would be so much easier to rely on a mechanism to identify trusted teams, but as yet no such mechanism exists.

We'll continue by discussing other issues involved in inter-team cooperation. These issues are more closely related to the CSIRT's operational procedures.

Mandatory Information

The issue of dealing with incoming information was described previously. There is critical information that a team must have before it can process a report. If this information is not supplied in the initial report, a delay will be incurred until the team obtains the information. The delay can be significant in some cases, such as if the report was sent just before a weekend or if extreme time zone differences are involved. A team can attempt to ensure that another team reports the mandatory information through the use of an inter-team reporting form, but such a process will need to be developed before that type of situation arises.

For example, a team may choose to share or report information at a peer-to-peer level (CSIRT to CSIRT); they may choose to report up to a high-level coordinating team (e.g., a Government or country-level CSIRT), or the team may be mandated (through a regulation) to report **because** of the organizational structure (e.g., some teams in the U.S. are required to report "up the command chain" within a specified time). Not all inter-team relationships are at the peer-to-peer level. Some teams elect to participate in a voluntary hierarchy as the situation may dictate.

Example: Two or three CSIRTs may share the roles in analyzing a piece of malicious code or new vulnerability, with one team collecting the analysis from each and coordinating the overall analysis. This would be a peer-to-peer relationship with a voluntary hierarchy.

Even if teams interact as peers for one activity, it does not preclude them interacting in different ways on other occasions. Less frequently, teams exist within a mandatory hierarchy (e.g., some of the military service branches report "up the command chain" to a coordinating CSIRT).

Example: Take the above example. Consider that three CSIRTs agree to share responsibilities for the analysis of some activity that has been reported. Each takes a piece of the work: one examines the vulnerabilities, another the logs, and a third works on preparing the information that will be disseminated to coordinate all the pieces involved in the response and follow-up with affected sites. This latter team is the team that has the external-facing role, and all information flows from them to other parties. At some later time, other teams may participate in similar activities, but taking on different

roles. Even within such roles, each of the CSIRTs may have other roles they play within their own hierarchy (e.g., reporting to sponsors, other coordinating teams, etc.).

The point to make here is that teams can have different roles in different situations, handling functions as a peer, coordinator, leader, liaison with law enforcement or media, or some other role.

Who Has the Lead?

Even if teams normally interact at the peer-to-peer level, transient, voluntary hierarchies often evolve for the duration of a single incident. When multiple teams are involved within one incident there is a need for coordination; otherwise, duplication of effort will take place, such as multiple teams contacting the same sites with the same information. Rather than waste the time of the teams and sites in this way, one team will usually take the lead for a given incident. Deciding who takes the lead in coordinating response to an incident is usually done on a case-by-case basis. Generally, the coordination is undertaken by the CSIRT receiving the first report or handling the largest part of the incident. Coordination can also be agreed upon in advance through a predefined arrangement (e.g., subscribing to a coordination service with mandatory subordination).

3.7.4.3 Sites Outside the Constituency

As a team becomes well known, it will receive requests and information from almost everywhere, especially if it is dealing not only with the local aspects of a single corporation or organization.

Example: CERT-NL might be incorrectly assumed to be the Dutch CSIRT (judging from the name alone). If people do not know anything else other than there is a CSIRT in the Netherlands, and if they have an incident involving a Dutch host/site, they may report it to CERT-NL. This is true even if the site involved is not within the formal constituency of CERT-NL, the customers of the Internet service provider SURFnet.

If such information should be received, CERT-NL would pass it on to one of the other CSIRTs that exist within the government, universities, and commercial organizations. When a team receives an incident report, they will have to deal with it at some level, whether they were the right team to report to or not. Only teams with a very specific constituency or service will perhaps opt for not dealing with this kind of report at all. The least the reporter can expect is a short message indicating that he should resend the report to another team.

Example: Consider a medical analogy. If you experience a health problem and ask for help, there is no way that a doctor or nurse can ignore you (at least in many parts of the world).

Note, however, that the nature of the help that you provide in such situations may be different from what you offer to your own formal constituency. Another factor that might affect the response that you offer to a site is the trust level. If you don't know the source of a report, it is difficult to assess the quality and relevance of the report (except that the data provided may verify authenticity, correctness, and relevance).

Example: The CSIRT receives a wide variety of calls; some of these could be related to vulnerabilities reported by anonymous callers. The level of credence provided to one call may be limited; however, if the team received the same call about vulnerabilities from trusted members of the community (say another CSIRT or trusted expert), the latter call will be given a higher level of confidence that the report has some validity (or possibility of quality and relevance).

When setting up a team and allocating resources and responsibilities, it is important to understand that requests originating from outside of the declared constituency will occur and must be handled. In most of the cases, a simple reply containing more appropriate addresses to report to will help the reporting site to contact the right parties (and limit the occurrence of future similar requests). To be able to give such answers, the team must prepare the necessary information in advance and establish policies as to what reply is adequate for what questions or reports.

In the past, some teams, especially if they were responsible for a large constituency, provided reporting sites with more adequate addresses; and in addition to relaying this information, they also provided the reporting site with some kind of "first aid." This often resulted in the reporting site receiving the same service as a constituency member. This approach gives a team a good reputation, but requires additional resources and might lead to the following problems:

- Other CSIRTs do not like their constituents to receive help from another team. (There may be information that the CSIRT needs to obtain or provide to the site, but if the site receives preliminary information from another CSIRT and assumes this is all that is needed, they may never contact their own CSIRT.)
- Upper management (or sponsors) generally do not like their team's resources spent on "outside" (non-constituent) parties.
- The service to the declared constituency might be adversely affected due to resource limitations.

One special case might arise when the reporting site does not fall into the declared constituency of any CSIRT.

Example: Approximately 90% of all European nations now have a funded (not volunteer-based) CSIRT. In many nations such CSIRTs were established for research networks, so depending on their policies, they may or may not handle incidents involving commercial

sites in their countries. In other nations up to 15 teams exist, still not providing complete coverage. Most likely the private users will not be served by a CSIRT.

Therefore each team should set clear expectations and establish understandable and enforceable policies to deal with external requests. Whenever there is another CSIRT responsible for the reporting site, the reporter should be directed to report to that CSIRT or that team should be notified about the report. If the reporter requests complete confidence they should be encouraged to contact the responsible CSIRT directly and should explain the benefits for doing this. For example, telling them it is more appropriate to report to their own team as they'll receive more appropriate response or assistance, indicating that there may be mandatory regulations requiring they report to their own team, and/or indicating additional benefits their own team will attain in seeing this report potentially as a piece of a larger activity (the "bigger picture" we discussed previously).

Additionally, since the existence of the report from the non-constituent member together with the request for confidence, in and of itself is valuable information to the responsible CSIRT, the team who originally received the report might choose to provide sanitized information to the responsible CSIRT about the report; without revealing any details on the origin of the request. (Armed with this knowledge, the responsible CSIRT might be able to identify why the original reporter opted for confidence and may then be able to improve their service or change some of their procedures to eliminate similar situations from occurring in the future.)

3.7.4.4 Parent Organizations

A team's parent organization might be upper management, a funding body, or shareholders. Like any other member of the team's constituency, the parent organization may request its services, from incident response to consultancy, training, or presentation delivery.

This is an important and political topic. In most cases the parent organization will receive a higher priority than is normally assigned to identical service requests from other constituents. In the case of incidents, if the parent organization consists of operational units that are also possible targets of attack, a team might handle the report and immediately escalate the incident involving those units to the CSIRT management's attention as well.

3.7.4.5 Law Enforcement

Whenever an incident is related to a crime, law enforcement will become a major issue. Law enforcement agencies will try to

- learn more about the incident itself
- learn more about the technical issues involved
- identify/contact sites involved

- obtain information on recent activities and/or damages related to the incident

A team is in a delicate position between confidentiality provided to its constituency and the need to cooperate with law enforcement. A team's policies will determine the amount and type of information a team will voluntarily supply to law enforcement. If required to with a legal order (via a subpoena or other court order), a CSIRT must provide specific information as requested by law enforcement. Policies and procedures should define the services provided to law enforcement and should clearly state the circumstances under which information is revealed.

To ensure good cooperation between a team and law enforcement, mutual understanding leading to mutual respect is necessary. Teams should be encouraged to develop a relationship with appropriate law enforcement contacts as early as possible to initiate these interactions.

The policies of a team should define who is responsible for talking to law enforcement agencies. This includes requests from other non-local or international law enforcement agencies. Such requests are difficult to address and should be redirected to local law enforcement. Therefore it is in the interest of each team to know their legal and law enforcement points of contact and prepare in advance for such requests.

Another benefit in cooperating with law enforcement agencies is the exchange of statistics and helping to raise awareness within the law enforcement community regarding the types of activity being seen by (or reported to) the CSIRT. Since the CSIRT will have first-hand knowledge not only about computer crimes but incidents not considered as crimes, they can substantially enhance the ability of law enforcement agencies to build the bigger picture. At the same time, law enforcement agencies may be able to share sanitized feedback with the CSIRT, as appropriate, regarding activity that could be of interest to the CSIRT in correlating other incident activity the team is seeing reported.

3.7.4.6 Media

Since the media has the power to influence public opinion, each team should have a media policy and establish associated procedures. The objectives should be

- provide reasonable feedback and information
- maintain the interests of sites
- speak for yourself and let the sites speak for themselves

The media has its own goals and reasons for obtaining information regarding an incident. These goals often conflict with those of a CSIRT.⁴² Consequently, the media are often seeking

⁴² McGillen, Terry. "CERT Incident Communications." 5th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, St. Louis, MO. August 1993 and

to obtain more information than a team is willing to provide. A team should make known to the media the team's point of contact for media requests. Prior to their first contact with the media, the point of contact should be suitably trained in media interactions, including what to expect and how to appropriately handle situations involving the media.

This topic will be discussed in greater detail in Section 3.8.8, "Information Disclosure."

3.8 Information Handling

Handling incidents is always related to handling information. Information is always the key, regardless of whether specific information relates to a site, a product, a new vulnerability, an ongoing attack, or a password.

Information must first be collected and entered into whatever type of system the CSIRT is using to record and track incident reports and other relevant data. Every piece of information must be stored and protected for the entire time it is held by the CSIRT. Tagging the information according to its type and sensitivity will facilitate its further handling. Before the information continues to be processed it must be prioritized to ensure that the most important information is worked on first. After the information is reviewed and analyzed, the information itself or an aggregation of multiple information pieces may be disclosed to provide guidance and support for the parties involved, usually the team's constituency.

3.8.1 Information Collection

While much of the information that a CSIRT handles will be received directly, there is also a need to collect information, such as proactively searching for information on the Web or retrieving information from other sources (technical reports, analysis, news, trusted experts, etc.).

Before collecting information, it is advisable to establish a dedicated policy and suitable procedures to determine

- what kind of information sources are acceptable
- what kind of quality controls to conduct
- how to recognize errors, omissions, or imprecise data

If information is actively collected, it may come from one of the following two sources:

McGillen, Terry & Fithen, Katherine T. "Public Communications in the World of Incident Response." 9th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, Bristol, U.K., June 1997.

1. **Open source information:** This includes any kind of publicly available information. The options range from more traditional services such as news or mailing lists to search engines or the Web.
2. **Exchanges with other parties:** Since other people may already possess the information that a team needs, exchanging information with others can directly benefit the team. The main problem here is knowing who has the information and establishing trusted relationships, so that the person/team is willing to share the information. (This highlights the importance of good partnership with others; see Section 3.7, “Interactions.”)

Since the information available is continually changing, information collection and other related policies and procedures must be frequently evaluated and verified to ensure the most current information sources are being reviewed.

Incoming information from other parties will have to pass through the team’s triage function, as described in Section 3.3, “Triage Function.” To stimulate the reporting of information related to events, vulnerabilities, and potentially interesting discussion threads, the reporting users must be provided with appropriate support, such as reporting forms. This support could also include contact information, such as team phone numbers or email addresses, that can be used to report other types of information or submit requests to the CSIRT.

Standardizing across policies and procedures will help the team collect information in a more consistent format. Having standardized the format used, further actions on the information such as storage, verification, categorization, and prioritization of the information will be much easier to carry out.

3.8.2 Information Verification

Before any information can be used, some kind of verification has to take place. Usually the process involved will at least consider the following three issues:

1. **Origin:** The source of the information and related factors like the knowledge, experience, role, and function of the reporter. As with all communication, the origin may substantially affect the further processing of the information provided.

Example: If DFN-CERT reports to CERT-NL an IP port scan over large IP address ranges originating from a Dutch university network, CERT-NL will assign a higher priority to this report than a report that originates from some unknown individual, although the report is still double-checked.

If a report comes in from a trusted source, it might make the follow-up a bit easier, but there are times when the identity of the caller may make the situation more difficult or complex.

Example: If the funding/sponsoring body calls, in most cases more time will be spent on the follow-up in comparison to other callers, regardless of the real priority.

2. **Content:** Is the information likely to be true, or is it obviously wrong or misleading? The presence or absence of technical correctness of the content may affect the subsequent processing of the information.

Example: Concerned constituents who have received hoax virus reports from other parties often send the reports to CSIRTs for verification. Hoax reports commonly contain information that is technically incorrect or even impossible. Although CSIRTs may need to alert their constituency to the fact that a hoax report is circulating, this may receive a lower priority than a virus report that does appear to be technically correct and warrants further analysis and investigation.

3. **Distribution:** This relates to the channel used for the report and possible impact on the authenticity of the incoming report. The possibilities range from digitally signed and verifiable reports to those that may have been received via an anonymous telephone call or even a letter via the postal service.

3.8.3 Information Categorization

Information entering organizations must be categorized in some way. All information enters a CSIRT through the triage function; this facilitates initial categorization. Examples of well known categories are private vs. business, and urgent vs. non-urgent vs. garbage; usually such simple categories are not even formally described.⁴³ Although categorization is implied by prioritization (handled below in Section 3.8.6, “Prioritization Criteria”), it is more appropriately considered as a separate and independent activity.

The category in which information is placed affects how the information is further handled (e.g., storage, dissemination, distribution, and disposal). Without differentiation, all information must be protected to the highest level and similarly disposed of.

Even if no explicit categorization is used, the person reviewing the incoming information will apply their own perceptions on the type and importance of each piece of information. Since these perceptions may differ among individuals, clear and concise procedures (as explained in Section 4.2.2, “Information Categorization Policy”) must be available to standardize and guide this process.

Many CSIRTs have a separate process or procedure for handling contact information that is different than the process used for incident reports and information requests. Contacts (people, organizations) are usually sheltered from exposure, even to trusted fellow teams. Therefore specific statements for how to sanitize this type of information may be included in the procedures. Some teams may even have contact information as a category of its own.

⁴³ Some guidelines refer to an information classification policy. We decided to use categorization instead. The word “classified” is used throughout this section and in Section 4.2.2, “Information Categorization Policy,” in its general context, not in its more restricted military and/or governmental context.

Categorization is often based on the information itself. Sometimes the categorization is obtained following an interactive dialog with the information provider. At other times the information provider also specifies a category for the information (e.g., a constituent reporting a vulnerability or an incident).

Information may also have to be (logically) cut into pieces—e.g., incident logs, where only specific names or IP addresses might need to be protected from disclosure, whereas the rest can be freely transferred to other parties.

Example: The CERT/CC handles contact information categorization by requesting that the reporter of an incident state the information disclosure restrictions on the data that they provide in three categories:

- other sites involved in the incident
- other response teams
- law enforcement

If the reporting site does not provide this information (requested in the CERT/CC's incident reporting form [CERT/CC 1997a]), then the CERT/CC uses a default of “no-disclosure,” which requires the CERT/CC to contact other sites or response teams without attribution to the reporting site.

3.8.4 Information Storage

Whenever information is stored (whether it is hand-written or stored in a computer system), security is of major importance. Without security, a team cannot realistically expect to protect the interests of its constituency and the confidentiality of the sites involved.

This is particularly true if information is stored collectively, such as in large databases. In such cases the value of the collected information is greater than the sum of its parts. For the same reason that collected information is a great benefit to the CSIRT (to help the team see the bigger picture), it is also a weakness. A CSIRT might survive the consequences if a small

activities have been identified and reported to a CSIRT, identify information about vulnerable sites, gain information on new vulnerabilities, etc.

Use of multiple logical databases is one useful approach to information storage. It allows information to be accessible, easy to use, easy to change, and flexible enough to support various services.

However the data is stored, access to the following must be possible:

- contacts
- actions taken (or to be taken)
- incidents (current information about the activity, status and ongoing summaries as changes occur)
- vulnerabilities and patches
- artifacts (scripts, tools, file remnants, etc.)
- logs or other data associated with the stored information

3.8.5 Information Sanitizing and Disposal

Information sanitizing and disposal is an essential component of information handling. This is particularly true for a CSIRT that often has sensitive information referencing a (possibly very large) group of people and organizations. As discussed previously in Section 3.8.3, “Information Categorization,” information in a given category should be sanitized and disposed of in a consistent manner appropriate for the level of sensitivity attached to the information.

Information can often be sanitized to prevent inappropriate disclosure of sensitive information without any adverse effect on the usefulness of the information provided to a recipient.

Example: Site A finds a copy of a password file in the various intruder artifacts found on one of its systems that has been compromised. The CSIRT at Site A is not sure of the origin of the password file. Even if incomplete information exists about the origin of a password file, there may be enough information included in the file to indicate a likely source. If so, then Site A’s CSIRT may be able to contact the suspected source site (Site B) so that they can determine if their site has also been compromised. To do this the CSIRT sends the Site B a copy of the file with all encrypted passwords removed. This avoids creating further potential risks through clear text transmittal of the passwords. Specific information like user names and home directories remain intact, allowing for a high degree of assurance, without further distributing information likely to be misused if captured by other parties. It also protects the passwords, in case Site B is not the source

of the password file; the CSIRT has not given confidential information to a non-related third party.

The storage of user and site-related information and the relationship between incidents and specific organizations have associated privacy concerns. It may be in the best interests of a CSIRT to keep a complete log of information, but this also could potentially affect every party for which information is stored.

Example: If there is a legal requirement to provide specific information about an intruder, law enforcement might request all the media on which data about the intruder is stored. As a consequence, the team can no longer assure confidentiality of other information that is not related to the intruder attacks that are also stored on the same media. Knowledge of such instances might result in reluctance of constituents to report future problems to the CSIRT. To avoid such perceptions—as well as the real consequences—the team needs to apply technical means to ensure that all data about the various incidents are independently stored and accessible. If law enforcement requests information be provided on one incident, only data related to that particular incident will be handed over.

To limit possible exposure, a team might choose to store only sanitized information after a specified period of time or to rely on a summary containing only statistical and technical points. By deciding to do this, the team must expend a considerable amount of effort to dispose of all information that is no longer needed (notwithstanding the actual effort that will be needed to prepare the summaries of those statistical and technical points in the first place). This is particularly difficult in the case of backups, because the whole purpose of a backup scheme is to ensure that information is available in the long term. It is unlikely that older backup tapes can be easily rewritten to dispose of information that is no longer needed.

Example: One way to handle this issue would be to use two different backup schemes: one for operating system and user data, another for incident-related information. This implies that no incident-related data is stored in the users' data area. While normal backup tapes are reused when needed, the incident related tapes are overwritten several times before reuse to avoid later recovery of previously stored information. If tapes are no longer used, they should be physically destroyed, not just thrown away.

3.8.6 Prioritization Criteria

Although many types of incidents are “critical” or “serious,” even within these individual categories, CSIRTs will need to assign a priority to determine which to handle first. The importance of an incident might depend on many factors, and the priority can also change if new information is discovered or reported. So trying to establish and maintain a priority list of ongoing incidents is not easy and can in fact be a dynamic activity.

Different schemes exist for selecting the most important incident or for ranking several incidents:

- resources needed to deal with it
- impact on constituency
- type of incident
- type or extent of damage
- target or source of an attack

As always, exceptions will arise that are not directly accommodated within the scheme selected. The scheme will need to provide some flexibility to allow for such exceptions. This might include initially giving an incident a default priority at the middle or top of the priority scale until sufficient information is available to prioritize it more appropriately. Any policies that affect the prioritization process must be regularly reviewed and refined over time to accommodate items that were once considered exceptions but are now common and reflect other changes in trends and needs.

Example: Several years ago, the CERT/CC listed root compromises of systems as a high-priority incident. However, over time (as the Internet has grown and the mission of the team has changed) other types of events have been given higher priority. The CERT/CC Tech Tip on Incident Reporting Guidelines⁴⁵ lists the following priorities:

- possible life-threatening activity
- attacks on the Internet infrastructure, such as:
 - root name servers
 - domain name servers
 - major archive sites
 - network access points (NAPs)
- widespread automated attacks against Internet sites
- new types of attacks or new vulnerabilities

Example: New incoming incident reports are reviewed by an experienced staff member. His insights provide him with the ability to assign any incident to the appropriate staff member(s). For example, he can assign more well-known incident activity reports to novice CSIRT staff (such as port scans, UBC/UCE reports, or attacks using well-known security vulnerabilities); while more complex incidents that require more in-depth expertise and knowledge will be assigned to the more senior members of the CSIRT staff for handling.

Continuous reprioritization of incidents must occur. Whenever new information on a given incident comes in, its overall priority might change. Since a change of priority also affects the

⁴⁵ http://www.cert.org/tech_tips/incident_reporting.html

reporter and affected sites, these parties will also need to be informed accordingly. This is most important whenever incidents are downgraded to a lower priority. On the other hand, a seemingly lower priority incident that is suddenly upgraded to a higher priority—due to the discovery of new information—could be equally as important a reason to notify the reporter and affected sites.

Since almost all teams operate with limited resources, there will be times when a team cannot handle all incidents reported to it. In rare cases it might be possible to hand off such incidents to other teams. When incidents cannot be handled, the reporter must be informed. Without such communications, the users are left in the dark, and rumors will arise about the team's apparent lack of response. This might damage the reputation of the team and negatively affect the overall operation.

Most teams select some combination of prioritization schemes to generate their overall prioritization criteria. Commonly, teams prioritize on one scheme and then refine the priority by application of one or more other schemes. Depending on the scheme chosen for use, there are tradeoffs to be considered. The tradeoffs must be defensible and communicated, as there will always be individuals who claim that their incident should deserve the highest possible priority. We will identify some of these tradeoffs as we discuss some of the possible prioritization schemes in the remainder of this section.

3.8.6.1 By Target or Source of Attack

When prioritizing based on the target of the attack, a value is assigned based on the role, mission, or importance of a target site or system. Targets within the constituency boundary can be viewed as more important than targets outside the constituency, since the CSIRT's function is to serve that internal constituency. Given multiple targets within the constituency, the team needs to be able to discriminate between different possible targets and associate corresponding priority values. A target's value might be determined by the type of data held on it, the role it plays within a network's infrastructure, or some other factor.

Example: Consider a CSIRT whose constituency is a manufacturing company. Using a "by target of attack" priority scheme, higher priorities would be assigned to incidents targeting systems that hold proprietary information (e.g., research or production systems) or personnel data than to those holding less-sensitive data.

It is not always possible to determine the real source of an attack, because intruders can hide the source of their activity. Often, intruders weave a path through many systems (often crossing international boundaries) before launching an attack. As a result, the only information about the apparent source of an attack in an initial incident report is the site being used to launch the attack. This attacking site is not necessarily the real source of the activity.

If used, this approach is similar to that approach taken for attack targets. Values are assigned to possible classes of attack sources based on the perceived threat.

Example: Consider a CSIRT whose constituency is military. Using a “by source of attack” priority scheme, higher priorities would be assigned to incidents involving attacks from “unfriendly sources,” militant organizations, or overseas sites, particularly those considered as hostile nations.

3.8.6.2 By Type or Extent of Damage

The extent of actual loss or damage resulting from an incident is sometimes difficult to assess. Not only can it be hard to collect this data after the fact but it is even more difficult to predict such data with any accuracy. The assessment will be influenced by the personal experience of those undertaking it, the correctness of incoming reports, and the type of information available to the team. A CSIRT with direct constituency authority is likely to have access to detailed information about an incident involving its constituency. A team with less authority is unlikely to have access to information at the necessary level of detail to make a reasonable assessment. As a result this type of scheme is more commonly seen in teams that have some constituency authority.

Even if the damage is known and can be described, the same metric must be used across different incidents to enable their comparison.

Example: Hospitals and emergency teams have similar prioritization schemes:

1. loss of life
2. injuries of humans
3. loss of money / violation of rights

If “loss of money” is used to determine priority, some model to determine the financial damage will be required. However, it is very difficult to calculate financial damage for some incidents. For example, it is difficult to estimate the amount of money that an organization might lose due to public knowledge or information disclosure relating to an intrusion. As a result, this criteria may be of limited use in prioritizing incidents.

3.8.6.3 By Incident Type

Using this criterion, known incident types are ranked depending on their overall (potential) technical impact; for example, a denial of service versus a privileged compromise. Prioritizing incidents by type can often result in too many “top priority” items being identified. Additionally, technical impact alone is not usually of interest except when a new, uncommon or not fully understood type of attack is discovered. As a result, this scheme is normally used in combination with other types of priority schemes.

Example: Five new incidents are reported with root compromises. All should be handled as soon as possible, since they are considered “top priority.” Two are from a major university and involve less than five hosts at sites where the staff is experienced in responding to incidents. One is a denial of service attack at a hospital affecting two hosts that hold a medial records database and laboratory test results data. The other two incidents involve hundreds of other hosts in your defined constituency, since the attacker is running attack scripts on other sites from compromised hosts.

The issue is to determine which incident to respond to first. Do you drop everything and deal with the biggest number of hosts? Do you drop the two major universities since they have experienced personnel available? Are other resources available within the team that can be utilized in the short term to help with the current reports? Can other teams be called on to help? Can you provide initial “first aid” to the hospital and follow up in more detail after you have dealt with the incidents that have greater technical impact?

Now, change this scenario a bit: What if the medical hospital were actually your sponsor’s site instead, and there were still the other incidents involving hundreds of other hosts within your constituency—how would this change your prioritization?

3.8.6.4 Feedback Request Prioritization

Generally requests for feedback can be handled differently from incident reports. The principle of “first come, first served” applies, but even in this case there may be a requirement to prioritize because of workload or other factors (such as available workforce and expertise). One method of prioritizing feedback requests is determining who gets a response first based on who is making the request. A request from a high-ranking official in the constituency or the team’s funding body will usually be sufficient to move the request to the top of the priority list.

3.8.7 Escalation Criteria

Escalation is often confused with prioritization. Although the activities are similar, escalation is concerned with raising the importance of an activity regardless of its priority. Escalation invariably requires at least one level of management to become involved for decision-making purposes. When escalation of one or more activities occurs, it is usually a sign that a team is experiencing an unusual or high workload and is under even more pressure than usual.

Escalation criteria and associated processes, procedures, and guidelines should be defined in advance in preparation for use. Here, too, there is a continuous need to regularly review the criteria and to adapt to changing needs and new developments, such as new attack styles, incident types, or even sponsoring organizations.

Note that escalation criteria can also be applied to the overall CSIRT service or the incident handling service.

Example: In late 1993 the CERT/CC began receiving incident reports related to network monitoring attacks and the capture of user name and password information. As isolated events, these reports were not difficult to analyze and handle. However, as the days passed and more reports were received with logs that captured thousands of user account and password combinations, the incident handling staff became overwhelmed with the scope of the activity. The resources for handling reports affecting many root-compromised systems and tens of thousands of captured user account/password combinations stretched the team's resources to its limits, and it was escalated to the director of the program. A strategy was outlined and approval given by management for obtaining and utilizing additional staff from within the program to assist with response activities.

Example: In 1999, the Melissa virus outbreak was an incident that was escalated in the CERT/CC. Prior to that activity, the CERT/CC did not focus their response activities on virus reports. Up to that point, viruses were generally not as widespread (e.g., spread by sharing infected files or affecting limited numbers of users or single entities). It wasn't until viruses leveraged the Internet as a propagation mechanism that this type of threat warranted escalation. This is an example of an event that caused the CERT/CC to change the way in which a service was provided to their constituency.

3.8.7.1 Individual Incident Escalation

Regardless of priority, it may still be necessary to escalate an individual incident. The escalation of an incident is normally the result of an incident handler being unable to address one or more aspects of the incident appropriately. The incident handler ends up needing additional support or management oversight, or to offload other work in order to appropriately handle the escalated incident. As an incident evolves and new information comes to light it may become apparent that the person to whom the incident is assigned does not have the technical expertise required to handle it appropriately, causing the need for escalation. By its very nature, incident escalation is driven by issues similar to those involved in incident prioritization.

Example: An email bombing incident is being handled by a novice staff member. During correspondence with the sites involved, new information is identified that indicates that the account being used to launch the attacks is itself compromised. The account contains password files from over 1,000 different systems. Given both the number of hosts involved in the incident and the staff member's lack of experience, the incident will require escalation.

Commonly used criteria for individual escalation include

- number of sites and systems under attack
- type of data at risk
- severity of attack
- state of attack
- source or target of attack
- impact on the integrity of infrastructure or cost of recovery
- attack on seemingly “secure” systems
- public awareness of incident
- new attack method used
- communication breakdowns
- technical ability, knowledge, and/or expertise of the individual CSIRT staff

It is also common for a team to have escalation criteria in place to simply notify management of unusual or potentially important situations.

Example: A local network service provider outside of your constituency sends details of an incident report to a public newsgroup. The report identifies connections that were made from the compromised system at their site to 1,000 remote systems. Some of the connections are believed to be the result of unauthorized activity. Due to limited resources and an inability to contact the registered users of the compromised system, the reporting site is unable to differentiate the legitimate connections from the unauthorized ones. Over fifty of the remote systems listed fall within your constituency. The incident should be escalated to management immediately due to the possibility of media attention related to the activity.

Communication breakdowns normally result from a complaint (whether valid or not) by a constituent or other party to the team. The constituent may not be happy with the way an incident is either technically or procedurally being handled or may have a specific complaint about a staff member. In such cases where the team’s reputation is at stake, escalation to management is advisable.

When incident information is missing, a team may be unable to make progress. In many cases this is not a concern and the team will follow up on the incident using the partial information available. However in other cases, lack of critical incident information is cause for concern. If in such cases the team believes that the critical information exists, but has just not been passed to the team, the incident may be escalated to seek additional techniques (or leverage) to try to gain the information.

Example: A site within the team's constituency is an ongoing source of intruder activity. The team has repeatedly made email and telephone requests to the site for more information, but none has been provided. Escalation may allow the team to exceed its usual levels of service by sending a team member to the site. Another approach to escalate this incident would be to inform the site's management about sanctions, such as blocking their network connectivity if they do not react in a reasonable way as outlined in their acceptable use policy.

3.8.7.2 Multiple Incident Escalation

From an incident handling service perspective, escalation criteria must also take into account additional factors, including the overall workload a team is experiencing, the requirement to meet its mission, the need to obtain and retain how the incident fits into the bigger picture of overall incident activity, and the additional resources available to the team.

There are times when a team has more incidents than it can possibly handle or it is unable to meet its published response timeframes. These situations sometimes arise gradually as the rate of incident reports increases. At other times there is a sudden sharp peak in incident reports. In either case escalation is applicable to enable the team to cope appropriately with the situation.

The actions (often applied simultaneously) resulting from escalation are for each team to determine. Possibilities include applying additional resources (e.g., extending staff working hours or calling upon others to support the CSIRT) or reducing the level of service provided.

Example: During the Y2K rollover, the CERT/CC manager negotiated with other members of its parent organization for additional technical and/or administrative support. Additional technical staff was used to help catalog information, contact sites by phone, and develop content for Web and email documents. The additional administrative support helped answer the CERT/CC hotline. This additional staff enabled the team's incident handlers to continue to focus on any incident analysis or response efforts that were needed.

When a CSIRT escalates an incident by seeking additional resources it should follow established and agreed upon guidelines for obtaining such resources. Plans and procedures should have been previously discussed with the other business areas that will provide the additional staff. The staff should be identified in advance and trained so that they are ready to help out whenever needed. Having pre-established contacts and methods for requesting assistance in place will facilitate the escalation process. CSIRT staff will be able to continue to focus on the ongoing incident activity instead of trying to determine who to call and how to train them. Possible resources might include

- other employees within the CSIRT, but external to the IR group
- other employees within the parent organization, but external to the CSIRT
- other teams, external consultants, or experts
- constituents or volunteers

Depending on the skills and expertise and demand for the group(s) chosen, obtaining help might be easier or more difficult to arrange or negotiate.

Teams are often faced with the prospect of reducing the level of service provided in response to incidents as a result of escalation. In such cases it is important to decide if the escalation should apply to all incidents or if incidents of a particular type can be excluded. In some cases the level of service may be reduced to a team providing direct, immediate assistance to the victim(s) and no more. Although this may be a necessary step to enable the team to recover to a steady state, it will also have other implications. In particular, it will adversely affect any attempts the team might normally undertake to identify the perpetrator of a particular incident and might also limit efforts in the analysis of techniques and mechanisms used by the intruder.

One major benefit of coordinating the response to incidents is that the CSIRT is able to develop, see, and interpret the bigger picture, as discussed in Section 3.4.2.1. This picture by itself is an important service to the constituency. But it also serves as an indicator on which to base immediate and future resource management decisions. Therefore, a detriment in reducing normal services and losing your grasp on the bigger picture is especially serious during escalation, at just the time when the bigger picture is critical to both the team and its constituency. Wherever possible care should be taken to retain the necessary level of analysis on those incidents that could be critical to maintaining the bigger picture.

When the team is in crisis mode and all resources are consumed by the workload or some other unexpected event, it is important to return to normal operations as soon as possible. Fixed criteria should be established to determine when a crisis is over. This will relieve the stress levels in team members and allow them to regroup, reprioritize, and get back on track with the regular tasks at hand that may have been suspended during the crisis.

3.8.8 Information Disclosure

For a team to be able to operate at all, it must disclose information. However if disclosure is conducted inappropriately, this routine activity can result in the team's demise. To prevent inappropriate (wrong, not allowed) disclosure, all information disclosed must be in line with the team's disclosure policy. Since this policy is critical for the perception and success of the team's operation, it is handled in much more detail in Section 4.2.3, "Information Disclosure Policy," while more general practical issues are discussed here.

There are different reasons that information may need to be disclosed and different groups or organizations that will receive this information. The process for disclosing information will differ depending on the group receiving the information and their plans for the information. The following is an example of the different types of groups who may receive information and their reasons for receiving the information.

Information may be disclosed to

- other teams about a new vulnerability that has been discovered
- other teams who are collaborating on incident analysis or response efforts
- sites that are the target or source of an attack
- management for budget purposes
- management for statistical reporting purposes
- a risk management group, to help in planning infrastructure and security improvements
- the funding body or shareholders for justification of CSIRT activities
- law enforcement for investigation or prosecution
- governmental organizations for notification or further reporting
- everybody who has a vested interest, so they are aware of ongoing activity and recommended mitigation or prevention strategies

The need for disclosure can result from requests or reports. Disclosure can also result from events that force specific actions, such as the publication of alerts or advisories. The CSIRT's disclosure policy will need to take into account the circumstances relating to both the type and reason for the disclosure.

Example: Whenever there is a large-scale attack targeting sites in the constituency, the CSIRT will inform the whole constituency rather than only the known victims. Usually the source of such an attack is not identified (nor might the sites targeted be identified). However, sometimes there is justification for disclosing the origin of attack. Examples include when knowledge of the origin is essential to stop the attack, when the origin is not willing to take corrective actions, or (in case of a real emergency) when the team's resources are so stretched that the only way to minimize or contain the damage is to provide as much information as possible about the attack (including preventive and reactive measures) to the constituency and let the sites handle it themselves. An important rationale for informing the whole constituency is to alert system, network, and security administrators; along with general users, to watch for and report suspicious activity that might otherwise have gone unnoticed.

When defining policies, a minimalist approach should be used. For most interactions and disclosure, it is not necessary to reveal the whole set of information because only part of it is

really needed. Therefore, the policy statements should decide between “need-to-know” as a default and full disclosure in justified and closely defined exceptions.

Example: Even if a new artifact is given to the CERT/CC by DFN-CERT, no information is disclosed relating to the site from which it was collected. On the other hand, the source may be disclosed if no need exists to hide it or if the source was public such as a USENET newsgroup. If the CERT/CC needs more information from the source (if it is a site) to analyze the artifact, it will request this information (providing reasons why this is necessary). If the reasons are valid, DFN-CERT will contact the site, explain the situation to them, and seek their permission to disclose the requested information prior to divulging the site’s identity. Usually permission is given, but still, it is important to ask first. Certainly it would be much more useful to have such information as early as possible.

The disclosure of information can take many different forms, each with different associated tradeoffs or benefits. In Section 3.5.1, “Announcement Types,” we discussed the announcement types (heads-up, alert, advisory, for your information, guideline, and technical procedure). These are generally public announcements. Information disclosure is clearly broader in scope than these examples might suggest. One could add many items to the list, including incident reports (aimed at specific incident-involved constituents or fellow teams) and internal reports (e.g., for management).

Because the policies will also affect others, the best way to avoid misunderstandings and problems is to define information disclosure defaults suitable for all situations. If there is a choice, the data required to make the choice should be requested before the actual situation demands it. This will avoid any additional delays.

Example: AusCERT initially implemented a registration process in which sites were asked if AusCERT could pass their contact information to other CSIRTs whenever the sites were involved in an incident. If a site did not wish to be contacted in regard to a specific incident, the site informed AusCERT, and the contact information was not passed on to other CSIRTs. If the contact information was needed at a later date for whatever reason, it was possible for AusCERT to go back to the site and request that permission.

Privacy issues relating to a site’s contact information and information about victim sites are obvious. Defining suitable policies and being sensitive to local laws will help to avoid many problems.

Some CSIRTs provide forms for submitting information to them. Usually forms make it easier for both the reporter and the team to obtain the relevant data, although there are some tradeoffs. While the reporter of an incident or new vulnerability will be asked for answers to

many questions, this is much more information than they would normally provide without the use of a standardized form.

Example: The CERT/CC makes its incident and vulnerability reporting forms available on its Web server⁴⁶ [CERT/CC 1997a, CERT/CC 1996].

Sometimes teams or organizations place specific requirements on their constituency to fill out forms or to report a defined set of information. Depending on their policies, the team or organization may accept incomplete or informal information.

The generation of statistics and trends is one of the most interesting services provided by CSIRTs beyond merely incident response. Because of their specific role, they are able to

- develop an overall picture of incident activity for the constituency
- provide the funding body with additional background information
- provide a better service to the constituency
- raise awareness with pragmatic projections

Part of the mission of any CSIRT is to make the best use of the information it has collected to serve the interests of its constituency. It is important for a CSIRT to think about the information it will collect and strategically plan how it will use that information, to whom the information will be distributed and disclosed, and what information disclosure policies and procedures will apply to its various interactions or collaborations with others. For example, the disclosure of CSIRT statistics to FIRST is discussed as a possible requirement with each prospective member. The prospective member will need to think about what information can be disclosed and in what format.

One last issue related to disclosure of information is standardization. As the disclosure process can be the most visible task of each CSIRT, great care must be taken to provide a unified and high profile interface to the “world,” especially the constituency and other CSIRTs. The way that information is distributed should be consistent over time (e.g., so comparisons with previous statistics can be made). Additionally, standardization will ensure a consistent “corporate identity” for the CSIRT. (If a CSIRT is located within another organization, the requirements of this parent organization will have to be considered.) Items to consider as part of this consistent interface are

- format of text provided, regardless of whether the text is distributed via mailing lists or through online information servers (headers, outline, footers, logos)
- authenticity, through formal signatures
- content and style guidelines

⁴⁶ http://www.cert.org/nav/index_red.html

4 Team Operations

In Chapter 2, we discussed the main functions that an incident handling service is built on, their interactions, and the handling of information. In this handbook we set out to explain what it takes to build an incident handling service. Compare this with building a house. We showed you a drawing of the house. We described the rooms and their purposes. We discussed staircases, corridors, electricity, telephone, heating, and water systems. What we haven't covered yet is how the building is operated and secured: maintenance of the heating and other systems, the annual chimney sweeping, the insurance procedures, and fire-alarm and burglar-alarm procedures. These are the "operational" components of the house, and in this chapter we describe these CSIRT operations in more detail.

Beginning with an introduction to the main operational elements, this section will also cover four essential operational issues: fundamental policies, continuity assurance, security management, and staff issues.

Many of these topics are not exclusive to incident handling services. Therefore it is not surprising that some aspects of these issues have already been covered in Chapter 3. Where appropriate, we will refer to that section instead of repeating the information. However, in this section we give a more practical approach than was possible in the "policy" level in Chapter 3. This chapter will cover useful general *procedures*. (Remember, procedures are the implementations of policy statements.)

4.1 Operational Elements

Operational elements are the building blocks of operations that span a wide range of ideas, from email systems to work schedules. We limit ourselves to those elements that bear a direct relationship to incident handling services, thus excluding all operational elements belonging to overhead, such as salary systems or coffee machines. The list of elements covered in this section is far from exhaustive. We will only discuss a selection of the most important practical operational elements. Where appropriate we provide real-life examples and include more detail on particularly important issues.

4.1.1 Work Schedules

A work schedule must differentiate between normal hours and out-of-hours; it includes such things as work shifts (including associated personnel), out-of-hours arrangements (like guards or operators providing answering services), backup, and all-hands-on-deck arrangements.

When considering work shifts, a good rule of thumb is to remember that after about two hours of routine work, you definitely need a break; but that after only one hour of concentrated stressful work (such as being in the midst of a big incident), you are devastated. When planning work schedules, continuity assurance (discussed in more detail in Section 4.3, “Continuity Assurance”) is the most important goal in relation to the quality of service provided.

4.1.2 Telecommunications

This includes “traditional” telecommunications like telephone, fax, cellular (mobile), pager and automatic response facilities. You will need this kind of technology (and other communications) to ensure that your organization and its members can be reached in accordance with your requirements *and* that your staff members have the technology available to initiate communications to the constituency and/or other parties as required. Implementation depends on the team’s mission and service specification.

Remember though that there is no such thing as guaranteed communication. Even the telephone system fails every now and then. If you cannot hear the telephone ringing, even the most costly and fancy technology won’t help. Similarly, way down in the Grand Canyon (or some long, dark tunnel), you are not likely to have a working cellular phone. Be aware that constituents will generally be displeased if they have to wait for more than four consecutive rings when they try to access your service by telephone, and if they consider their need urgent they will be even more displeased to reach a voice message system. If they have left a message and do not receive a quick response (for example, within 15 minutes) they are likely to still be displeased. A voice box might be useful to provide a first acknowledgment and further information about what to expect. Modern devices will contact a predefined number after receiving a new call. In this way the caller does not need to know the number to reach the CSIRT team member.

4.1.3 Email

The need for a good email system in today’s networking environment needs no advocating. It is possible to create an easy-to-use, robust email environment that is compliant with up-to-date standards for multimedia (MIME) and security (PGP, S/MIME). However it is by no means a simple exercise for a CSIRT, as an incident handling service will have additional

demands, such as good filtering capabilities, advanced search facilities and automatic response tools.

Usually CSIRTs build their own email environment based on a few standard tools, gluing these together and adding to them using scripts because there are no products available to fit their specific needs. Additionally they use a variety of converters (such as MIME, binhex, uuencode, zip, or gzip) and word processors because some users might use a PC office package to write “pretty” emails that are definitely not ASCII-compatible. As technology evolves, it is likely that such compatibility issues will become more transparent for seamless use. Nevertheless, it is important to consider the need for an interface between the email environment and other environments to handle the workflow. Without such an interface, most of the incoming information will not be integrated autonomously and automatically. Email provides an easy-to-use technology to exchange information asynchronously; and by prioritizing incoming email, CSIRT staff are able to handle their work more efficiently. In fact, it is not as time consuming as a telephone call in many cases; but in some cases, electronic means cannot replace direct communication. In any case, be aware that constituents will expect feedback in a timely fashion, whatever communication methods are used.

4.1.4 Workflow Management Tools

In any operational environment with heavy workloads and people working in shifts, tools that help to manage the workflow and hand-over of ongoing tasks are essential. The hand-written logbook is a classic example. With the complexity of today’s problems and the sheer amount of information involved, this kind of logbook should really be obsolete (but it is not, unfortunately). CSIRTs need a workflow management software tool (essentially a database with a program on top) that enables you to follow and add to the flow of events (such as incidents, requests, or ongoing analysis). Excellent workflow management tools are on the market now, working with the usual databases. However, security of these systems is normally lacking; so as a rule they are only useable within a secure intranet, which may be a problem for off-site coverage or distributed CSIRTs. Integration with email tools, the Web, and the telephone system (and pagers) is necessary to collect all incoming information and to interconnect events.

4.1.5 World Wide Web Information System

The World Wide Web (WWW) is ubiquitous and currently the hottest medium in use for retrieving information. Certainly no team could do without it. Existing anonymous FTP directories are still useful to provide access to large archives with programs and documents. One improvement, however, is that they are at least made accessible through the Web, and that Web-based information can link to them. Clearly Web servers and any public information server for a CSIRT (providing public information) must be implemented in a secure fashion

to avoid the information being manipulated by unauthorized parties. The latter requirement, however, opposes public availability. One possible way of avoiding this contradiction is by placing the Web server within a DMZ (demilitarized zone) protected by firewalls and maintaining its security by good maintenance and control measures. To secure the authenticity and integrity of the information maintained, it might be useful to maintain the information on a master server located inside and download it to the Web server on a regular basis, such as every night [Kossakowski 2000]. Additional checks based on cryptographic checksums (such as Tripwire or MD5) are useful too. If internal information is made available to team members, each of these pages and all links pointing to this information must be removed prior to any public dissemination.

4.1.6 IP Addresses and Domain Name

By separating your internal network from all other networks for security reasons, you will require ownership of IP address space dedicated to the team. With a Classless Inter-Domain Routing (CIDR) block of the Internet IP address space, it is of no consequence if only your team uses these numbers and not some other part of your parent organization.⁴⁷ Alternatively, a team might choose to use some private address space (e.g., 10.0.0.0) and either network address translation (NAT) or a firewall for all external connections.

The Domain Name Service (DNS) shouldn't list sensitive information such as the type of operating system that a particular host is running or give out a complete list of all internal hosts, because this might reveal information useful for a technical or social attack. In most cases, it is appropriate and helpful to claim an Internet domain for the team to promote the existence of the team and to provide an easy-to-remember interface for email or Web. Your domain space will typically be of the type `company-csirt.some-org.tld` or `company-csirt.tld`.⁴⁸

4.1.7 Network and Host Security

An incident handling service's internal computers, network, and the connection(s) to other networks must be securely configured and protected against attacks. This means splitting the internal network into compartments with different functions, with the interface to the outside world through a mature firewall. At least two compartments should exist: an operational network, where all service tasks are handled and the data used is stored, and a testbed (unless you perform no testing at all). Compartments should be separated, only connected to one another through a firewall and only when data transfer is necessary. Should a temporary connection be established with the testbed, be careful that it is truly temporary. Most of the

⁴⁷ That is, the IP address space for your CSIRT would be different than the IP address used by the parent organization. This would, of course, require associated effort for obtaining and registering the CSIRT domain name and IP address, as well as any associated fees.

⁴⁸ The "tld" stands for top-level domain, such as .com, .edu, .org, .nl, .de, or .uk.

time it is not necessary to connect the testbed to other networks at all. If it is too dangerous to connect the test network to other machines or networks, data can be transferred by using removable media.

The firewall selected will undoubtedly be influenced by the budget available. Typically, a dual-screened firewall will provide a high level of security; this type of firewall consists of one router serving the outside world, one router serving the inside compartments, and one or more bastion hosts to interconnect internal clients to external servers by application-level gateways (proxies) to prevent inside clients or servers talking *directly* to their outside counterparts. In addition a DMZ, protected by the firewall to the inside as well as to the outside, should be used for any server providing an accessible service to the public (WWW server, ftp server) or act as a gateway or transfer system (proxy server, email gateway).

It hardly needs arguing that of all organizations a CSIRT must have its systems more than up-to-date with regards to security patches and updates. Logging facilities, wrappers, and a variety of other defensive tools should help to identify and prevent intrusion attempts. But even the security of home systems and access from home systems and laptops must be considered, if they are used for sensitive work.

Denial-of-service attacks form a special category of attacks that should be considered carefully, since they affect the ability of the team to perform its tasks. Having network connectivity available through more than one service provider can be part of the answer to this problem. At least that way, when the main entrance is blocked, you can use the emergency exit to maintain at least minimal communications such as email. Section 4.4, “Security Management,” will deal in more detail with security management for a CSIRT.

4.2 Fundamental Policies

A number of policies are “fundamental” (i.e., independent of the set or level of service(s) chosen by a team) and must be in place. Basic issues were discussed previously in Chapter 2, and some examples of service-specific policies were discussed in passing in Chapter 3. This section will discuss in more detail fundamental policies for the team’s operation. Keep in mind, however, that since it is most likely that service and quality specifics will affect the content of fundamental policies, the discussion below is generic and includes only some examples for clarification.

4.2.1 Code of Conduct

A code of conduct for an organization is a set of general rules indicating how to behave in a way that benefits and supports the intent of the organization’s mission statement and the organization’s character. A code of conduct applies to every staff member at all levels in the

organization; it is an attitude, and attitudes should be classless. It provides basic direction about how to react in certain situations and sets the foundation for interactions both within and external to the team.

The code of conduct is a policy that one can fall back on when all other policies, rules, and procedures don't seem to apply, or when one is left without time to consider. It should become a natural professional behavior of the experienced incident handler; novice CSIRT staff will need mentoring to internalize the accepted code of conduct for the team.

A code of conduct need not be more than one page of text, but there is no harm in it being longer and explained by example. If it's too long, it probably contains procedures, which it definitely should not. The advantage of small size is also ease of communication internally and externally. Since one cannot be ashamed of one's code of conduct, why not publish it for the constituency and fellow teams? This will also help form the type of basic understanding needed for collaboration between teams.

A very simple code of conduct example (that complements the CSIRT policies and procedures) is shown below:

Demonstrate due curiosity, but at the same time ...
show proper restraint.

Thoroughly inform those who need to know, but ...
do not gossip.

Take due care, but ...
do not forget priorities.

Always be polite and constructive, but ...
trust nobody without proper verification.

Know the procedures and follow them, but ...
never forget that the mission comes first.

This example is almost poetic in nature; form and choice of words totally depends on the type of organization. Remember that it's the organization's mission statement and character that decide the code of conduct. Another interesting example is the CSIRT Code of Conduct. This code of conduct (illustrated in Figure 6) is adapted from a list of behaviors that were developed in January 1991 by Rich Pethia, the first Manager of the CERT Coordination Center.⁴⁹

⁴⁹ CERT Coordination Center. "CSIRT Code of Conduct." Materials from the course *Managing Computer Security Incident Response Teams (CSIRTs)*.

4.2.2 Information Categorization Policy

CSIRTs must have a policy on information categorization. Without one, CSIRT staff will apply their own perceived categorization to each piece of information, or not attempt to differentiate at all. As individual perceptions may differ, resulting in inconsistent and possibly inappropriate service, a policy must be available to guide categorization.

The complexity and size of this policy will depend on the team's mission and constituency. For instance the simplest case would be just a division between "sensitive" and "all other" information. All sensitive information should be treated with extra care while all other information is considered public.

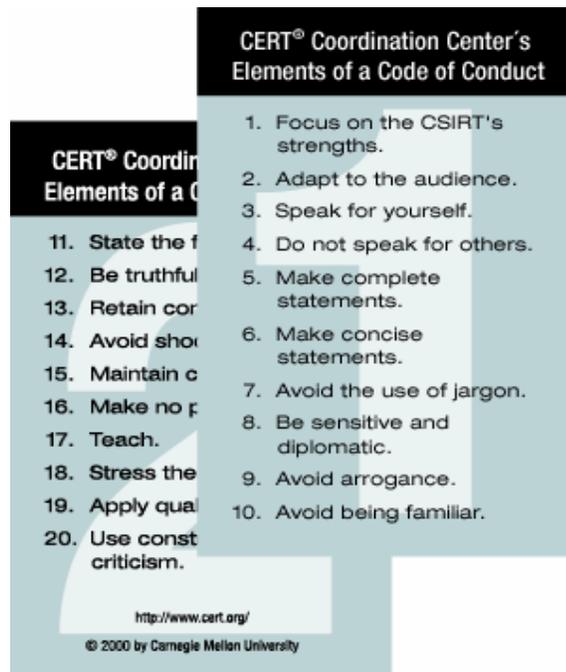


Figure 6: CERT/CC Code of Conduct

A slightly more elaborate scheme could define the categories "internal classified" for use only within the team; "internal unclassified" for exchange on a need-to-know basis with fellow teams; "external partner," for interaction with the constituency and fellow teams; and finally "external public," for public information. This is the approach taken by CERT-NL and detailed in their operational framework [CERT-NL 1992].

The CERT-NL scheme has the disadvantage that the distinction between "internal classified" and "internal unclassified" is not always clear in real life. A better approach might be to change the terms into "fully classified," "partly classified," and "unclassified." The main difference is that the strictest category really only allows communication within the team, whereas the "partly classified" category features the "need-to-know" principle coupled with

an enumeration of more-or-less trusted communication partners including other CSIRTs, listed in order of level of trust. But this discussion should not imply that this issue is a matter of names. The only thing that really matters is that everyone follows the same method of categorization. A pragmatic way of setting an initial scheme in place might be to ask the team members separately to categorize some documents, review and evaluate how each team member categorized or ranked the documents (to understand the rationale for what each did), then reach consensus on developing a scheme that everyone can support and use.

Military teams, for example, are expected to have the entire range of military information security grades (up to “top secret” or “state secret”) in place, complete with extensive procedures for every category on how to deal with information.

Note also that the category selected will affect the way the information is handled (e.g., storage, disclosure, and disposal). As a result, policies and procedures must be developed for each category. Then, regardless of the content of the information, this consistent set of policies and procedures applies to all instances of that category. All policies and procedures for operational tasks should include statements on how to deal with each category. This will include specifying default categorization values. It makes a big difference whether a default is “public” or “internal.” The defaults may differ for different types of information, and how they are handled within each category will also necessarily differ.

Example: The default categorization for contact information may be “internal” and differ from the default of “public” for publicly released advisories issued by other CSIRTs.

Sometimes it is not clear what category information should be placed in, since it might be considered as a candidate for more than one category. The old adage “Better safe than sorry” applies here. Within the CSIRT environment the category chosen is normally that which ensures that the information has greatest protection. If at a later time new details become available to indicate the information has been incorrectly categorized, it can easily be recategorized.

4.2.3 Information Disclosure Policy

One of the most important issues that a CSIRT needs to pay attention to is how it is respected and trusted by its constituency and other teams. Without that trust and respect, a team will not be able to function successfully and effectively, as people will be reluctant to report information to it. It is important to define an information disclosure policy for the realm of incident response and beyond. Without such a policy, CSIRT staff will have no guidance on what they can say to whom and when as they handle calls and respond to email.

Most teams treat all information reported to them in strict confidence and do not share the information beyond the scope of their immediate team members. Exceptions to this guideline

include using generic information for trend and statistical purposes or in cases in which the sites and parties involved have given their consent to disclosing the information about themselves or their site to other specific parties (such as other sites involved in the incident, law enforcement, or other response teams coordinating the response to an incident).

This policy should take into account the information disclosure restrictions that might be placed on information provided to a CSIRT by other organizations and the parent organization, which might have its own requirements (in some cases, even legal requirements for external audits). For example, if another CSIRT reports an incident, what can its constituents expect regarding the disclosure of the information reported? Will it be reported to law enforcement or the CSIRT management? The policy should specify limitations, which should be made publicly available (to the constituents and other interested parties). Under what circumstances must a team pass sensitive (even contact) information to law enforcement or a court? At the current time, we are not aware of any requirements for CSIRTs to have a similar legal status regarding client confidentiality (such as those for doctors or lawyers).

Example: Consider a scenario in which CERT-NL provides the CERT/CC with information about an incident. Say the incident took place at an educational site in the Netherlands from which the intruder launched a successful attack against a system in the U.S. CERT-NL will pass logs and timestamps to the CERT/CC and request that they forward the appropriate details to the U.S. site. They will also indicate whether the information can be passed to other sites involved in the incident. Additionally, CERT-NL may provide the name of the Dutch educational site and the contact information for the system administrator at that site to the CERT/CC, with the understanding that the name and contact information is for the CERT/CC's use only and not for further distribution.

This additional information helps the CERT/CC to understand the bigger picture and related activities. As a result of this information, CERT/CC is then able to correlate this incident with other possibly similar and/or related events involving Dutch hosts (some perhaps not yet known to CERT-NL). After CERT-NL receives this additional information, they in turn are able to help three other sites in their constituency to recover root compromises that were not previously detected (and that were subsequently identified based on the information received from the CERT/CC).

In addition to information disclosure restrictions on information provided to a team, the information disclosure policy also needs to take into account requests from others to receive information. Commonly, such requests are for detailed technical or sensitive information.

There are essentially three factors that determine if, to what extent, and how information will be disclosed. These are the *purpose*, *target*, and *category* of the information.

1. Disclosing any chunk of information needs an underlying purpose; in other words, someone has a “need to know” this information. This need-to-know principle can be applied to all information.

Example: Warning managers of a site that their machines may have been compromised by an http-daemon vulnerability because they are using a vulnerable software version requires only a bare minimum of information. No break-in-specific information is available, and the information needed relates to the vulnerability itself, available workarounds, or patches.

Example: However, if an incident involves break-ins through a backdoor installed by the vulnerable software, it may be necessary to provide the relevant log, timestamp, and the originating IP address information, thus revealing some contact information.

The purpose and extent of information disclosed is different in these two cases.

2. The target of the information is whom it concerns, e.g., members of the CSIRT constituency, other CSIRTs, internal management, law enforcement, media, visitors, experts, or the public.

Clearly one is going to be much more restrictive when handing information over to the public than one would be when communicating with a trusted fellow CSIRT.

3. The category of the information is decided by the information categorization policy (as discussed previously).

When it comes to deciding whether or not to disclose the information, it clearly makes a difference whether a bit of information is “internal” (e.g., contact addresses of constituents) or “public” (e.g., advisories). This category will affect the way that the information is protected. For example, public information might be relayed through normal email, which is only protected by the authenticity of a digital signature, whereas internal information would prescribe the use of encryption or a secure channel.

Suppose there is a clear purpose in disclosing some particular information. If a decision is made to disclose, the information, category, and target factors will determine how disclosure will take place, what pieces of the information will be disclosed, and to whom.

Example: Consider a large-scale incident, with intrusions involving thousands of hosts all over the world. As a result of the incident, several detailed log files have been provided to your team by sites involved.

- For the CSIRTs and sites you have a trusted relationship with, you might hand over those parts of the log files that relate specifically to them or their constituency.
- To other victims, you pass on the relevant (for their site) log entries to enable them to check their own logs, together with guidelines on how to protect against future attacks of this kind.

- You might warn law enforcement by telling them about the size and spread of the event, plus generic exploit details.
- You may tell the media about the size and spread of the activity, with a warning and some comforting words.
- To trusted experts you might give all the gory details (sanitizing the information to site-specific information) so that they might learn more about exploits, trends, and signatures.

4.2.3.1 Second-Level Disclosure

When one entity discloses information to another, it is likely that the latter will spread the information further. In some cases (e.g., the media) this is obvious; in other cases, less so (e.g., internal management). It is important to agree with the target of disclosure on what this target is allowed to do with the information. Once the information is handed over, it is out of one's control. And even if a binding contract exists stating what the target is allowed to do with the information, it can still leak out (e.g., through a security breach), and the originating party can still be affected by the repercussions (damage to reputation or even lawsuits).

Example: With the media you can request/require that a draft is sent back for your comment/approval before publication.

Fellow teams are often given detailed information, under the (often tacit) assumption that the information will only be used on behalf of the teams' constituencies and is not to be spread beyond.

One helpful approach for others is to place a label on disseminated information clearly stating the expected use of it (for example: "For internal use within the CSIRT only"). This is particularly helpful when exchanging sensitive information with others.

4.2.3.2 Timing of Disclosure

When is one going to disclose certain information, or *how soon*? On the one hand, it is nice to be certain of the facts before disclosing anything, which often takes a lot of time. On the other hand, likely victims should be warned as soon as possible, even if the information is not yet quite complete or correct. Interestingly enough, both extremes may lead to lawsuits, especially if a team has very explicit contracts with its constituents.

Example: The constituents of a commercial CSIRT may become very upset when they experience problems that might have been prevented had their CSIRT acted more quickly and given them a heads-up. Being given inadequate information that leads to systems going down, or still being vulnerable in spite of the CSIRT's warnings may also cause the constituent to file a complaint or lawsuit. This kind of behavior from the constituency is less likely to strike a team who has no authority over, or contract with, its constituency

and is primarily providing advice to others (such as a national team, the CERT/CC, or even some large, distributed corporate teams).⁵⁰

4.2.4 Media Policy

The media can be a powerful and useful tool to publicize and disseminate important CSIRT information and you'll want to develop a good rapport with them. Nobody needs convincing that it's good to have a media policy. Even if a very detailed information dissemination policy exists, handling the media is especially difficult.

The main issue to consider is *where* will the primary interface to the media reside? Will it be internal or external to the CSIRT? For teams dealing with both highly technical and sensitive data, as CSIRTs and related teams do, it is advisable to have the media spokesperson external to the team. This way the spokesperson has little or no access to sensitive data, since they only know as much as they need to know to fulfill their function to notify the media, according to the information dissemination policy and media policy. Usually this information is heavily sanitized. Such a situation avoids the danger of too much being told to the media and potential lawsuits. If the spokesperson is external to the team, someone within the team must be responsible to ensure that the spokesperson receives continuous updates about what's going on.⁵¹

Establishing a List of Media Contacts

To avoid having publications written by disreputable or poor-quality journalists, or appearing in the "wrong papers," it is useful to screen several media contacts and their papers or magazines before putting them on a list of media contacts that you're willing to work with. You should actively target good technical journalists and publications that you would like to work with. Many publications have good people on these jobs; however, security is still often a weak spot. Part of the collection of contacts must be devoted to means for strong authentication and understanding the (technical) background of the journalist and her/his agenda.

Providing Rules of Engagement

These rules inform media contacts of what they can expect from you and how you expect to interact with them. Do not hesitate to make clear what you expect from them, such as

- Only contact the CSIRT's designated media spokesperson(s).
- Do not falsify quotations or citations.

⁵⁰ It should be recognized, however, that having no authority does not preclude someone deciding to file a lawsuit if the team is in a particularly litigious part of a country or region.

⁵¹ McGillen, Terry. "CERT Incident Communications." 5th FIRST Workshop on Computer Security Incident Handling. St. Louis, Michigan, August 1993.

- Provide the spokesperson a chance to comment on, edit, or approve an article before its publication.
- Any violation of these rules will result in removing the media contact from the media contact list.

Briefing the Media in Advance

Taking the lead instead of waiting for the media to come to you can save a lot of time not having to explain actual developments over and over again. Advance briefing also allows you to prepare for questions that might have otherwise unnerved you. Going one step further, ensure that the media knows the mission of your team and give them a global sense of how this role is performed. Also use these opportunities to spread proactive messages.

Specifying Out-of-Habitat Behavior

Team members and their media spokesperson are likely to appear in public. They do not suddenly become invisible when there is media attention. So they should be prepared to face the media at any time. When unexpectedly faced with the media, the simplest solution is the “no comment” approach. While this solution is acceptable for all team members, it is not a feasible option for the designated spokesperson. A more elegant (and difficult) approach is to train the team members in media interactions and help them understand what they *can* say in public, instead of what they *cannot*. This is a more positive approach, and as a result they will project a more positive image for the media, even if they were not briefed in advance for a particular situation.

Providing Outreach Through Announcements

Using the predefined contact list, up-to-date briefings can be distributed before other public dissemination to provide media contacts with background information about ongoing developments, as appropriate. Additionally, this list can be used to send a heads-up (say for a new service offering) or invite them for a detailed briefing to alert them to an upcoming event or conference being organized by your CSIRT.

4.2.5 Security Policy

These days, every self-respecting organization has or claims to have a security policy, embracing all security aspects ranging from locks on the doors to backups, passwords, firewalls, and encryption. Handbooks have been written about how to write security policies [Wood 1998, RFC 2196].

Instead of doing a bad job at emulating those efforts, we will only highlight those aspects of security policies that are especially relevant to the readers of this handbook.

First of all, one must consider the fact that CSIRTs and the like cannot choose but to operate in networked environments, which make them fundamentally vulnerable to attack. Add to

this the fact that CSIRTs are also very popular targets for intruders, and the prime risk factor is clearly outlined: a team that's suffering from an intrusion loses its ability to (re)act, and also the trust invested in it if the situation is not controlled in a swift and professional manner.

Attacks on the systems of CSIRT might be motivated by the fact that as CSIRTs are high profile, they are sought-after targets for a wide variety of intruders. Novice intruders see them as an attractive target, and they are of great interest for the professional intruder, since the intruder might find information on companies that have experienced everything from denial-of-service attacks to mission-critical intrusions, and much more.

The security policy is heavily affected by other policies, because their goals must be protected by the security.

Example: The information categorization policy defines variables that also occur in the security policy and that set the level of protection for files and documents, which must be implemented using appropriate technology and established security procedures.

The security policy should cover all aspects relevant for the team's computer and network and also consider the connection to other networks:

- physical security
- recovery planning (backups, etc.)
- local network security
- local information security
- external communication security
- handling of local security incidents
- disaster handling, business continuity

4.2.6 Human Error Policy

We are all human; therefore we all make mistakes. It would be nice to think that CSIRT staff could be immune to this trait. However, they are particularly vulnerable as a result of the high-stress situations in which they are placed and the responsibility associated with the nature of the information that they handle.

Unfortunately, a human error policy that can be referenced and used when such mistakes occur is often neglected or not considered. A human error policy can help minimize and contain the damage inflicted by human errors. At the same time, it can give both the erring staff member and his/her management an opportunity to solve the problem in a professional and constructive way, instead of the all-too-usual strife and fear, which are counter-

productive. A human error policy should *not* say “Be as stupid as you want; we will always be nice to you.” It should clearly state what possibilities a staff member should consider if he has made an error that may have bad results; it should clearly state the proper reactions from management; and it should outline the consequences.

The following scenario might be seen as a general guideline for handling such occurrences: A staff member who did something that may have bad results should report it as soon as possible to the appropriate manager. Having an escape hatch to a trusted “third party” can be beneficial. With the error noted, managers and the staff member alike should put aside their sentiments for the moment and work *together* on containing the situation; keeping the wrongdoer aboard clearly is important (unless the act was obviously malicious). After the immediate problem has been addressed, an appointment between the staff member and manager (plus the trusted third party) must be made for the *next* business day. In that conversation the cause of the error must be jointly analyzed to avoid similar mistakes from happening in the future. If some bad habit or wrong perception of the staff member is the cause, it should be agreed on to change that habit or perception; checkpoints can be jointly defined to see if that agreement works out in the near future. Depending on the cause, training or educational measures might be most beneficial to allow the staff member to adapt to the position.

Here’s a more specific example:

Example: It’s the end of a hot week, and pressure and workload are high. A staff member incorrectly cut-and-pastes information about Site A into an email message intended for Site B. As a result, information is inappropriately disclosed. Action is taken promptly to inform management and Site A and Site B of the oversight. All parties are understanding. Methods are then sought to decrease the chance of this happening again (shorter shifts, easier-to-use tools, refresher training, or more coffee).

If mistakes by any staff member start becoming regular occurrences, then additional steps outside of the human error policy will be necessary.⁵²

4.3 Continuity Assurance

The continuity of consistent and reliable services is essential to the successful operation of a CSIRT. This directly reflects on the perceived competence and level of trust of a team by its constituency. Assuring continuity is a general operational issue covering many important

⁵² In risk management, there are known principles to deal with such situations. This includes “separation of duty” and the “for eyes only principle, where only those that need to know have access to certain information.” See, for example, Botha, R .A. and Eloff, J. H. P., “Separation of duties for access control enforcement in workflow environments,” *IBM Systems Journal* Vol. 40, No. 3, 2001, pp. 666-682.

aspects of operations, three of which will be dealt with below in separate sections: workflow management, out-of-hours coverage, and off-site coverage. Before embarking on this, it is useful to recognize the fact that the length of time for which one seeks to assure continuity may make quite a difference for the kind of problems encountered, the services that can be provided, and (thus) the measures to be taken. Here a division into three rough categories is used. Threats to the continuity of the team's operation are therefore reviewed before the more practical topics are addressed.

4.3.1 Continuity Threats

From a practical point of view, we make a division into three main categories to differentiate the threats that each team faces in relation to its continuity: short-term issues ranging from days to weeks, medium term within months, and long-term issues in years.

4.3.1.1 Short-Term Issues

Operational topics are mainly responsible for threats to the continuity within days or weeks. Four topics can be identified, which provide their own challenges and are responsible for most of the short-term issues: lack of time, unavailability of critical personnel, transitions between shifts, and unavailability of infrastructure elements.

Lack of Time

Lack of time can be incidental or structural. If it's structural (usually caused by lack of funding), it is outside the scope of this handbook and normally not a short-term issue. Incidental lack of time (e.g., due to an unforeseen workload by a new incident with widespread attacks) is dealt with primarily through prioritization. Prioritization has been dealt with in Section 3.8.6, "Prioritization Criteria." What it means for an incident handling service is that you let a sniffer log that is two weeks old wait a bit if at the same time all your attention is focused on an acute case of intrusion. If you don't have a predefined prioritization scheme, you will prioritize anyway, but it takes you more time to think about it, and you may be less consistent. Extreme lack of time may result in the need for crisis management, as it affects the team's service. When you have a lot of work at hand, it is helpful to make notes about what is going on. When the time comes to transfer the work to a colleague on the next shift or to a person outside your team, such as a guard or operator who will take over part of your work during the night, these notes will be of crucial value. Just taking notes on a piece of paper is the oldest form of workflow management, but it is still workable. Workflow management is treated below in more detail.

Academic CSIRTs are particularly vulnerable to incidental lack of time, which is caused by informal and not very precise resource planning. In addition, the time needed for dealing with the workload is underestimated and there are not enough spare time slots and no preassigned tasks to allow the team a break or to complete some unresolved tasks.

Unavailability of Personnel

Unavailability of critical personnel can arise at any time, because illness, accidents, and unforeseen events are inevitable. To avoid a single point of failure, backup arrangements for personnel should be made in advance. Team members should back up one another (e.g., buddy system). All members of a critical team should not be allowed to have the same day off. Regular job rotations may be considered to help spread knowledge and thus risk. Training to fit other needs gives personnel a perspective and helps to avoid such situations. Lack of critical personnel may arise during the time just before and after business hours. During that time most of the critical team members may be commuting to or from home. They may be reachable but still will have a hard time performing specific actions. This can be avoided by having team members “stagger” their business hours (one staff member works from 7 a.m. to 3 p.m.; while another works from 9 a.m. to 5 p.m., for example). If personnel are on a business trip, they might (or might not) be available to help out if their specific expertise is needed for some task. It is not much fun when your staff has to conduct critical business from a remote site like a conference, even if it might be seen as “thrilling” by an outsider. It raises a lot of problems, the impact on security being just one of them. Off-site coverage is discussed below, since it raises a separate set of issues. Another reason for unavailability of personnel is, by definition, out-of-hours. This topic is also addressed below.

Transition of Shifts

The transition between shifts poses special problems, even in the case where a good workflow management system is available. Depending on the circumstances, two cases should be considered: transitions between regular shifts during business hours, and transitions between business-hours and out-of-hours coverage. In the first case, some time must be reserved for a verbal transfer between shifts; “gut feeling” is often essential but hard to capture in any database. Sometimes events are not finished and open topics must be handed over. Additional explanations are necessary in these cases.

Example: Some teams conduct daily briefing sessions at shift changes to identify current activity (new and ongoing incident reports) and provide status updates (what has been done, when tasks need to be completed, etc.).

In the second case (for out-of-hours transitions), more facets of the same problem arise, because of the difference between both types of coverage. These differences include staff (e.g., regular staff vs. out-of-hours answering-service staff such as operators or guards). It may well be, for example, that the guards do not have access to the workflow management system, meaning that reporting forms will have to be transferred and analyzed the next business day. In addition, it is also unlikely that such guards will have the technical skills to respond to calls in the same way as the CSIRT staff.

Unavailability of Infrastructure Elements

Unavailability of critical communication paths and operational elements such as email

servers or information servers (WWW, anonymous FTP, etc.) will lead to the inability to provide specific services in a timely fashion. This could lead to complaints being lodged against the CSIRT or lawsuits for failure to meet contractual requirements and/or services, and in some cases, may have serious effects on the survival of the CSIRT and/or constituents.

4.3.1.2 Medium-Term Issues

For the medium term, the useful thing to do to help continuity is to get people together and analyze what has been going on, what went wrong, what went right, and how to use this information to make the service better. Post mortems after different types of incidents, to review activities, should be held to examine both policies and procedures and to determine where (or whether) improvements are needed. Other brainstorming sessions and meetings should be planned at regular intervals. This will highlight failures in policies or procedures. Another medium-term issue is the lack of funding and its influence on the team's operation and the level of service provided to the constituency. Staff burnout is also a serious risk to consider, especially in the strenuous incident handling arena (and whenever there is a lack of funding). Good work conditions will help to ease the burden on the individual. Encouraging staff to take their holidays and vacation time to help "recharge" and invigorate them is essential for the health and well-being of your CSIRT staff. Job rotation will help too; the latter will also help against staff boredom, which can also lead to staff burnout. Boredom is not unusual in handling incidents, not because of lack of work, but because of the repetitive nature of incident response tasks. Job enrichment and continued education also are good ways of motivating an individual. These can have positive benefits for the team, since your staff will develop new capabilities and further enrich the team's services.

4.3.1.3 Long-Term Issues

The ability to adapt to changes (e.g., in technology or service agreements) will affect the ability of the team to survive over the long term. Training of staff is therefore a long-term investment in continuity. Training more team members for the same functions lessens the impact of single points of failure, changing trends, or a team member leaving or falling ill. Section 4.5, "Staff Issues," covers this topic in more detail. One factor that is becoming more important over time is working habits, especially if the team hasn't changed much over time. By falling into some kind of routine drill, the situation is stabilized, but this doesn't ensure continuity. Stabilization may limit the team's ability to adapt to change; the team may be vulnerable to common mistakes that are ignored since the established procedures are accepted as is. The ability to react to the dynamic environment of incident response is a continuous learning process for both the team and its staff; flexibility is a necessity because change must become a way of life. So, although policies and procedures are critical to operations and must be in place to ensure consistent behaviors, they must also be routinely reviewed and validated to ensure that they are still viable and that the team is still following them, and to determine whether there is any need for change.

4.3.2 Workflow Management

Workflow management is just what it says: managing the flow of events that are part of work—your own work, a team’s work, a company’s work. Workflow management is applied at all possible levels, with all kinds of sophistication. A househusband will usually use only his agenda and his wits for managing the workflow. A company building cars will need a bit more workflow management. It is not surprising that much of today’s practice in workflow management stems from the logistics area, where this has been an issue for years.

Incident handling continuity problems arise as CSIRTs have to deal with a lot of problems over longer periods of time, because of the continually changing events and also changing team members working on these problems (due to shift changes, holidays, job rotation, and people leaving). For all problems, incidents, and related issues such as information on artifacts or vulnerabilities, the related information should be available, as appropriate, to all team members on duty at any time. In addition to the information about the problem itself, a track record of subsequent actions taken by the team, and what more needs to be done, should also be available. This facilitates the hand-off of ongoing incidents to the team member(s) taking over the problem.

Consider the common prime carriers for information coming into a CSIRT: email, files, faxes, telephone notes, and letters. How to make this available to all at any time is not an easy task. Applying numbers to incidents and tagging all information on the incident with some type of tracking number or ID is the very first thing to do; this point has been extensively covered in a previous section. That done, one could opt for the classical system: All paperwork (faxes and letters, possible telephone notes too) is indexed and archived, and all electronic files are numbered and stored (email and files are usually stored in different places). Then there may also be Web pages to consult, which makes the number of archives to consult (with regard to one incident) a discouraging four. Some teams use a fifth archive (some sort of database) for tracking records of the incident.

Though the aforementioned classical solution may assure continuity, it is hardly an efficient way of doing this, and it may backfire on the team in rough weather when every minute counts. Therefore the ultimate goal should ideally be to have no more than *one* archive, at least one archive that meets the eye. Any supporting structure should be hidden from the team member using the archive.

Eliminating paperwork is not that difficult. Scanning techniques are quite sophisticated today and relatively inexpensive. Incorporating these tools and techniques into everyday incident handling practices would be a good thing to do; recognizing, however, that it is not one with a high priority, since the vast majority of information is electronic to begin with. If documents are only maintained in electronic form, it is important to consider that legal requirements or rights are often affiliated with the “original” document or the signature of the

person sending a letter, etc. So all hard copy materials that have been transferred to the electronic archive should be maintained a for requisite length of time, as appropriate.

Perhaps the best current practice for integrating email, files, and access to the Web is using a Web browser. Converting email archives to the Web is possible most of the time (certainly in UNIX environments). Accessing files from a browser is easy. Search functions and indexing are also easily implemented. The Web solution is still one that a team will most likely have to devise and customize.

One further degree of complexity now has to be added: how to properly keep the history of an incident. In the above example one could simply write notes as they arise into some file or database and make it accessible through the Web or groupware system. But this still means that the majority of the actual management of the *flow* of events itself is left to the person on duty. This person has to do all the routine work, like checking on open incidents regularly, (more than likely a manual process, possibly helped with the use of some in-house-developed tool or script), updating information, annotating the record with new information, sites, etc. Where possible, any routine work should be carried out by the machine. There is enough good software around to undertake workflow management.

The groupware vendors (Lotus Notes, for example) are working hard on offering these kinds of solutions in a single software package. This development is of great interest to CSIRTs. But a common problem of workflow management software that is mainly developed for internal networks is a lack of security. This lack of security usually makes it unfit to use in a distributed environment. Teams might adopt secure tunnels over the Internet to undertake distributed work. Using a Web browser to access the workflow software (and other tools such as an email client) through the secure tunnel may solve the security problem in an elegant manner.

Essentially, workflow management software uses an underlying database in an intelligent way to keep track of changes occurring in the database (or changes *not* occurring!).

Example: A new incident is identified. It is assigned some type of tracking ID and is stored in the database. From that point on, all related events associated with it are logged. Every incident has a status field ranging from “new” to “closed.” Lack of status change may trigger alarms. (Note that this is just the core functionality; many additional possibilities and other associated actions would be tracked and recorded throughout the life cycle of the incident.)

However, integration between such tools and Web or groupware archives is still lacking in most cases, which is a serious problem. Full-text search engines are available, but must be used in addition to other products. There is light at the end of the tunnel; Web gateways for these tools are beginning to appear, and ideally these will enable the use of a workflow

management system through a Web browser. Groupware suites are starting to incorporate workflow management capabilities, though this is not yet mature from the CSIRT perspective.

In conclusion, workflow management is important to consider for helping to assure a CSIRT's continuity of work. Many practical solutions for pieces of the problem exist, but there is no single, comprehensive solution to date. Some tools can be excellent, but need tailoring and programming to adapt databases and workflows to local needs.

4.3.3 Out-Of-Hours Coverage

If your service specification calls for out-of-hours coverage, it should be quite clearly outlined what is expected during out-of-hours and what is not. Once that is clear, one can identify the functions that need to be available during out-of-hours, and the level of service expected. The quality parameters (such as response times) may well be different between business hours and out-of-hours. Without clear descriptions and guidance (along with appropriate policies and procedures), constituents will likely call for help even if they have minor problems. Each of these functions should then be analyzed with regard to the continuity aspect; what works trivially during the day in the office may well be a big problem in the evening at home. Any problems occurring should be eliminated.

Examples of typical out-of-hours problems are given below. Off-site coverage is handled in the next section.

4.3.3.1 Hotline Coverage

There are different choices for how to implement the coverage of a hotline.⁵³ The most important issue is to define who will answer the hotline calls during out-of-hours: a person from the team on duty, another staff member, or an answering service such as voicemail, a guard, or a call center of a telecommunication provider.

That decided, there are several possible ways to relay the calls to the person who will actually handle the call. If a team member will answer calls directly, this can be implemented using a call-relaying mechanism. Alternatively, a hotline number for out-of-hours calls can be disseminated to the constituents, pointing to an alternative contact number. Last but not least, the person on duty might physically stay in the office during their shift.

If hotline calls are relayed through other staff members or external parties, they can have a list of home telephone numbers for each team member, or a hotline number can be provided to call or page the staff member.

⁵³ Sometimes referred to as a help desk, customer service line, call desk, etc.

Depending on the choices made, there will be constraints to quality parameters (such as response times) that have to be considered. Issues such as provision of home equipment vs. time to travel to the office to respond to a call will also need to be considered.

4.3.3.2 Escalation

If things go awry in the daytime, escalation is usually as easy as asking other team members to help; but what happens out-of-hours? Thought should be given to this issue. For most teams, it might be a good approach to consider having at least one other team member available as a backup on short notice. Alternatively, a backup might be chosen in a crisis situation by finding out who is available. Since a crisis situation will affect team operations, the position of a “manager-on-duty” who decides and addresses conflicts might be appropriate.

4.3.3.3 How to Reach Other Teams or Customers

Your team is not the only one undertaking out-of-hours coverage. Evaluate your existing working relationships with other teams, customers, and others based on their availability outside business hours and build on these relationships, as they might be willing to provide you with other emergency numbers that they would not normally disclose. Note the time-zone problem: what is out-of-hours here may be business hours elsewhere and vice versa. National holidays differ across the world. Even the observance of public holidays may differ within a single country.

Example: The time-zone problem can be an advantage too. Cases have been reported where U.S., European, and Australian teams have used their geographical separation, covering many time zones, to enable continuous work on a problem (like an incident or vulnerability analysis and resolution). As the business day of one team came to an end, it would hand off the problem to another team whose business day was just beginning, and so on.

Example: Independence Day (celebrating the independence of the U.S. from the U.K.) is traditionally observed in the U.S. on July 4 of each year. If this holiday falls on a weekend (Saturday or Sunday), some companies in the U.S. may choose to observe it on the Friday before or the Monday after. Clearly this holiday is not observed in other parts of the world.

Example: The U.S. Veteran’s Day holiday is traditionally observed by only U.S. government (local, state and federal) and military agencies. Banks, other businesses, and organizations in the U.S. may or may not observe this holiday.

4.3.4 Off-Site Coverage

Off-site coverage is different from out-of-hours coverage because the regular services must be provided from a remote location. Usually there have to be good reasons (such as a disaster or some other crisis situation) to continue your business-hours service, with on-duty personnel being off-site (at a conference venue, at a constituent's site, or even a backup facility). This results in most of the same problems as out-of-hours coverage and more, because the level of service expected will be the same as that provided in business hours from your normal base of operations. The constituents need not—
—
aware of your specific situation. The focus should be addressing their problems and not concerning them with the steps that you have to take to provide them with service. Obviously, due to the complications that it presents, off-site coverage should be reduced to an absolute minimum.

The location (e.g., their homes during out-of-hours, or hotel room at a conference location) from which people on duty work is not necessarily known in advance. This poses extra security problems that usually have to be evaluated in a very short period of time. Depending on the circumstances, a decision must be made either to reduce the level of security necessary to provide a specific service or to keep the high security level but prevent access to the internal CSIRT network due to lack of necessary security measures. In such cases the security of the team will outweigh all other considerations.

There is obviously a good reason for the team members involved to be off-site in the first place. They will have additional tasks to undertake (e.g., a presentation at a conference or a customer meeting) in addition to any incident handling work they are requested to conduct off-site. The priorities associated with the tasks must be clear and determined in advance. These priorities determine which tasks take precedence and which can be left until the next day, when they return to the office, or until another person is available.

4.4 Security Management

A CSIRT must clearly place great emphasis on guarding its own security, but to cover all relevant aspects is beyond the scope of this document.

However, the specific CSIRT issues addressed in this section lead to the need for additional comments. The following factors (which are generic for the majority of installations) must be taken into account when considering the goals for CSIRT security management:

- confidentiality: to get what you are allowed to get and nothing more
- availability: to get what you want when you want it
- integrity: to ensure that information stays the way it was intended to be

- authenticity: to know for sure the origin of the information
- exclusivity: to ensure that only the intended recipients can use the information
- privacy: to guarantee that the interests of persons and organizations are protected
- obligation: to guarantee that due diligence requirements are fulfilled

4.4.1.1 Use of Encryption and Digital Signature Applications

The use of encryption and digital signature applications is unavoidable for any CSIRT. They offer good possibilities for securing data on the team's computer systems and during data transfer through unsecured networks. Cryptographic methods can also ensure authenticity to protect connections (especially from outside) into the team's internal network. (See below for more considerations.) Between the team and cooperating partners, common encryption tools, such as S/MIME and PGP or GPG,⁵⁴ enable secure communication of sensitive data (such as the analysis of an incident, a new artifact, or a summary of recent trends on a routine basis). Log files related to intrusions can be encrypted and transferred using email to and from constituents to keep sensitive information about victims and the systems involved private. With regard to internal encryption, one can choose proprietary standards. Several good possibilities exist but these will not be discussed further in this handbook. When dealing with the outside world, you have to opt for (de facto) standards such as PGP and GPG. S/MIME is becoming more prevalent and may also become a de facto standard, judged by the support it receives from Microsoft, Netscape, and the rest of the user community. In addition to confidentiality, authenticity can be achieved; however, there are other issues that arise as a result of this (see key management and certification issues below).

Serviceable programs such as S/MIME and PGP/GPG have been available for years. Using these programs, the user is often confronted with the program and technology directly (including the integration with the email client), but strong measures are available. S/MIME is now integrated in popular commercial email clients (Netscape, Microsoft) as well as open source software (Mozilla). PGP Version 6.x/7.x, as well as other OpenPGP implementations like GPG, have brought more user friendliness, graphical user interfaces, and better integration with email clients.

4.4.1.2 Key Management and Certification

Use of cryptography introduces a key management and certification problem. S/MIME and PGP/GPG use asymmetric encryption (also known as private key encryption) for providing strong authentication. This avoids the weaknesses of symmetric (also known as single or secret) key encryption schemes, since the secret key must be known to all communication partners. Hence it is impossible to provide authenticity of the origin and destination (since the

⁵⁴ GPG or GnuPG is an open source implementation of PGP supporting the OpenPGP standard developed in the IETF. GPG is compatible to newer PGP releases, but does not support PGP 2.6.x versions.

key is “shared”). However, asymmetric encryption makes use of two keys (which are interrelated) for each person/role. While the public key can be distributed to everyone without compromising the authenticity, the private key must be protected like a password.

Example: In the asymmetric encryption scheme, if Moira wants to send Don an encrypted email, Moira uses Don’s *public* key to safeguard the text that she writes and transmits it to Don, who then is the only one who can decrypt it using his *private* key. In addition, Moira uses her *private* key to sign the written text she sends, so Don is able to verify the origination using her *public* key.

The key management problem touches both public and private keys. The private keys have to be stored safely; if somebody controls your private key, he can decrypt everything you can decrypt. Unlocking a private key is done using a type of password called a passphrase. This passphrase must be well protected to ensure security. Some people carry their private key on a diskette or other removable media like a USB token, though one might wonder about the security that such an approach provides, especially if the diskette is not equally well protected.

Beware, though; using strong cryptography without common sense is no cryptography at all. If you use a three-letter passphrase for unlocking your private keys, and you’re not doing this on a totally isolated system, then you break the chain of security. Therefore the availability of a *strong* program is not the only critical issue; it must also be used in the *right* way. Similar problems are related to the storage of passwords and passphrases to unlock the private keys on computer disks or within programs and scripts.

One problem with public keys is the need to check whether the public key you obtained is authentic and really belongs to the person that the key is attributed to. This is why certification authorities for S/MIME or PGP/GPG are so common and necessary today. Trusted third parties (TTPs) like Verisign Inc. will sell users a key pair (public and private key). (One has to trust that the TTP does the appropriate checks to verify identities before issuing such keys.) From a user’s perspective, it is absolutely necessary to be the only person with access to the private key. If you buy a key from a TTP, the TTP will sign your key, thus making it more trustworthy. Caution should be exercised when relying on TTPs, however, because they rely on proprietary policies relevant only for their users and customers. None of these systems currently provides a digital identity for network citizens worldwide to reliably compare personal ID systems (as with, for example, visual inspection of a passport⁵⁵).

Where users can sign or accept each other’s key themselves, the same problem arises: how to check the authenticity of keys. If no direct relationship to a person exists that can be used to

⁵⁵ Although, of course, an argument can be made that such comparisons are also not 100 percent foolproof, as crafty criminals have certainly, in the past, falsified passport credentials.

verify the fingerprint, users have to rely on a “web-of-trust,” which means another user has certified that he verified the binding between a key and its user.

Example: Moira signs Don’s public key; her colleague, Ann, who does not know Don, but wants to send secure email to him, obtains Don’s public key and sees Moira’s signature on his public key. As Ann had (on some previous occasion) verified that Moira’s public key is really *her* key (after some personal meeting where both exchanged the fingerprints of their keys), Ann can compare Moira’s fingerprint/signature on Don’s key. Ann might then also trust Don’s key without needing to personally check with Don. This is the web-of-trust concept of user certification of PGP, which builds and expands the trust relations within smaller user communities. Stated differently: someone you know and believe has done the right thing to verify a key will be the link to someone you don’t personally know.

If and when a CSIRT should sign keys (either with their team key or with the key of an individual team member) is a question to be addressed. One might argue that if a CSIRT has signed the key of somebody who then proves to be untrustworthy, this action reflects poorly on the CSIRT itself. Although technically speaking this is not true, a CSIRT will try to avoid any problem and might choose not to sign any key with their team key.

Example: The CERT/CC has chosen not to sign any keys from the outside world with the CERT/CC team PGP key.

Example: CERT-NL uses only its master certification key to sign “very trustworthy” people. CERT-NL members have a less restrictive policy on what they can sign using their personal keys.

4.4.1.3 Firewalls and Network Security

Ideally the team’s network is separated from the outside world by a well-designed firewall. The outside world *includes* the team’s host organization [Chapman 1995]. Firewalls are not the ultimate solution and must be supplemented by appropriate authentication and authorization throughout the network. To recognize attacks and possible breaches of security, adequate administration and control must be ensured. Firewalls are useless if, for example, log files are not regularly checked for suspicious activities (at least daily). Tools such as Swatch⁵⁶ and logsurfer⁵⁷ allow for online recognition of suspicious log-file entries and help with analyzing and data reduction options.

Consider redundancy issues when building the local network. Critical components are not only the firewall and related hosts, but also servers (shadow file servers, shadow disks,

⁵⁶ <http://swatch.sourceforge.net/>

⁵⁷ <ftp://ftp.cert.dfn.de/pub/tools/audit/logsurfer>

surplus workstations, and hot spares). To protect against interrupted power supplies, backup power arrangements should be made.⁵⁸

Do not forget to have extra backup tapes in another location; think of fire or other disaster situations, for instance. But also recognize that the other place may be less secure than your own, so a good practice is to encrypt the backups. Encryption might be also an option by providing specific services like a file server. Consider the use of cryptographic systems like Kerberos or something based on it (like AFS) on your local network, or at least use file encryption for sensitive data. This will give additional protection if the firewall cannot block each attack.

4.4.1.4 Isolated Networks for Tests

Any testbed (e.g., for viruses, artifacts, programs with unknown behavior) must be separated from the operational or live production network of a CSIRT. This ensures that the availability and integrity of the mission-critical computer, communication, and network systems will not be affected by any testing, analysis of malicious code, contamination, or denial-of-service events that could affect the ability of the team to perform its function and that could ruin the public standing of the team. This is even more true if, for example, a virus escapes or an attack involves other systems of the Internet.

Example: In 1998, due to a flaw in the INND news daemon software, unintended break-outs of USENET news “control messages,” created for testing purposes by a CSIRT, caused thousands of /etc/passwd files to be sent from vulnerable news servers all over the world. It was very embarrassing for the CSIRT and could have been prevented had the team members performing the tests used an isolated testbed or ensured that they had secured their testbed properly.⁵⁹

4.4.1.5 Providing Off-Site Access to Local Facilities

When working at home or on the road (using a mobile computer), special care must be taken to offer secure access to local systems holding email and workflow management tools. The firewall design should not be punctured to allow for this kind of access. Thinking about this paradox of security versus outside access, one soon arrives at essentially only two possible

⁵⁸ Uninterruptible power supplies (UPS).

⁵⁹ Another subsequent effect of this testing was that the CSIRT now had all these password files to cope with. What should they do? Contact the sites to tell them that the test script went awry and the sites' INND software was vulnerable? Should they send the password files back to the sites? If so, certainly the CSIRT would not forward these files without sanitizing them first (to remove the password field content), so then the problem would be how to do this mass sanitization for all the files and get them sent to the appropriate contacts. That in itself also poses a problem. Does the team have all the (correct) contact information? What effort would be needed to look up contact information? Should the team even undertake this task? Are the sites within or external to their constituency? Is follow-up contact with external sites part of their mission and goals? As can be seen, there are lots of issues to address in this circumstance.

solutions. The first one is dialing into the network. This is in itself relatively secure, especially when the access procedure uses strong authentication such as one-time passwords or challenge-response cards. Other protection schemes rely on dialing back to a predetermined number. Even then, strong authentication is mandatory. Tapping remains possible; however, this too can be secured using encrypting devices or secure telephones and equipment (in the U.S., for example, using devices such as aSTU III or STE). Physical protection against loss or theft of a mobile (laptop) computer and the data stored on it or associated media is another security issue with which to be concerned.

The second solution is using public networks, probably the Internet. The only right way to do this is using end-to-end encryption to build a tunnel with strong authentication and encryption—a virtual private network (VPN). The neatest solution is application-level encryption, but this is often not feasible or not good enough.

As an alternative a tunnel can be built from a laptop (or other device) to the team's network on the network level. Products like SSH (Secure Shell) are built for this purpose. Experience teaches however that all these tools should be implemented very carefully and thoughtfully, otherwise the cure can be worse than the disease. As many tools are relatively new, efforts should be made to ensure adequate testing and protection; in addition, some of these new tools can contain software vulnerabilities.⁶⁰

But just protecting the communication link between a home system and/or notebook and the team's network is not enough, as the security of the systems involved might also affect the network directly ("escape" of a computer virus into the team's network) or indirectly (sensitive data is copied from a home system without notice). Therefore many of the security considerations must be applied to such systems as well. It might be easier for remote access to restrict the necessity to a minimum or disallow the handling of categories that are especially sensitive.

4.4.1.6 Physical Security

A CSIRT may not have full authority to implement all aspects of physical security itself. Physical security is usually provided by the parent organization, and must be enhanced to meet the requirements of the CSIRT if possible. Physical break-ins can be at least as damaging as intrusions over the network. Lock regimes, clear desk policy, authorization of personnel and visitor arrangements should be taken into account. In addition, consider document handling: lockers, safes, litter deposit, shredding. Do not forget to consider the physical location of faxes and printers, or even hotlines inside the "safe" environment.

⁶⁰ See, for example, the following CERT Advisories: "Trojan Horse OpenSSH Distribution," <http://www.cert.org/advisories/CA-2002-24.html> and "Multiple Vulnerabilities in SSH Implementations," <http://www.cert.org/advisories/CA-2002-36.html>.

Telephone conversations should not have the possibility of being overheard by other persons, such as guests.

Concern should be raised over wiring schemes in the building, location of hubs, etc., with regard to the possibility of eavesdropping. Distrust all other public communication mechanisms, especially mobile ones, which are susceptible to eavesdropping (although this is not a distinctive physical security problem). Consider using encryption for connections in question. Encryption can also be applied to protect file systems and backup media, and by that provide more security in case the physical security cannot be guaranteed 100 percent. Be aware that besides technical means, information “leakage” can also occur—visitors or guests might (covertly or overtly) seek information, for example, in the normal course of small talk or by just being in the room when incident-related information is discussed.

Consider janitorial or other cleaning staff, employees of the electricity company, or anyone else who might have access to your facility. Often these people are overlooked as they are low-profile and mostly invisible; however, they can completely ruin your security design. Ensure that your physical security plans take them into account.

Example: The CERT/CC offices are not accessible by cleaning staff unless a CSIRT staff member is present or the cleaning staff is escorted by a security guard.

4.4.1.7 Disaster Handling

In case of disasters, be it a highly destructive network intrusion, sabotage, fire, or other natural disaster, priority schemes and escalation procedures should be in place: what to do first (and what to neglect) and whom to warn.⁶¹ Some definition should exist of when to enter “disaster mode” (and when to return to normal operation). When disaster mode is in effect, people who do not normally belong there will tend to crowd the office. Even then, security still counts, and these disaster-induced risks should be taken into account accordingly.

When a fire is raging, the fire fighters will be everywhere, including inside your superbly secured control room or computer room. Are the consoles locked? What sensitive documents are lying around? What’s the printer printing? It’s virtually impossible to impose the strictest security even in such places, but make it a part of disaster handling to assign somebody to look after these security issues when a disaster strikes (without endangering them, of course). This should include gathering sensitive and critical information, including hardcopy documents and electronic media. It is also useful to conduct a post mortem to review and evaluate the CSIRT processes and procedures to identify what worked well and where improvements might need to be made.

⁶¹ Although disaster recovery and business continuity operations (in the parent organizational sense) are not covered in this handbook, it is important for the CSIRT to ensure that it has previously established relationships with those groups in the organization who are responsible for these roles (e.g., disaster recovery plans and/or continuity of operations plans).

If the constituency relies on the operation of the CSIRT, take precautions to provide a backup in times of crisis and disasters. After the critical event, measures must be in place to allow for a quick recovery.

4.4.1.8 Handling Internal Security Incidents

Organizations like to keep quiet about internal incidents. If nothing is articulated (written) in the security policy, nothing will be said on this topic; “keeping quiet” is the natural reaction. But it is often the wrong reaction. With the possible exception of internal attacks, incidents will have some involvement with systems outside the CSIRT’s network. As a result, other external people may be aware of the activity and may disclose the information publicly. Certainly if the perpetrator knows of the attack, and as intruders may brag and publicize their activities (supplying proof to support their claims); the incident may become public even if internally no one reports it. This is a good example for the rationale behind having in place a human error policy (see Section 4.2.6) and the supporting procedures, so that in such circumstances, the individual can (is encouraged to) report such activity. CSIRTs must “practice what they preach,” and internal attacks against the CSIRT cannot be ignored. If an incident occurs internally in the CSIRT, it should be reported, just as any other internal (to the organization) incident. CSIRTs must prepare for and address such incidents, not just for the obvious reason of containment but also because if they try to hide them and someone subsequently exposes the activity, then the team’s reputation may be irreparably damaged.

This holds true for organizations whose business it is to deal with security incidents. If you admit a problem, people will ask you “how come your security is not good enough,” and you have to explain what happened. If you hide a problem and it leaks out, you may find yourself out of business; people will not trust you any longer (because of your silence and your insecurity) and people will not trust your expertise any longer because you were not able to protect your own systems.

Clearly one should attach high priority to internal incidents, however not to the extent that other high-priority issues are ignored. A careful balance must be struck here.

4.5 Staff Issues

Regardless of the provision of appropriate documented policies and procedures, CSIRT work is essentially service based. As a result, there is an inherent reliance on competent and trustworthy staff to effectively execute a team’s policies and procedures and to exhibit diplomacy when dealing with constituents. Hence CSIRT staff play a pivotal role in ensuring the mission and service of the operation. In this section we will discuss the issues related to identifying, hiring, training, and retaining suitable CSIRT staff. We will also discuss arrival and exit procedures and extension of staff. Additionally we will discuss possible alternatives

to consider when the core CSIRT staff are insufficient either in numbers or technical skill to address situations that might arise.

4.5.1 CSIRT Staff

Many people incorrectly consider the most important attribute in CSIRT staff to be their technical experience. Although technical experience is a desirable attribute, by far a more critical criterion is an individual's willingness and ability to follow procedures and to provide a professional interface to constituents, customers, and other parties interacting with the CSIRT. It is a more desirable approach to hire individuals with less technical experience and good interpersonal and communication skills, and then train them in CSIRT-specific technical skills, than vice versa. Certainly this handbook itself provides a good start for educating and enhancing the understanding that all staff members will need in order to interact with other teams and provide a suitable service.

Having a wide range of interpersonal skills is important, because team members are constantly communicating with each other, their constituency, and other parties, such as other response teams. The reputation of a team relies on the professional interactions that its team members undertake. Interactions of a team member who is a technical expert but possesses poor communication skills may severely damage a team's reputation and standing in the community, while those interactions that are handled professionally and competently will serve to enhance the CSIRT's reputation as a valued service provider. Hence attention to an individual's interpersonal skills is extremely important.⁶²

The following interpersonal skills are important for incident handling staff and are listed here (in no specific order):

- common sense to make efficient and acceptable decisions whenever there is no clear ruling available and under stress or severe time constraints
- effective oral and written communication skills (in native language and English) to interact with constituents and other teams
- diplomacy when dealing with other parties, especially the media and constituents
- ability to follow policies and procedures
- willingness to continue education
- ability to cope with stress and work under pressure
- team player
- integrity and trustworthiness to keep a team's reputation and standing

⁶² Fithen, Katherine T. "Hiring IRT Staff Interview Process." 8th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, San Jose, California, June 1996.

- willingness to admit to one's own mistakes or knowledge limitations about a topic
- problem solving to address new situations and efficiently handle incidents
- time management, in order to concentrate on priority work

From a technical perspective, each incident handler requires a basic understanding of the underlying technology and issues on which the individual will base their expertise. The nature of these skills is similar, regardless of the underlying software and hardware technologies in use by the team or constituency.

The following technical foundation (with a few general examples in parentheses) is important for incident handling staff:

- public data networks (telephone, ISDN, X.25, PBX, ATM, frame relay)
- the Internet (aspects ranging from architecture and history to future and philosophy)
- network protocols (IP, ICMP, TCP, UDP)
- network infrastructure elements (router, DNS, mail-server)
- network applications, services and related protocols (SMTP, HTTP, HTTPS, FTP, TELNET, SSH, IMAP, POP3)
- basic security principles
- risks and threats to computers and networks
- security vulnerabilities/weaknesses and related attacks (IP spoofing, Internet sniffers, denial of service attacks, and computer viruses)
- network security issues (firewalls and virtual private networks)
- encryption technologies (TripleDES, AES, IDEA), digital signatures (RSA, DSA, DH), cryptographic hash algorithms (MD5, SHA-1)
- host system security issues, from both a user and system administration perspective (backups, patches)

It is imperative that some subset of the team has an in-depth understanding of the full spectrum of technologies and issues in use by the team and constituency. This additional level of expertise is a resource that will be used to broaden and deepen the technical resource and capability of the team and educate other team members through training and documentation. It also ensures that the team can cover smaller subsets of a constituency's technology base and can provide a full range of services. The following specialist skills to consider are in addition to an in-depth understanding of each of the technical skills listed above:

- technical skills such as programming, administration of networking components (e.g., routers, switches) and computer systems (UNIX, Linux, Windows, etc.)
- interpersonal skills such as human communications, experience in presenting at conferences, or managing a group
- work organization skills

A team may be unable for some reason to fund, find, or hire staff to provide the necessary specialist skills considered appropriate. Section 4.5.6, “Extension of Staff,” discusses possibilities for addressing such situations. Section 4.5.4, “Training Staff,” highlights other means to build upon and maintain strong skills and to support the continuous improvement to reflect changes in constituency, technology, service offerings, etc.

No single set of skills will be applicable for every position on a given team. It will be necessary to look at the constituency served and the range of technologies used to determine what skills are appropriate for the specific team’s composition. Wherever possible, individuals with a mix of skills should be hired to ensure that no single team member in the organization is indispensable. On the other hand, smaller teams should have at least one person experienced in the skills named to ensure such issues are handled in a professional way, although this can lead to other problems when such a person leaves the team. While it might seem contradictory, it is much simpler to replace even the most experienced team member than a person serving as an interface to the sponsoring/funding organization and to other teams.

4.5.2 Hiring Staff

When considering applicants for a given staff vacancy, it is important to decide in advance the hiring process that will be used to identify the most appropriate candidates. Observations from operational experience show that even a candidate who appears on the surface to have the appropriate skill set still might not be able to cope with the CSIRT working environment. In addition, when a crisis arises, some candidates may not have the ability to handle their tasks and do the job. It is better for all concerned to submit a candidate to a hiring process that is designed to identify candidate strengths and deficiencies. Armed with that information, the team can decide if they are able to train the candidate in the specific skills that the candidate may need or choose not to hire the candidate.

Every CSIRT will be bound to specific requirements based on the requirements of their parent organization, local and national laws, and culture. However, where possible and appropriate, the following steps should be included in any CSIRT hiring process:

- pre-interview document check
- pre-interview telephone screening
- interviews that cover topics from technical abilities to social skills and team fit
- candidate technical presentation
- reference checks, including criminal records, as appropriate

Depending on specific organizational needs of the parent organization or the constituency, more steps such as security clearances and/or background checks may also be necessary or required.

The overall hiring process should be designed to ensure or identify that the candidate has the suitable interpersonal skills for the position and has (or can be trained) in the necessary technical skills. As many of the team members as possible should have the chance to interact with the candidate, whether that be as an interviewer, through a lunch meeting, or as a participant at the candidate's technical presentation. Additionally it is important that during the interview process the CSIRT staff effort involved in the interview process is kept to a minimum, yet is used to the maximum effect.⁶³

The pre-interview document check and telephone screening with the candidate will help to ensure that the candidate is worth bringing in for a personal interview. This step can cover issues as wide ranging as the candidate's general level of interest in computer security to obtaining more specific detail on items covered in their resume. But most importantly, this is an opportunity to obtain a good impression of the candidate's oral communication skills.

To make the best use of the CSIRT staff interviewing candidates, it is worthwhile deciding in advance what particular issues (ranging from technical issues and ethical issues to social skills) you would like to gain through the interview process and which existing staff are most suited to cover those issues with the candidate. Each of the various interviewers can then cover specific topic areas and save any duplication of effort. Interviewer feedback on the issues covered can then be consolidated and discussed by the team members. In some cases, another approach may be taken: similar topics may be covered by other team members involved in the interview process—but from different perspectives—to determine the candidate's depth or breadth of knowledge about a particular topic and/or point out any weaknesses or inaccuracies in the candidate's understanding.

⁶³ Michele Crabb, "How To Find and Hire Good Technical People," *Proceedings of SANS 1996 Conference*, Washington, D.C., May 12-18, 1996; Katherine T. Fithen, "Hiring IRT Staff Interview Process," *8th Workshop on Computer Security Incident Handling*, Forum of Incident Response and Security Teams, San Jose, Calif., July 1996.

The requirement to have a candidate give a technical presentation provides the CSIRT with an opportunity to gauge other technical and interpersonal qualities of the candidate. The team can understand how much common sense the candidate has and how the candidate copes under somewhat stressful situations. They can quantify other attributes such as general presentation skills, attention to detail, technical accuracy, and ability to answer questions on the fly.

Once an individual is hired, of course, there is much more work to be done to integrate them into the CSIRT. The hiring part is just the beginning; the new staff member will need to undergo training for some period of time (see 4.5.4, “Training Staff”). Some new staff members may be given access only to limited information until appropriate certifications or relevant clearances are obtained (government or military clearances, for example). All new staff will also, of course, be required to go through some type of training period to acclimate them to the CSIRT as well as other specific policies and procedures for the team and/or organization.

4.5.3 Arrival and Exit Procedures

Due to the sensitive nature of the information handled by a CSIRT, it is important that special procedures are in place to handle the arrival of new staff as well as the departure of staff from the team. New staff members might be expected to sign CSIRT-specific agreements in addition to any standard employee agreements (such as non-disclosures or intellectual property rights) required by the parent organization. The CSIRT-specific agreements might include issues ranging from information disclosure to network connectivity and media interactions.

Prior to the departure of a member of the CSIRT (even if they are simply moving out of the team but staying within the same parent organization), exit procedures should be followed and would involve actions to be taken by other appropriate CSIRT members (such as a team’s system administrators). Exit procedures might include

- change of passwords (both personal and system passwords)
- return of any physical security devices and other media (telephone, pagers, backups)
- revocation of keys (both physical and digital)
- debriefing to review her/his past experiences and to collect ideas for improvements
- exit interview to remind the departing person of responsibilities, which may include additional agreement signing
- an announcement to the constituency and other parties with which the CSIRT regularly interacts
- action to be taken with future correspondence (email, postal) addressed to the individual

If a person leaves the team of their own will, it is worthwhile to understand the reason for their decision to leave. This might enable the team to recognize circumstances that need further attention to avoid similar departure by other team members.

Example: Due to long periods without job rotation, the person left the team as another organization offered a much more interesting job in the area of multimedia security.

If a team member is fired (terminated), different exit procedures might apply, since there are underlying reasons for the decision that affect the trust placed in the employee. This procedure might include escorting the person to remove their personal effects (under the watchful eye of the CSIRT management), an exit interview with the human resources department in the organization, and escorting/moving the terminated employee off the premises (this might be handled by the human resources manager, a site security officer, or other security official). Changes of passwords and other measures certainly still apply in such cases. In particular, if specific items like keys or hardware tokens cannot be collected, appropriate contingency plans must be implemented.

4.5.4 Training Staff

Staff training is necessary from three perspectives: bringing new staff members up to the necessary skill level to undertake their work; broadening the abilities of existing staff members for personal development and overall team benefit; and keeping the overall CSIRT skill set up-to-date with emerging technologies and intruder trends.

When looking at the overall training needs of a team, it is important to identify the overall skills needed for each team member, as well as the general skill coverage required for the team as a whole. New staff members should be trained immediately in any mandatory skills required to make them effective as soon as possible. From a broader perspective, the team should be evaluated as a whole to identify training that will expand or increase coverage of skill sets in the team and that at the same time addresses a given individual's skill set. Policies and procedures should be in place to cover at least initial training and to ensure ongoing training as policies and procedures change. Sometimes a refresher course is necessary to maintain a steady awareness as to why it is important to follow the established policies and procedures, as well as to exercise situations in which personnel must apply their own common sense if a gap in the policies and/or procedures is identified.

In addition to the interpersonal and technical skills discussed earlier in this section, it will be important for every member of the team to be trained in areas specific to the incident handling functions and the local team environment. Training should include coverage of the following issues:

- new technical developments
- local team policies and procedures
- understanding and identifying intruder techniques
- communicating with sites
- incident analysis
- maintenance of incident records
- team building
- work load distribution and organizational techniques

Initial training is strongly related to on-the-job-training and deserves further discussion. Initial training in many professions is generally in the form of background reading, observation, and then learning by experience. This holds true for incident handling, but there is no formal educational path for CSIRT staff and limited documentation in the literature, and most written material comes in the form of workshop reports or presentation slides. Since the written material through which people can learn to handle incidents is limited, on-the-job-training becomes a necessity.

The CSIRT Development Team is working to address this gap in available CSIRT training and has developed and expanded a suite of training courses and other supporting materials that can help train new (or existing) incident handling staff. These training materials target managers and technical personnel in areas such as creating and managing CSIRTs, responding to and analyzing security incidents, and improving network security [CERT/CC 2002a].

It is also important for new CSIRT staff to review and study internal documents, such as policies and procedures, case studies, or past incident summaries that have been archived by the team.

Even the most experienced staff members feel some level of stress when dealing with sensitive information. Some of that stress results from their understanding of the magnitude of the consequences if they handle the information inappropriately. New staff can be overwhelmed with the sheer volume of information, policies, and procedures that they encounter in a CSIRT. As a rule, it is inappropriate to submit such new staff to tasks where they might inadvertently disclose sensitive information without some initial training. Try to ensure that the trainee can learn the profession without making costly mistakes. A commonly used approach is one where existing CSIRT staff mentor new staff in the teams' policies and procedures through on-the-job training. A new staff member might gain proficiency in the areas of triage and request handling before moving on to small-scale incidents. In each area, the approach could take the form of the new staff member first observing the actions of an experienced staff member and undertaking follow-up discussion to address any areas of

confusion. Then as they become more familiar with the environment, the new staff member drafts email for review and edit by an experienced team member. In this way, progression can be made until the new staff member is suitably proficient and considered able to handle such tasks without assistance.

Other approaches prior to dealing with real-life incidents, such as role-playing games, might be appropriate and show the new member how policies and procedures affect the handling process [Smith 1994].⁶⁴

On-the-job training can also be used for existing team members who need to be trained to maintain their knowledge base. This is vital for the team, since the technical world is changing rapidly. In addition, attending conferences, technical exchanges, “brown-bag lunches,” and/or work in appropriate international task forces and working groups provides knowledge not only to the team member involved, but to the team as a whole, when such knowledge is shared across the team.

4.5.5 Retaining Staff

As discussed in the introduction of this document, experienced CSIRT staff are in short supply and expensive to hire and train for your CSIRT environment. So having invested in the time and resources to identify, hire, and train staff, it is most important to try to retain them. The two main reasons for turnover of CSIRT staff are burnout and low salary.

Many CSIRT staff suffer from burnout (the authors of this handbook are not exceptions), where the constant pressures and stress from daily (and often nightly, if a 24-hour service is offered) incident handling tasks become a burden and intrude into the private life. Staff can become bored with routine incidents, are physically tired, lack attention to detail, and make costly mistakes. Large salaries are now becoming available in the incident response world, mostly by way of fee-for-service CSIRTs. But not all teams, especially in the research and education community, will have the budget to pay a competitive salary. On the other hand these teams do not necessarily provide 24-hour coverage. The pull of large salaries will inevitably be enough to immediately draw certain people, but for others, incentives such as job satisfaction and personal growth possibilities will encourage them to stay. The following approaches should be considered to address both of these issues:

- rotation of duties related to routine work and incident handling
- no more than 80 percent of any individual’s effort dedicated to incident handling service

⁶⁴ Longstaff, Thomas A. “Incident Role Playing: An Exercise to Develop New Insights Into the Process of Investigating a Computer Security Incident.” 5th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, St. Louis, Mo., August 1993.

- attendance at technical conferences/workshops/tutorials (such as the FIRST Conference)⁶⁵ or other security venues that are applicable (e.g., training courses)
- participation at technical working groups (like the IETF)
- development of in-house training courses
- attendance at in-house training courses

Teams that have the greatest success in retaining quality staff have strong team environments where staff mix socially as well as in the work environment. They are also organizations in which the contributions of all team members (technical and non-technical, new and experienced) are considered and valued.

4.5.6 Extension of Staff

A team may be unable (for some reason) to find, fund, train, or hire appropriate staff to provide the necessary specialist skills required by the team. In such cases, the team can consider developing relationships (and clear agreements of understanding) with experts in the field to provide the necessary skills. When a situation arises where in-house expertise is insufficient, these experts can be called upon to fill the void. Because workload in the CSIRT environment is unpredictable and often driven by external events that are not under the control of the CSIRT, there are times when existing CSIRT staff will be insufficient to cope with the level of demand for its services. It may be appropriate for the CSIRT to have procedures in place to reach out for assistance to individuals previously identified as backup or extensions to the core CSIRT staff. This will enable the team to cope when the incident load peaks above given thresholds, or in other circumstances defined in the team's escalation policies and procedures. These additional staffing resources might be drawn from

- other areas of the security teams in the CSIRT parent organization
- other groups in the CSIRT parent organization
- other groups in the CSIRT's constituency
- other CSIRT organizations
- external trusted experts and service providers

When considering staff to serve in this role, the same hiring principles should apply for them as for any CSIRT member. Also, the following processes should be established in advance so extension staff can be activated as quickly as possible:

- agreed-on criteria for calling in extension staff participation
- non-disclosure agreements, service level agreements, memoranda of understanding, etc.
- up-to-date contact information

⁶⁵ <http://www.first.org/conference/>

- prior agreements from management
- procedures to establish secure communications
- initial and regular training

It is essential to provide extension staff on-the-job-training before they are allowed to participate in the actual incident handling process. This will give all personnel the chance to socialize with each other and become familiar with the way policies and procedures are executed through the day.

5 Closing Remarks

5.1 Closing Remarks from the First Edition

Writing this document took much longer than expected and required a considerable amount of effort. It wasn't always easy to decide when to provide more detail and when not to. What started out to be a short report soon took on a life of its own. Finally we decided that a handbook would be a more appropriate term for this document.

One issue that we struggled with continually was how useful the information would be to someone implementing a CSIRT for their own environment and how to provide information that would still be applicable a year or more from now. As is true for security in general, the needs of each CSIRT are unique and the CSIRT environment is dynamic. There is no chance of long-term stability, since technology, the constituency base, and the intruder community can change any time. To ensure successful operation, a CSIRT must have the ability to adapt to the changing needs of the environment and exhibit the flexibility to deal with the unexpected. In addition, a CSIRT must simultaneously address funding issues and organizational changes that can affect its ability to either adapt to the needs or provide the service itself.

Throughout the years that we have worked in and influenced the area of CSIRTs and incident handling, we have found it rewarding work (despite being hard, sometimes frustrating and demanding work). The rewards come from believing in the work, getting the chance to interact with other dedicated members of teams from around the world, experiencing the willingness of the CSIRT community to share lessons learned and support each other, and having a general interest in the work and the underlying technology that supports it.

One of the main motivations for writing this document is to help others. Collectively, we have helped many teams across the world to form, and we learned a lot, made new friends, and had fun in the process. But we wanted to document as much of the information that we'd learned as possible so that others can benefit from it. We hope that we have succeeded in not only documenting the information, but also providing it in a form that is both meaningful and useful for others. The area of CSIRT development, practices, and processes is still in its infancy, and it is still struggling to find its place within the computer security realm—which in turn is finding its niche in the computing arena. We hope that this document will be seen as a major contribution to the continued development and maturity of CSIRTs and this field. If not, we hope that it can at least be a starting point for further discussion, refinements,

improvements, and initiatives to develop better documents, policies, and perhaps standards. We will be happy to be involved with or contribute to other efforts of this nature.

5.2 Closing Remarks for the Second Edition

Near the latter half of 2002, the CERT CSIRT Development Team decided that it was time to take another look at the content in the CSIRT Handbook with a view towards publishing a second version that incorporated updated information and lessons learned since the first version was published in 1998.

The updates were undertaken by the CSIRT Development Team (CDT), which is part of the Networked Systems Survivability Program (where the CERT/CC is also housed) at the Software Engineering Institute in Pittsburgh, PA. Since the first hotline call that was handled in 1988, CERT staff have been handling incident reports, learning, and transitioning knowledge to others. Early in CERT/CC history, staff developed training materials to assist the U.S. Army's Land Information Warfare Activity in building their incident handling capability. From that time, training has been an important component of the activities in the CERT/CC. In 1996, the idea of a CSIRT Development Team germinated and grew from within the incident handling team of the CERT/CC, and this team was formalized as a separate component. Our CDT mission is focused primarily on the development of CSIRTs world-wide. We achieve this through training courses, direct customer relationships, and developing products that provide awareness, education, training, and knowledge in the area of computer security incident response. The CDT has trained hundreds of technical and management staff to provide them with the knowledge and skills they will need to build their own CSIRT.

Being a part of a world-class organization such as the CERT/CC has taught us new things, stretched our minds, and broadened our vision to reach out to others—and to return to the community the lessons that have collectively been experienced after nearly 15 years of operation. It has provided each of us with an opportunity to meet new friends, share experiences, work with other leaders in the CSIRT community, and interact with other dedicated members of CSIRTs throughout the world.

There is a growing consensus that having a CSIRT with the ability to detect attacks and incidents, as well as initiating appropriate responses, is expected in many of today's organizations. In fact, for some sectors there are mandates or regulations that require an organization to have an incident response capability in place. In this regard CSIRTs that are part of an enterprise are seen as a partner that is integrated with their risk and security management infrastructure. CSIRTs that are supporting large constituencies of independent organizations are seen as support centers that provide a variety of benefits to the communities being served—through support for incident handling (response, notification, education, training, awareness) and as dependable links between and across organizations affected by

threats, attacks, and computer security incidents. In addition, the protection of critical infrastructures is becoming more important today than ever before, and organizations that have robust CSIRTs in place with well-defined incident handling processes, as well as related management processes, will be able to quickly and effectively respond to threats and attacks when they occur.

Other things have changed over the years as well. In 1998, the original authors stated that “the area of CSIRTs is still in its infancy;” in 2003, we recognize that, in many ways, CSIRTs are still the pioneers in developing and promoting a number of incident handling and management practices.

CSIRTs can serve an important role in supporting best practices for management of risk and security issues in the technologies that our society relies on today. By incorporating feedback and lessons learned—from knowledge gained in responding to incidents, vulnerabilities, and attacks—CSIRTs can help improve security quality management processes for the long-term security efforts across the enterprise. In our collective view, it would be fair to say that the CSIRT has an established role and place within the computer security realm. New interest in the work of such teams results from the need to analyze what actually happened during an attack or incident, understand the threats and risks, implement effective responses, and prevent future attacks. Computer forensics is an area that is also becoming more important in the CSIRT environment and could be seen as a natural progression for additional services that CSIRTs might provide to their constituency (and in fact are already being offered by some teams).

Reading through the original handbook and working on the revisions has shown us that we remain as passionate about the work, find it rewarding, and want to help others. At the same time, we see that the same demands and frustrations that Moira, Don, and Peter experienced when they wrote the first handbook still exist today. These past four years have reaffirmed for the CSIRT Development Team that the rewards come from believing in the work and seeing new teams forming around the globe.

In closing, the CSIRT Development Team would like to paraphrase the original authors’ words from 1998: We hope that this document will be seen as a major contribution to the continued development and maturity of CSIRTs. The original CSIRT Handbook has been widely used as a starting point for new teams, as an evaluating and benchmarking tool by existing teams, and as a mechanism to suggest further refinements, improvements, and initiatives to develop even better policies, processes, procedures, and standards of practice in this area.

Finally, we believe strongly that this new release, together with other documents the CSIRT Development Team will publish,⁶⁶ will further add to the body of knowledge and support in this important area of activity. Each CSIRT environment is a little different from the next and each has something to add to the overall knowledge we seek. The importance of sharing lessons learned, the exchange of ideas, experiences, and processes with other new or existing teams can have a tremendous effect on “raising the bar” on CSIRT activities. We have knowledge about a large number of teams that exist (through our own interactions and discussions with individuals in FIRST and TERENA), but we suspect that there are many additional teams being created that we have not yet heard about. If you have established a CSIRT, we’d like to hear from you.

We welcome your comments on this handbook. If you want to share your opinions or have suggested additions (or other comments) to this handbook, please contact us. We regularly attend FIRST conferences, and we can be contacted in person or reached as a group by sending email to csirt-handbook@cert.org.

⁶⁶ For example, the *State of the Practice of CSIRTs*, *Organizational Models for CSIRTs*, and other CSIRT-related publications that will appear on our Web site at <http://www.cert.org/csirts/>.

Appendix A: About the Authors

The authors of this handbook have extensive experience in the formation, documentation, and operation of their own teams' incident handling services and in assisting many different computer security incident response teams (CSIRTs) around the world from their inception through formation and operation. The authors are leading figures in the CSIRT community. They are frequently invited to give presentations on a wide range of Internet security topics, from critical infrastructure issues to social impact.

Moira J. West-Brown <*moira@west-brown.com*>

Moira J. West-Brown was a senior member of the technical staff within the CERT Coordination Center (CERT/CC) based at the Software Engineering Institute (SEI) until 1999. Prior to leaving the SEI, West-Brown led the group responsible for facilitating and assisting the formation of new CSIRTs around the globe. This group assisted a wide range of different organizations in the formation of national, government, Internet service provider, and academic CSIRTs.

West-Brown joined the CERT/CC in 1991 as a technical coordinator, responding to computer security incidents and vulnerability reports. For several years she managed the CERT Operations team, which focuses on reactive tasks aimed at responding to computer security attacks and vulnerabilities. She successfully led the team through a period of dramatic increase in the rate of intruder reports, and she established and developed many of the operational standards adopted for use by other CSIRTs today.

Prior to her tenure in the CERT/CC, West-Brown had extensive experience in system administration, software development, and user support/liaison, which was gained at a variety of companies ranging from academic institutions and industrial software consultancies to government-funded research programs.

West-Brown has been an active figure in the international CSIRT community. She developed a variety of tutorial and workshop materials focusing mainly on operational and collaborative CSIRT issues. She was elected to the FIRST Steering Committee in 1995 and served as the Steering Committee Chair from 1997 to 1999.

West-Brown holds a first-class bachelor's degree in Computational Science from the University of Hull, U.K.

Klaus-Peter Kossakowski <kpk@presecure.de>

Klaus-Peter Kossakowski is managing director of an independent German company providing senior consultancy on International Information Infrastructure, IT security and Incident Response since early 2000. Kossakowski's work currently involves risk and security management, incident response services, public key infrastructures, intrusion detection, network security, forensics, and security improvement.

Kossakowski has worked in the security field for more than 15 years. In 1988 he was one of the first members of the Virus Test Center in Hamburg, where he focused on malicious network programs. He was involved with DFN-CERT (the first German CSIRT for an open network) from its inception. From January 1993 until he left DFN-CERT at the end of 1997, he managed the DFN-CERT team, which was modeled after the CERT/CC. He successfully led the team from a research effort to a functional and well-respected entity in the CSIRT community. From 1998 to 1999 he was a senior consultant and project manager at secunet Security Networks AG, a German IT security provider, where he founded the internal secu-CERT team. He was also a visiting scientist at the CERT/CC from 1998 to 2003.

Kossakowski continues his collaborative efforts with the CSIRT Development Team to develop technical articles, documents, training materials, and other CSIRT-related publications.

Kossakowski's particular interests in the CSIRT arena are international issues, cooperation, and establishing a CSIRT infrastructure. As the co-chair of the IETF working group "Guidelines and Recommendations for Incident Processing" (GRIP), he has been involved with the development of several RFCs since 1994. Together with Don Stikvoort, he initiated a closer cooperation among European CSIRTs and organized several annual meetings to support these. His vocal role in the European CSIRT community resulted in him becoming chair for a TERENA task force called "CERTs in Europe." This task force outlined the concept and service definition of a European CSIRT Coordination Center. As a result of this effort, EuroCERT was implemented in late 1996. He was elected as a member of the Forum of Incident Response and Security Teams (FIRST) Steering Committee in 1997 and re-elected in 1999 and 2001. In this role he actively supports international CSIRT cooperation and the move of FIRST toward a new organizational structure.

Kossakowski has defended his doctoral thesis in "Information Technology—Incident Response Capabilities" at the University of Hamburg. He holds a first-class degree in Information Science from the University of Hamburg. Kossakowski is a member of the Internet Society (ISOC), the Information Systems Security Association (ISSA), and the German "Gesellschaft fuer Informatik e. V." (GI).

Don Stikvoort <don@elsinore.nl>

Don Stikvoort is managing director and co-founder of the Dutch companies STELVIO and S-CURE, offering senior consultancy services in the areas of Internet and intranet security.

He has worked in the security area for more than 15 years. After his academic years he embarked on his working life as Infantry platoon commander in the Dutch Army. He joined SURFnet, the Dutch national research and educational network, in 1989. During his 9 years at SURFnet, Stikvoort had a variety of responsibilities. He started out as consultant but soon became responsible for setting up the SURFnet backbone. Later he managed subcontractors responsible for the SURFnet Helpdesk and other user-oriented services, and led several development projects. He was involved in the formation of CERT-NL in 1991 and was its chairman from 1992 to 1998.

Stikvoort is an active participant internationally in RIPE, TERENA, IETF, and the FIRST community, particularly in regard to security issues. Together with Klaus-Peter Kossakowski, he initiated the closer cooperation of European CSIRTs in 1993 and has contributed since that time to efforts leading to a more structured European incident coordination. From 1996 to 1998 he was actively involved in helping FIRST to evolve its organizational structure and funding model. In 1998 he finished the first version of the *Handbook for Computer Security Incident Response Teams (CSIRTs)* together with Klaus-Peter Kossakowski and Moira J. West-Brown. Stikvoort was chairman of the Program Committee for the 1999 FIRST conference in Brisbane, Australia.

Stikvoort holds a degree in experimental low temperature physics from Leiden University, The Netherlands. He is a member of ISOC and participates in several national security fora. He is end-responsible for the operation of the international FIRST secretariat and the European Trusted Introducer service (which offers an independent accreditation process for European CSIRTs), both of which are subcontracted to S-CURE. Recently Stikvoort has been engaged in consultancy on behalf of kennisnet, the national Dutch school network, and on behalf of the Dutch Government GOVCERT.NL . He currently participates in the set-up and execution of several CSIRT courses, and is the coordinator of eCSIRT.net, an EU-funded research project that aims at implementing the IODEF standard for exchange of incident data and statistics between CSIRTs.

Georgia Killcrece <georgia@cert.org>

Georgia Killcrece is a member of the technical staff in the CERT® CSIRT Development Team, part of the Networked Systems Survivability (NSS) Program based at the Software Engineering Institute (SEI) at Carnegie Mellon University in Pittsburgh, Pennsylvania. Killcrece joined the CERT Coordination Center (CERT/CC) in September of 1989, just 10 months after the CERT/CC was formally established.

Killcrece has worked directly with a variety of organizations (government, industry, and academia) across various stages of their CSIRT development: education/training, planning, implementation, operation, collaboration. She is also involved in the development and delivery of public and onsite training courses, and is one of a few instructors at the Software Engineering Institute who regularly teaches the suite of five CSIRT courses. At the end of 2002, Killcrece and her team completed a program to license the CSIRT training materials. Through transition partners licensed to deliver the CSIRT courses, CSIRT training can now be disseminated to a much broader global community. Jointly, the team has published documents currently available on the CERT Web site (<http://www.cert.org/csirts/>), to extend the availability of information to organizations establishing their own CSIRT.

Prior to becoming leader of the CERT CSIRT Development Team, Killcrece was a technical coordinator and incident response coordinator in the CERT Coordination Center from 1994 to 1999. In those roles, she gained first-hand knowledge of the processes involved in forming, operating, and evolving incident response teams, including the dynamics of working in a fast-paced team environment. Her primary responsibilities included handling computer security compromises and vulnerability reports. As a part-time member of the Information Services team in the Networked Systems Survivability program, Killcrece also contributed to the development and maintenance of CERT/CC materials published on the CERT Web site.

Killcrece is a co-author or contributor to papers and reports relating to computer security incident response teams, including *Creating a Computer Security Incident Response Team: A Process for Getting Started*, *CSIRT Services*, *CSIRT Frequently Asked Questions*, and this revised edition of the *Handbook for Computer Security Incident Response Teams (CSIRTs)*. Other publications include *Organizational Models for CSIRTs* and *State of the Practice of Computer Security Incident Response Teams*, scheduled for release in 2003.

She can be reached directly by email at georgia@cert.org or via the CSIRT training alias at csirt-info@cert.org.

Robin Ruefle <rmr@cert.org>

Robin Ruefle is a member of the technical staff in the CERT CSIRT Development Team, part of the Practices, Development, and Training group in the Networked Systems Survivability program at the SEI.

Ruefle's focus is on the development of management, procedural, and technical guidelines and practices for the establishment, maturation, and operation of CSIRTs worldwide. Ruefle develops and delivers sessions in the suite of courses offered to CSIRT managers and incident handling staff, including *Creating a CSIRT*, *Managing CSIRTs*, *Fundamentals of Incident Handling*, and *Advanced Incident Handling for Technical Staff*. She also participates in the Train-the-Trainer program that licenses these products to existing CSIRTs. As a member of the CERT CSIRT Development Team she provides guidance in the development of implementation strategies, policies, standard operating procedures, response plans, and training programs for new and existing CSIRTs. Ruefle is a co-author or contributor to papers and reports relating to computer security incident response teams, including *Creating a Computer Security Incident Response Team: A Process for Getting Started*, *CSIRT Services*, and *CSIRT Frequently Asked Questions*.

Prior to coming to the SEI, Ruefle worked as a consultant and trainer in the Academic Computing department at the University of Pittsburgh. During her time with Academic Computing she managed the academic computing public labs, training program, and consulting office. She also served as project leader for the wireless laptop pilot program and Microsoft Windows NT conversion project. Before that, she developed database programs, training courses, and documentation for the Central Services Comptroller's office in the Office of the Budget of the Commonwealth of Pennsylvania.

Ruefle received a BS in political science and a Master of Public and International Affairs from the University of Pittsburgh. She is also an adjunct faculty member in the MBA program at Chatham College and the Graduate School of Public and International Affairs at the University of Pittsburgh, where she teaches courses in information technology, management information systems, and information retrieval and analysis.

She can be reached directly by email at rmr@cert.org or via the CSIRT training alias at csirt-info@cert.org.

Mark Zajicek <mtz@cert.org>

Mark Zajicek is a member of the technical staff at the Software Engineering Institute (SEI) at Carnegie Mellon University.

Zajicek's current work is focused on helping other organizations to build their own computer security incident response teams (CSIRTs). As a member of the CERT CSIRT Development Team <<http://www.cert.org/csirts/>>, part of the Practices, Development, and Training group in the Networked Systems Survivability program at the SEI, he is responsible for providing guidance to new and existing CSIRTs worldwide. He has co-developed a variety of documents and training materials, and is an instructor for a suite of courses that provide training for CSIRT managers and technical staff.

Previously, Zajicek was the Daily Operations team leader for the CERT Coordination Center (CERT/CC), after having joined the CERT/CC's incident handling staff in 1992. Prior to joining the CERT/CC, he was a user consultant for the Computing Facilities group at the SEI. Zajicek also helped support the CERT/CC during its initial start-up in 1988.

Zajicek holds a bachelor's degree in Electrical Engineering and Biomedical Engineering from Carnegie Mellon University

Appendix B: Glossary

This glossary lists acronyms and abbreviations that are used throughout the handbook and contains a short list of definitions of the most important terms relevant to the objectives of the handbook.

Acronyms and Abbreviations

24x7	twenty-four hours a day, seven days a week
AFS	Andrew file system
BCERT	Boeing CERT
CERT/CC	CERT Coordination Center
CERT-NL	Computer Emergency Response Team Netherlands
CIDR	Classless Inter-Domain Routing
CIRC	Computer Incident Response Capability
CIRT	Computer Incident Response Team
CSIRC	Computer Security Incident Response Capability
CSIRT	Computer Security Incident Response Team
CSRC	Computer Security Resource Center
DFN-CERT	Deutsches Forschungsnetz Computer Emergency Response Team
DNS	Domain Name System
FIRST	Forum of Incident Response and Security Teams
FTP	file transfer protocol
GPG	Gnu Privacy Guard
GRIP	“Guidelines and Recommendations for Incident Processing”
HTTP	Hypertext Transmission Protocol
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IHT	Incident Handling Team
INND	Internet news daemon

IP	Internet protocol
IRC	Incident Response Team
IRT	incident response team
ISP	Internet service provider
MCERT	Motorola Computer Emergency Response Team
MD5	Message Digest 5
MIME	Multipurpose Internet Messaging Extension
NTP	Network Time Protocol
PGP	Pretty Good Privacy
POC	point of contact
RFC	request for comments
S/MIME	Secure Multipurpose Internet Mail Exchange
SERT	Security Emergency Response Team
SIRT	Security Incident Response Team
SMTP	Simple Mail Transport Protocol
SSC	site security contact
SSH	Secure Shell
STE	Secure Terminal Equipment
STU III	Secure Telecommunication Unit III
SUNSeT	Stanford University Network Security Team
TCP	Transmission Control Protocol
TERENA	Trans-European Research and Education Networking Association
TTP	trusted third party
UBC	Unsolicited Bulk E-Mail
UCE	Unsolicited Commercial E-mail
UDP	User Datagram Protocol
UNI-CERT	Unisource Business Networks Computer Emergency Response Team
WWW	World Wide Web

Terms

Artifact

The remnants of an intruder attack or incident activity. These could be software used by intruder(s), a collection of tools, malicious code, logs, files, output from tools, status of a system after an attack or intrusion. Examples of artifacts range from Trojan-horse programs and computer viruses to programs that exploit (or check for the existence of) vulnerabilities or objects of unknown type and purpose found on a compromised host.

Authenticity

If the identity of some subject or object can be checked and verified, the relationship between the subject/object and its identity is called authentic. In the realm of computer security this is usually associated with verifying that information which is sent or received has not been altered in any way during transmission.

Bugtraq

A mailing list for the discussion of security problems and vulnerabilities. Occasionally full disclosure reports of new vulnerabilities and exploit tools are distributed through this list.

Computer Security Incident

Any real or suspected adverse event in relation to the security of computer systems or computer networks. Examples of such events are

- intrusion of computer systems via the network (often referred to as “hacking”)
- the occurrence of computer viruses
- probes for vulnerabilities via the network to a range of computer systems (often referred to as “scans”)

In the computer security arena, these events are often simply referred to as incidents.

Computer Security Incident Handling

By providing the basic set of services (triage, handling, and request), a team offers a defined constituency support for responding to computer security incidents. In addition to this basic set, an announcement service might also be offered.

CSIRT

An acronym for “computer security incident response team.” This is a team providing services to a defined constituency. There are several acronyms used to describe teams

providing similar types of services (e.g., CSIRC, CSRC, CIRC, CIRT, IHT, IRC, IRT, SERT, SIRT). We have chosen to use the generic term “CSIRT,” as it has been widely adopted in the computer security community.

Depending on factors such as expertise and resources, the level and range of service provided might be different for various teams.

Constituency

A specific group of people and/or organizations that have access to specific services offered by a CSIRT.

Intruder

A person who is the perpetrator of a computer security incident. Intruders are often referred to as “hackers” or “crackers.” While “hackers” were very technical experts in the early days of computing, this term was later used by the media to refer to people who break into other computer systems. “Crackers” is based on hackers and the fact that these people “crack” computer systems and security barriers. Most of the time “cracker” is used to refer to more notorious intruders and computer criminals. Sometimes it is argued that the term “attacker” would be better, as an unsuccessful attack doesn’t constitute an intrusion. But because of the intention of the person responsible for the attack, the term “intruder” is used throughout this document.

Liability

The responsibility of someone for damage or loss.

Policy

A set of written statements directing the operation of an organization or community in regard to specific topics such as security or dealing with the media.

Procedure

The implementation of a policy in the form of workflows, orders, or mechanisms.

Remnant Files

Files left by intruders on compromised systems. These can range from Ethernet sniffer log files, password files, exploit scripts, and source code to various programs (may also be called “artifacts”).

Security Policy

A policy addressing security issues.

Site

Depending on the context in which this term is used, it might apply to computer system(s) that are grouped together by geographical location, organizational jurisdiction, or network addresses.

Site Security Contact (SSC)

A person responsible for computer security issues at a specific site.

Social Engineering

Social engineering is the art and science of getting people to do something you want them to do that they might not do in the normal course of action.

Instead of collecting information by technical means, intruders might also apply methods of social engineering such as impersonating individuals on the telephone, or using other persuasive means (e.g., tricking, convincing, inducing, enticing, provoking) to encourage someone to disclose information. Since these are based on the social interactions and habits of people, it is called social engineering.

Triage

The process of receiving, initial sorting, and prioritizing of information to facilitate its appropriate handling.

Trojan Horse

A normally trustworthy program or process modified to include unwanted and unknown functions that may (or can) compromise the security of the user, system, network, application, or protocol involved.

Vulnerability

The existence of a software weakness, such as a design or implementation error, that can lead to an unexpected, undesirable event compromising the security of a system, network, application, or protocol.

Bibliography

- [Aslam 1995]** Aslam, Taimur. "A Taxonomy of Security Faults in the UNIX Operating System." Master's Thesis, Purdue University, 1995.
- [Brand 1990]** Brand, Russell L. "Coping With the Threat of Computer Security Incidents: A Primer from Prevention Through Recovery." Version CERT 0.6. Pittsburgh, Pa., June 1990.
- [CERT/CC 1988]** CERT Coordination Center. "CERT/CC Advisories." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <<http://www.cert.org/advisories/>> (1988-2003).
- [CERT/CC 1996]** CERT Coordination Center. "CERT/CC Product Vulnerability Reporting Form Version 1.0." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <ftp://ftp.cert.org/pub/vul_reporting_form> (October, 1996).
- [CERT/CC 1997a]** CERT Coordination Center. "CERT/CC Incident Reporting Form, Version 5.2." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <ftp://ftp.cert.org/pub/incident_reporting_form> (December 1997; last revised April, 2000).
- [CERT/CC 1997b]** CERT Coordination Center. "The CERT Coordination Center FAQ." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <http://www.cert.org/faq/cert_faq.html> (Last revised November 21, 2002).
- [CERT/CC 1997c]** CERT Coordination Center. "CERT Security Improvement Modules." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <<http://www.cert.org/security-improvement/modules.html>> (Last revised July 9, 2001).

- [CERT/CC 1998a]** CERT Coordination Center. "Incident Reporting Guidelines." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <http://www.cert.org/tech_tips/incident_reporting.html> (May 11, 1998; last revised September 26, 2002).
- [CERT/CC 1998b]** CERT Coordination Center. "CERT Summary CS-98.05 - SPECIAL EDITION." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <<http://www.cert.org/summaries/CS-98.05.html>> (May 28, 1998).
- [CERT/CC 1998c]** CERT Coordination Center. "CERT/CC Incident Notes." 1998-2002. Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <http://www.cert.org/incident_notes/>. (Last revised December 17, 2002).
- [CERT/CC 1998d]** CERT Coordination Center. "CERT/CC Vulnerability Notes." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <<http://www.kb.cert.org/vuls/>>.
- [CERT/CC 1998e]** CERT Coordination Center. "Problems With The FTP PORT Command or Why You Don't Want Just Any PORT in a Storm." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <http://www.cert.org/tech_tips/ftp_port_attacks.html> (Last revised February 12, 1999).
- [CERT/CC 2000]** CERT Coordination Center. "Windows NT Configuration Guidelines." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <http://www.cert.org/tech_tips/win_configuration_guidelines.html>.
- [CERT/CC 2002a]** CERT Coordination Center. "CSIRT Development." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <<http://www.cert.org/csirts/>> (Last revised December 11, 2002).

- [CERT/CC 2002b]** CERT Coordination Center. "Computer Security Incident Response Team Frequently Asked Questions." Pittsburgh, Pa.: Software Engineering Institute, Carnegie Mellon University. <http://www.cert.org/csirts/csirt_faq.html> (Last revised July 2, 2002).
- [CERT-NL 1992]** CERT-NL. "CERT-NL Operational Framework, Version 2.1." Utrecht, Netherlands, June 23, 1992.
- [Chapman 1995]** Chapman, D. Brent & Zwicky, Elizabeth. *Building Internet Firewalls*, 1st ed. Sebastopol, Calif.: O'Reilly & Associates, 1995.
- [CIAC 1994]** Lawrence Livermore National Laboratories. "CIAC Bulletin: Computer Incident Advisory Capability." Livermore, Calif. <<http://ciac.llnl.gov/cgi-bin/index/notes>> (1994-1998).
- [Cormack, 2002]** Cormack, Andrew. *Writing Advisories*. JANET Guidance Notes GD/NOTE/007. London, 2002. <http://www.ja.net/documents/gn_advisories.pdf>.
- [Devargas 1995]** Devargas, Mario. *The Total Quality Management Approach to*

- [Gordon 1995]** Gordon, Sarah. "Social Engineering: Techniques and Prevention," 445-451. *Proceedings of the 12th World Conference on Computer Security, Audit and Control*. Westminster, London, U.K., October 25-27, 1995. Oxford, U.K.: Elsevier, 1995.
- [Greening 1996]** Greening, Tony. "Ask and Ye Shall Receive: A Study in 'Social Engineering.'" *ACM SIG Security, Audit & Control Review* 14, 2 (1996): 8-14.
- [Icove 1995]** Icove, David; Seger, Karl; & VonStorch, William. *Computer Crime: A Crimefighter's Handbook*. Sebastopol, CA: O'Reilly & Associates, 1995.
- [IETF 1997]** Internet Engineering Group Task Force. "An Open Specification for Pretty Good Privacy (openpgp), Charter 1997-1998." <<http://www.ietf.org/html.charters/openpgp-charter.html>> (Last revised July 31, 2001).
- [ISC 2003]** Internet Domain Survey, Internet Software Consortium. <<http://www.isc.org/ds/WWW-200301/index.html>> (January, 2003).
- Kaufman 1995]** Kaufman, Charlie; Perlman, Radia; & Spencer, Mike. *Network Security: Private Communication in a Public World*. Englewood Cliffs, N.J.: Prentice Hall, 1995.
- [Kossakowski 1994]** Kossakowski, Klaus-Peter. "The DFN-CERT Project." 6th Workshop on Computer Security Incident Handling, Forum of Incident Response and Security Teams, Boston, Mass., July 1994. <<ftp://ftp.cert.dfn.de/pub/csir/dfncert/papers/6csihw.dfncert.ps.gz>> (1994).
- [Kossakowski 2000]** Kossakowski, Klaus-Peter & Allen, Julia. "Securing Public Webservers" (CERT Security Improvement Module CMU/SEI-SIM-011). Pittsburgh, Pa.: CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University. <<http://www.cert.org/security-improvement/modules/m11.html>> (2000; last revised April 25, 2001).

- [Kossakowski 2001]** Kossakowski, Klaus-Peter. "Information Technology Incident Response Capabilities." Doktor Thesis at the University of Hamburg, Germany. Hamburg: Books on Demand, 2001 (ISBN: 3-8311-0059-4).
- [Longstaff 1993]** Longstaff, Thomas A. *Results of a Workshop on Research in Incident Handling* (CMU-SEI-93-SR-020). Pittsburgh, Pa.: CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University. <<http://www.sei.cmu.edu/publications/documents/93.reports/93.sr.020.html>> (September 1993).
- [NIST 800-12]** National Institute of Standards and Technology. *An Introduction to Computer Security: The NIST Handbook* (NIST Special Publication 800-12). Gaithersburg, Md.: National Institute of Standards and Technology.
- [NRL 1995]** Naval Research Laboratory, IS Security Group. *IS Security Incident Response Manual* (Code 1220.2). Washington, D.C.: Naval Research Laboratory, 1995.
- [NRL 1997]** Naval Research Laboratory, IS Security Group. *IS Security Incident Response Plan*. Washington, D.C.: Naval Research Laboratory, January 1997.
- [Olnes 1994]** Olnes, Jon. "Development of Security Policies." *Computers & Security* 13, 8 (1994): 628-636.
- [Pethia 1990a]** Pethia, Richard D. "Forming and Managing a Response Team." *Workshop on Computer Security Incident Handling*. Pleasanton, CA. June 1990.
- [Pethia 1990b]** Pethia, Richard D. "Developing the Response Team Network." *Workshop on Computer Security Incident Handling*. Pleasanton, CA. June 1990.

- [Pethia 1990c]** Pethia, Richard D. & van Wyk, K. R. "Computer Emergency Response: An International Problem." Pittsburgh, Pa.: CERT Coordination Center., Software Engineering Institute, Carnegie Mellon University, 1990.
- [RFC 1281]** Pethia, Richard D.; Crocker, Steve; & Fraser, Barbara. *Guidelines for the Secure Operations of the Internet* (IETF Request for Comments 1281). <<http://www.faqs.org/rfcs/rfc1281.html>> (1991).
- [RFC 1422]** Kent, S. T. & Linn, J. *Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-based Key Management* (IETF Request for Comments 1422). <<http://www.faqs.org/rfcs/rfc1422.html>> (1993).
- [RFC 1984]** IAB and IESG. *IAB and IESG Statement on Cryptographic Technology and the Internet* (IETF Request for Comments 1984). <<http://www.faqs.org/rfcs/rfc1984.html>> (1996).
- [RFC 2196]** Barbara Fraser, ed. *Site Security Handbook* (IETF Request for Comments 2196). <<http://www.faqs.org/rfcs/rfc2196.html>> (1997).
- [RFC 2350]** Brownlee, N. & Guttman, E. *Expectations for Computer Security Incident Response* (IETF Request for Comments 2350, Best Current Practice). <<http://www.faqs.org/rfcs/rfc2350.html>> (1998).
- [RFC 3067]** Arvidsson, J.; Cormack, A.; Demchenko, Y.; & Meijer, J. *TERENA's Incident Object Description and Exchange Format Requirements* (IETF Request for Comments 3067, Informational). <<http://www.faqs.org/rfcs/rfc3067.html>> (2001).
- [Schneier 1995]** Schneier, Bruce. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. Chichester, U.K.: John Wiley & Sons, 1995.
- [Shimomura 1995]** Shimomura, Tsumotu & Markoff, John. *Takedown*. London: Secker & Warburg, 1995. ISBN 0-436-20287-5.

- [Smith 1994]** Smith, Danny. *Forming an Incident Response Team*. University of Queensland: Brisbane, Australia, July 1994.
- [Stoll 1989]** Stoll, Clifford. *The Cuckoo's Egg*. Doubleday, 1989, 326pp, ISBN 0-370-31433-6.
- [TERENA 1995]** Kossakowski, Klaus-Peter, ed. "Final Report of the TERENA Task Force 'CERTs in Europe,'" Amsterdam, The Netherlands: Trans-European Research and Education Networking Association, October 1995.
- [West-Brown 1995]** West-Brown, Moira J. "Incident Trends." *Proceedings of the UNIX Network Security Conference*, Washington D.C., November 1995.
- [Wood 1998]** Wood, Charles Cresson. *Information Security Policies Made Easy*, 6th ed. Sausalito, Calif.: Baseline Software Inc., 1998. ISBN# 1-881585-04-2.

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE December 1998 (Revised April 2003)	3. REPORT TYPE AND DATES COVERED Final		
4. TITLE AND SUBTITLE Handbook for Computer Security Incident Response Teams (CSIRTs)		5. FUNDING NUMBERS F19628-00-C-0003		
6. AUTHOR(S) Moira West-Brown, Don Stikvoort, Klaus-Peter Kossakowski, Georgia Killcrece, Robin Ruefle, Mark Zajicek				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2003-HB-002	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) This document provides guidance on forming and operating a computer security incident response team (CSIRT). In particular, it helps an organization to define and document the nature and scope of a computer security incident handling service, which is the core service of a CSIRT. The document explains the functions that make up the service; how those functions interrelate; and the tools, procedures, and roles necessary to implement the service. This document also describes how CSIRTs interact with other organizations and how to handle sensitive information. In addition, operational and technical issues are covered, such as equipment, security, and staffing considerations. This document is intended to provide a valuable resource to both newly forming teams and existing teams whose services, policies, and procedures are not clearly defined or documented. The primary audience for this document is managers who are responsible for the creation or operation of a CSIRT or an incident handling service. It can also be used as a reference for all CSIRT staff, higher level managers, and others who interact with a CSIRT.				
14. SUBJECT TERMS computer security incident response team, incident handling, incident response, CSIRT, incident response service, team operations, information handling, continuity assurance, security management, risk management			15. NUMBER OF PAGES 222	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	