

Handbook for Computer Security Incident Response Teams (CSIRTs)

Moira J. West-Brown
Don Stikvoort
Klaus-Peter Kossakowski

December 1998

HANDBOOK
CMU/SEI-98-HB-001



Carnegie Mellon
Software Engineering Institute

Pittsburgh, PA 15213-3890

Handbook for Computer Security Incident Response Teams (CSIRTs)

CMU/SEI-98-HB-001

Moir J. West-Brown
Don Stikvoort
Klaus-Peter Kossakowski

December 1998

Unlimited distribution subject to the copyright.

This work was undertaken 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).

The SEI is sponsored by the U.S. Department of Defense.

Copyright © 1998 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 F19628-95-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 52.227-7013.

This document is available through Asset Source for Software Engineering Technology (ASSET): 1350 Earl L. Core Road; PO Box 3305; Morgantown, West Virginia 26505 / Phone: (304) 284-9000 or toll-free in the U.S. 1-800-547-8306 / FAX: (304) 284-9001 World Wide Web: <http://www.asset.com> / e-mail: sei@asset.com

Copies of this document are available through the National Technical Information Service (NTIS). For information on ordering, please contact NTIS directly: National Technical Information Service, U.S. Department of Commerce, Springfield, VA 22161. Phone: (703) 487-4600.

This document is also available through the Defense Technical Information Center (DTIC). DTIC provides access to and transfer of scientific and technical information for DoD personnel, DoD contractors and potential contractors, and other U.S. Government agency personnel and their contractors. To obtain a copy, please contact DTIC directly: Defense Technical Information Center / Attn: BRR / 8725 John J. Kingman Road / Suite 0944 / Ft. Belvoir, VA 22060-6218 / Phone: (703) 767-8274 or toll-free in the U.S.: 1-800 225-3842.

Table of Contents

Abstract	xi
Preface	xiii
Acknowledgements	xv
1 Introduction	1
1.1 Scope of the Document	3
1.2 Intended Audience	4
1.3 Use of This Document	5
1.4 Document Structure	5
2 Basic Issues	7
2.1 CSIRT Framework	7
2.1.1 Mission Statement	8
2.1.2 Constituency	9
2.1.2.1 Constituency Definition	9
2.1.2.2 Overlapping Constituencies	11
2.1.2.3 Relationship to Constituency	11
2.1.2.4 Promoting the CSIRT to the Constituency	13
2.1.2.5 Gaining Constituency Trust	13
2.1.3 Place in Organization	14
2.1.4 Relationship to Other Teams	16
2.2 Service and Quality Framework	18
2.2.1 Introduction to Services	19
2.2.1.1 Service Descriptions	19
2.2.1.2 CSIRT Services	19
2.2.2 Information Flow	21
2.2.3 Policies	24
2.2.3.1 Attributes	25
2.2.3.2 Content	25
2.2.3.3 Validation	27
2.2.3.4 Implementation, Maintenance, and Enforcement	28
2.2.4 Quality Assurance	28

2.2.4.1	Definition of Quality System	29
2.2.4.2	Checks: Measurement of Quality Parameters	30
2.2.4.3	Reporting and Auditing	31
2.2.4.4	Balances: Procedures to Assure Quality	32
2.2.4.5	Constituents' View of Quality	33
2.3	Adapting to Specific Needs	33
2.3.1	The Need for Flexibility	34
2.3.2	Legal Issues	36
2.3.2.1	Legal Issues Management	37
2.3.2.2	Liability	39
2.3.2.3	Disclosure of Information	42
2.3.3	Institutional Regulations	43
3	Incident Response (IR) Service	45
3.1	IR Service Description	45
3.1.1	Objective	46
3.1.2	Definition	47
3.1.3	Function Descriptions	48
3.1.4	Availability	48
3.1.5	Quality Assurance	49
3.1.6	Interactions and Information Disclosure	49
3.1.7	Interfaces with Other Services	49
3.1.8	Priority	49
3.2	IR Service Functions Overview	50
3.3	Triage Function	52
3.3.1	Use of Tracking Numbers	53
3.3.1.1	Unique Intra-CSIRT Tracking Numbers	54
3.3.1.2	Unique Inter-CSIRT Tracking Numbers	54
3.3.1.3	Tracking Numbers are Public Information	55
3.3.1.4	Tracking Number Life Cycle	55
3.3.2	Use of Standard Reporting Forms	55
3.3.3	Pre-Registration of Contact Information	56
3.4	Incident Function	57
3.4.1	Incident Life Cycle	58
3.4.2	Incident Analysis	60
3.4.2.1	The Bigger Picture	61
3.4.2.2	Analysis Depth	61
3.4.2.3	Log-File Analysis	63

3.4.2.4	Artifact Analysis	65
3.4.2.5	Analysis of Software Environment	68
3.4.2.6	Intra-Incident Web-of-Relations Analysis	69
3.4.2.7	Analysis of the Texture of Ongoing Incidents	70
3.4.3	Tracking of Incident Information	70
3.5	Announcement Function	72
3.5.1	Announcement Types	72
3.5.2	A Priori Considerations	73
3.5.2.1	Announcement Triggers	73
3.5.2.2	Categorization Criteria	74
3.5.2.3	Prioritization	74
3.5.2.4	Clearance of Information in Announcements	74
3.5.2.5	Distribution Channels	75
3.5.3	Announcement Life Cycle	75
3.5.3.1	Initiation	75
3.5.3.2	Prioritization	75
3.5.3.3	Development	76
3.5.3.4	Final Preparation	76
3.5.3.5	Distribution	77
3.6	Feedback Function	77
3.6.1.1	Life Cycle	78
3.6.1.2	FAQ and Other Default Feedback	78
3.6.1.3	Organization of Feedback Function	79
3.7	Interactions	79
3.7.1	Points of Contact	80
3.7.1.1	Incident-Related Contacts	80
3.7.1.2	Non-Incident-Related Contacts	81
3.7.1.3	Finding Contacts	81
3.7.1.4	Maintaining Contacts	82
3.7.2	Authentication	82
3.7.2.1	Social Engineering	83
3.7.2.2	Technical Possibilities and Limitations	83
3.7.2.3	Databases	84
3.7.2.4	Anonymous Information	84
3.7.3	Secure Communication	85
3.7.4	Special Considerations	85
3.7.4.1	Constituency Sites	86
3.7.4.2	IR Teams	87
3.7.4.3	Sites Outside the Constituency	90

3.7.4.4	Parent Organizations	92
3.7.4.5	Law Enforcement	92
3.7.4.6	Media	93
3.8	Information Handling	94
3.8.1	Information Collection	94
3.8.2	Information Verification	95
3.8.3	Information Categorization	95
3.8.4	Information Storage	97
3.8.5	Information Sanitizing and Disposal	97
3.8.6	Prioritization Criteria	98
3.8.6.1	By Target or Source of Attack	99
3.8.6.2	By Type or Extent of Damage	100
3.8.6.3	By Incident Type	101
3.8.6.4	Feedback Request Prioritization	101
3.8.7	Escalation Criteria	101
3.8.7.1	Individual Incident Escalation	102
3.8.7.2	Multiple Incident Escalation	103
3.8.8	Information Disclosure	104
4	Team Operations	109
4.1	Operational Elements	109
4.1.1	Work Schedules	109
4.1.2	Telecommunications	110
4.1.3	Electronic Mail (email)	110
4.1.4	Workflow Management Tools	111
4.1.5	World Wide Web Information System	111
4.1.6	IP Addresses and Domain Name	111
4.1.7	Network and Host Security	112
4.2	Fundamental Policies	113
4.2.1	Code-of-Conduct	113
4.2.2	Information Categorization Policy	114
4.2.3	Information Disclosure Policy	115
4.2.3.1	Second Level Disclosure	117
4.2.3.2	Timing of Disclosure	117
4.2.4	Media Policy	117
4.2.5	Security Policy	119
4.2.6	Human Error Policy	120
4.3	Continuity Assurance	121
4.3.1	Continuity Threats	121
4.3.1.1	Short-Term Issues	121
4.3.1.2	Medium-Term Issues	123
4.3.1.3	Long-Term Issues	123
4.3.2	Workflow Management	123

4.3.3	Out-Of-Hours Coverage	126
4.3.3.1	Hotline Coverage	126
4.3.3.2	Escalation	126
4.3.3.3	How to Reach Other Teams or Customers	127
4.3.4	Off-Site Coverage	127
4.4	Security Management	128
4.4.1.1	Use of Encryption and Signing Applications	128
4.4.1.2	Key Management and Certification	129
4.4.1.3	Firewalls and Network Security	131
4.4.1.4	Providing Off-Site Access to Local Facilities	132
4.4.1.5	Physical Security	132
4.4.1.6	Disaster Handling	133
4.4.1.7	Handling Internal Security Incidents	133
4.5	Staff Issues	134
4.5.1	CSIRT Staff	134
4.5.2	Hiring Staff	136
4.5.3	Arrival and Exit Procedures	138
4.5.4	Training Staff	138
4.5.5	Retaining Staff	140
4.5.6	Extension of Staff	141
5	Closing Remarks	143
	About the Authors	145
	Moira J. West-Brown <mjw@cert.org>	145
	Klaus-Peter Kossakowski <kpk@work.de>	146
	Don Stikvoort <don@elsinore.nl>	146
	Bibliography	149
	Glossary	157
	Acronyms and Abbreviations	157
	Glossary Terms	160
	Artifact (a.k.a. Critter)	160
	Authenticity	160
	Bugtraq	160
	Computer Security Incident	160
	Computer Security Incident Response (CSIR)	160
	Constituency	161
	Intruder	161

Liability	161
Policy	161
Procedure	161
Remnant Files	161
Security Policy	161
Site	161
Site Security Contact (SSC)	161
Social Engineering	162
Triage	162
Trojan Horse	162
Vulnerability	162

List of Figures

Figure 1: CSIRT Within an Organization	15
Figure 2: CSIRT Peer Relationships	17
Figure 3: Service and Quality Framework as Derived from Mission Statement	18
Figure 4: IR Service Functions	51

List of Tables

Table 1:	Examples of CSIRT Types With Associated Missions and Constituencies	9
Table 2:	Possible Authority Relationships Between a CSIRT and its Constituency	12
Table 3:	Service Description Attributes	19
Table 4:	Example CSIRT Services	20
Table 5:	Examples of Possible Information Flow to and from the IR Service	23
Table 6:	Basic Policy Attributes	26
Table 7:	Policy Content Features	27
Table 8:	Examples of Dynamic Environment Factors and Their Impact on CSIRTs	35
Table 9:	Examples of Liability Issues Arising From Omission	40
Table 10:	Examples of Liability Issues Arising From the Content of Signed Contracts	40
Table 11:	Examples of Liability Issues Arising From Information Disclosure	41
Table 12:	Range of Possible IR Service Objectives Based on Differing Team Types	46
Table 13:	Possible Instantiations of Incident Function Attributes	58
Table 14:	Analysis Depth Factors	62
Table 15:	Notable Characteristics of Log-Files	64
Table 16:	Incident Tracking Information	71
Table 17:	Possible Inter-Team Support Types	88
Table 18:	Considerations for Information Sharing	89

Abstract

This document provides guidance on the generic issues to consider when 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 response (CSIR) service, which is the core service of a CSIRT. The document discusses 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 often sensitive information. In addition, operational and technical issues are addressed, 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 consists of managers responsible for the creation or operation of a CSIRT or a CSIR service. It can also be used as a reference for all CSIRT staff, higher-level managers, and others who interact with a CSIRT.

Preface

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 electronic mail (email) discussion between the authors 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 and bullets. 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.

If you have comments on this document, if you want to share your opinions, or if you have suggested additions 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 the following address:

`csirt-handbook@cert.org`

Acknowledgements

We have many people to thank for their contributions that have made this 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:

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)
Hannes P. Lubich (SWITCH-CERT, now with Bank Julius Baer)
Jorgen Bo Madsen (NORDUnet CERT, now with Tele Danmark)

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

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.

Based on reviewer feedback, we significantly revised the structure of the first draft and incorporated many of their excellent comments. We were unable to incorporate some of their comments due to lack of time, but hope to address those and other constructive comments that we receive in a future revision.

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.

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.

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 incidents on the Internet. The CERT/CC was one of the first organizations of this type—a computer security incident response team (CSIRT¹).

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 electronic mail (email) address that you can contact for help if you have or suspect a computer security incident. A CSIRT service doesn’t

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

¹ Within the computer security arena, these teams are often simply referred to as incident response teams (IRTs).

normally provide response by showing up on your doorstep (some do offer that luxury); 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 are more than 36 million hosts on the Internet and a corresponding increase in intruder activity. 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. This has resulted in a surge of interest in the formation of CSIRTs.

Hundreds of CSIRTs around the world have since formed; and they, like newly forming CSIRTs today, face many challenges as they strive to become operational. There are various documents and tutorials available [Kossakowski 94a, Smith 94, Smith 96, Sparks 97] to help an organization to understand the need for a CSIRT, to obtain funding for it, and to define the main functional issues to consider. But little is available in the area of operational policies and procedures. 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. Unfortunately, they are rarely able to obtain documented guidance on establishing appropriate and reliable services. Either existing teams have nothing documented to share, or they are unwilling to share their documentation due to its sensitive nature. Seeking expert advice is difficult too because there is a shortage of experts in the field. Those existing experts are highly sought after, have little time to make available, and are expensive to engage. As a result, newly forming teams are left to fend for themselves, learning from their own experiences, and often making expensive mistakes rather than having the benefit of others' experience. This can be a slow, painful, and dangerous process where the absence of suitable knowledge of appropriate services and suitable policies and procedures can result in the demise of a team.

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 good 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.

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 CSIRT services, policies, and procedures appropriate for a given situation and identifying operational issues.

In particular, this document helps an organization to define and document the nature and scope of a computer security incident response (CSIR) service². This is the core service of a CSIRT. 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 and 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. Almost invariably, 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 the document.

² For simplicity, the CSIR service will be referred to as the incident response (IR) service throughout the remainder of this document.

Some CSIRTs have dedicated staff while others pull together part-time, volunteer staff and trusted security experts to address a given security crisis. A CSIRT's staff is its interface with the world, and the image that its staff members project through the way that they conduct themselves and the quality of service that they provide are paramount to the CSIRT's success. Finding appropriately qualified staff is difficult since they 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 IR service, such as vulnerability analysis and intrusion detection. However, a detailed description of those services and 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 CSIRTs has not kept pace with Internet growth and intruder activity. Many more organizations will recognize the need for a CSIRT to address their specific needs. Anticipating this need, we have targeted this document 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 IR service
- the operation of an IR 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:

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

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 stage 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 of use to 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 IR service.

In addition, an existing team can use this document to ensure that they have covered the main issues and options that they consider appropriate for their organization when developing their IR 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. However, it is important to note that this material is only provided for reference and guidance. We do not intend to dictate the range or content of services, policies, and procedures that any given team should implement. These must be determined on a per-team basis. Hence, we encourage you to use the material provided in this document to understand the issues appropriate for your team's unique environment and decide which approach that you should adopt based on your particular goals, needs, and situation.

1.4 Document Structure

The rest of this document is organized as follows. Chapter 2 presents the basic frameworks of the CSIRT model and describes the 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 IR service. This sets the context for the main focus for this document, the IR service, which is described in detail in Chapter 3. Chapter 3 describes the construction of an IR service and its functional components. Additionally, it discusses the range and nature of interactions that are associated with an IR 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 consid-

erations, are not all exclusive to an IR 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 IR service discussed later in this chapter and covered in depth in Chapter 3. Without providing the IR 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 core IR service, the team is not a CSIRT.

This chapter presents the basic frameworks 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, and 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, of which 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 in 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 incident response.

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 needs. Each team needs to define its own set of criteria and operating guidelines.

To obtain that goal in a structured fashion, it is best to start with and 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”:

- 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 they interact with. 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 if 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 unlikely to improve. However, the mission statement of every CSIRT must have the backing of senior management (the Corporate Security Officer, Head of Information Technology, Board of Directors, or equivalent) in the parent organization. Without such backing the CSIRT will struggle to obtain recognition and resources.

A mission statement is imperative 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.

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. 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 role, purpose and the framework within which it operates, by other parties that 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 parties. 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. 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 shows examples of how CSIRTs of different types may fulfill differing missions and serve differing 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

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

An essential CSIRT task 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 aspects of constituency are detailed below.

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. It can be difficult, or even impossible for a network service provider 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 the Pennsylvania State University response team can be defined simply as “Pennsylvania State University” and as the domain “*.psu.edu”.

However, even if 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.

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. The multiple constituencies might intersect, be sub- or supersets, or be totally separate. For instance a technical CSIRT might provide general security information on the product it specializes in to an unbounded constituency 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 the product.

Even if a CSIRT has a bounded constituency it will still have to deal with information associated with or coming from parties that do not belong to their constituency. For instance, a CSIRT providing an incident response service to a bounded constituency will undoubtedly wish to accept incident reports that directly affect its constituency from parties outside of its constituency and act appropriately with that information to ensure that it reaches appropriate points of contact and is coordinated within its constituency. Many CSIRTs act as a coordination point between their constituency and other parties external to it such as other CSIRTs, system administrators, vendors, law enforcement, legal counsel, and the media. These interactions can vary from pure relaying of requests to complete sharing of data and full cooperation. It is important that a CSIRT decides, documents, and states how these 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 (passing reports from other teams and sites to their constituents) to their constituency as other sites and teams do not know if a given constituent falls within the CSIRT’s constituency and so can not report the activity directly to the appropriate CSIRT. As a result their customers are contacted directly by other sites or CSIRTs involved in an incident, then as necessary the customers seek support of their own CSIRT. However, this is a good example of how weblike trust relationships come into play rather than teams being constrained by a true hierarchical model.

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 overlapping constituency situations will result in confusion between the CSIRTs themselves 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 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. These organizations fall into the constituencies of two teams:

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

If any incidents affecting German government sites are reported to DFN-CERT, they forward the information to BSI-CERT; and further follow-up and actions are coordinated between 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 German university called the CERT/CC in the U.S. for assistance with an incident at 9am local time in Germany, which was 3am local time in the U.S. A CERT/CC staff member was paged and provided the administrator immediate assistance. When asked, the administrator in Germany, new to his job, was unaware of the service offered by DFN-CERT in Germany. But once he knew of the existence of DFN-CERT and how it could offer them more appropriate service in terms of local needs, language, and time zone, he then contacted DFN-CERT.

2.1.2.3 Relationship to Constituency

The nature of the relationship between a CSIRT and its constituency will directly impact 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.

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. In other words, 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 in the role of an advocate or advisory capacity.

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

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 IR, or vulnerability analysis and response, or training services can be offered. However, services such as incident tracing and intrusion detection (listed in Table 4 “Example 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 patch availability for intruder exploitation of 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 appropriate patches to address the threat. 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 appropriate patches have been installed to address the threat. 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 apply the suggested changes. However, the CSIRT cannot force the constituency to apply the suggested changes.

2.1.2.4 Promoting the CSIRT to the Constituency

Once the constituency has been defined, it will be important (regardless of 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, a constituent needs to know which CSIRTs are offering them service.

For all practical purposes, 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 make reports to it
- others: those 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 news groups
- CSIRT or organizational information/Web server
- presentation, workshop, and tutorial materials
- general awareness materials and news letters (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. It 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 view 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. That 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 high-brow 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. 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 IR 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.

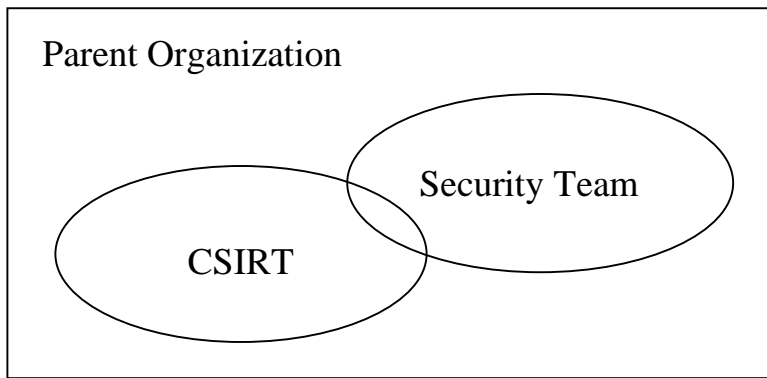


Figure 1: CSIRT Within an Organization

Within 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 organizations IT security department.

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

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 interact and coordinate their responsibilities.

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

Example: 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 and how to coordinate with other groups to ensure that it does not operate in isolation. 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 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 in the world, and a growing number are served by a CSIRT. So 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 and defining constituency and place within organization are not sufficient without also covering the coordination issue.

In the realm of CSIRTs as it exists today, there is some hierarchical structure between CSIRTs. There are teams providing service to clearly marked constituencies and others who serve in a coordination role across groups (commonly national or international) of CSIRTs. However the 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 be more cautious with other teams that they have less experience of.

Some formal hierarchies do exist, such as within the U.S. military. The U.S. Army, Air Force, and Navy (ACERT/CC³, AFCERT, and NAVCIRT, respectively) teams serve their own constituencies and the U.S. Department of Defense ASSIST team coordinates across the various 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 are informed of all activity

³ Plays a coordinating role itself across geographically dispersed U.S. Army teams

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 amongst other CSIRTs. The example in Figure 2 depicts both CSIRTs A and B as coordinating CSIRTs. In addition to coordinating amongst CSIRTs C and D, CSIRT B also has another constituency component that is not covered by C or D and is served directly by B. Whereas, 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 hierarchy because CSIRTs E and F communicate directly with each other.

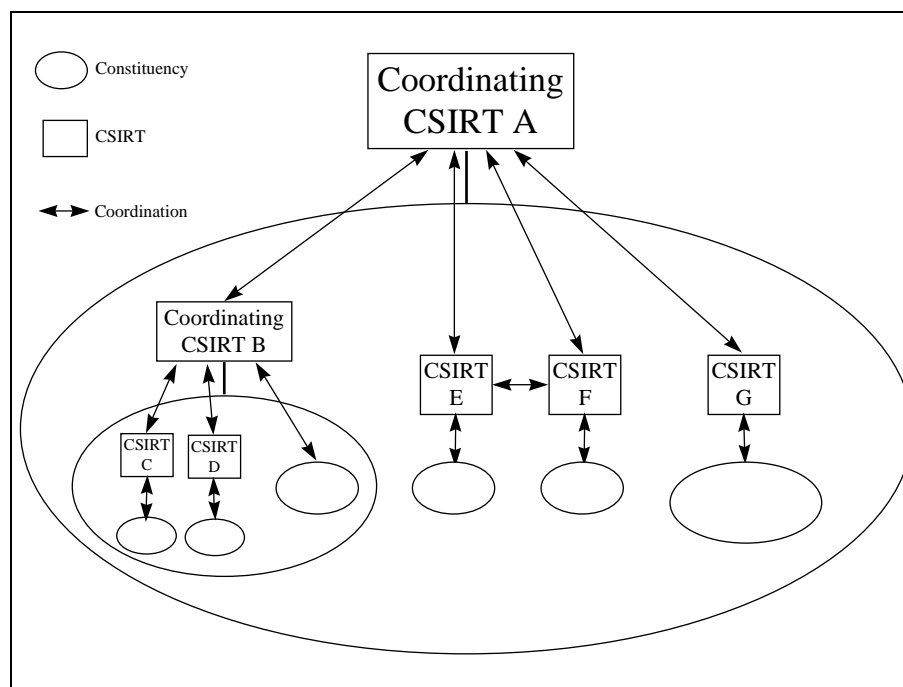


Figure 2: CSIRT Peer Relationships

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, AusCERT), fee-for-service response teams (e.g., IBM-ERS, Global Integrity's REACT), teams for commercial organizations (e.g., Motorola's MCERT, Boeing's BCERT), network service provider teams (e.g., MOREnet, ANS), and universities (e.g., Pennsylvania State University's PSU-CERT, Stanford University's SUNSeT) can all be represented using this approach.

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.

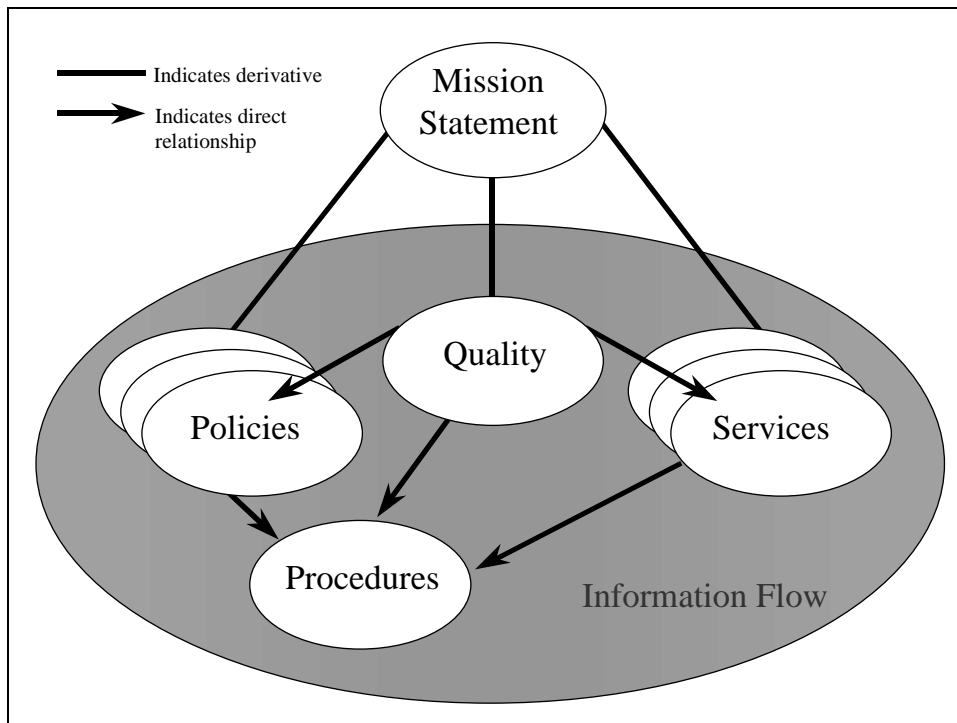


Figure 3: Service and Quality Framework as Derived from Mission Statement

Following this framework, the three derivatives of the mission statement will be discussed in more detail. Information flow would naturally be the fourth topic to discuss. However, for the purposes of this handbook, the flow of information within a team is not a basic CSIRT issue in its own right. Information flow is of basic interest only where it pertains to external communication. So information flow (the flows of information inside the team) are clearly not basic CSIRT issues; that topic will be discussed in relation to services in Section 2.2.2 “Information Flow” and wherever relevant in the subsequent treatment of CSIRT issues in Chapters 3 and 4.

2.2.1 Introduction to Services

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 IR service, which is the focus of this document, will be described in detail in Chapter 3. However to provide the necessary context for the discussion of the IR service, this section introduces issues that are generic to all CSIRT services and provides a brief discussion of other services that a CSIRT might offer.

2.2.1.1 Service Descriptions

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, the descriptions should include an explanation of the items listed in Table 3.

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 regards 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

Table 3: Service Description Attributes

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.

2.2.1.2 CSIRT Services

For a team to be considered a CSIRT, it only needs to provide the IR service. However, CSIRTs commonly offer other services in addition to the IR service, 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 IT security).

Detailed descriptions of other services that a CSIRT may provide are outside the scope of this document. In addition to the mandatory IR service, Table 4 lists some of the more common services that CSIRTs may provide and the form that those services might take. Of these additional services, some are more common (such as announcements, or vulnerability analysis and response) as they are closely associated to IR.

Mandatory CSIRT Services	
Service Name	Description
Incident Response	Provide a focal point for reporting computer security incidents that provides coordinated support in response (and indication to others) to such reports.
Common CSIRT Services	
Service Name	Description
Announcements	Disseminate information on protective measures to take against existing or upcoming security threats.
Vulnerability Analysis and Response	Serve as a focal point for reporting computer security vulnerabilities that provides coordinated support in response to such reports.
Artifact ⁴ Analysis and Response	Generate technical analysis reports pertaining to malicious code.
Education	Provide training to promote security awareness and improve expertise.
Incident Tracing	Support tracking and tracing intruder activity.
Intrusion Detection	Support active detection of intruder activities.
Auditing and Penetration Testing	Support security auditing or penetration testing of computer systems and networks.
Security Consulting	Provide expert advice for computer security and network issues both for operation and development procurement.
Risk Analysis	Undertake risk analysis assessments.
Technology Watch	Provide information on upcoming technology that may pose security threats.
Security Product Development	Design and develop security tools for incident detection and prevention.
Collaboration	Establish collaborative relationships with other entities such as law enforcement, service providers and the telephone company.
Coordination	Interact with both internal and external parties to develop and maintain trust relationships.

Table 4: Example CSIRT Services

⁴ Instances of malicious code: see Glossary.

Some of the services in Table 4 are clearly reactive, such as incident response and incident tracing, and others are clearly proactive, such as technology watch. However, depending on how the other services in Table 4 are either implemented or initiated, they could be proactive, reactive, or both.

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 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.

Each of the additional services can be addressed in a similar fashion to the IR service using service descriptions described above and in a similar fashion to the handling of the IR service presented in Chapter 3.

2.2.2 Information Flow

Whatever range of services is offered by a CSIRT, it is important to understand which of those services are in some way related to each other and what these interdependencies are. In particular, it is necessary to specify the interfaces between the services and any associated information flow between them. 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 helpdesk which directs the requests to the appropriate service, or each service could advertise their own contact information and directly handle requests specific to their service.

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 as policy and procedure changes occur.

It may be necessary to give different priorities to the same type of request depending on the source of the request. For example, the IR service could obtain simultaneous requests for incident statistics from both the vulnerability service (e.g., to assess the frequency with which a given vulnerability is exploited and prioritize further action) and the education service (in the process of updating public presentation materials). A higher priority would likely be given to the request of the vulnerability service. This example also raises the issue of information sharing again. The information provided to the vulnerability service would most likely include details on the frequency of incidents reported to involve specific methods of exploitation. The information to the education service would be sanitized for a public offering, at least in such a way as to remove details of yet unsolved exploitation methods.

Some basic examples of possible information-flow relationships between the most commonly provided CSIRT services and the IR 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 IR service.

Service Name	Information flow to IR service	Information flow from IR service
Announcements	Warn of current attack scenarios	Statistics or status report; New attack profiles to consider or research.
Vulnerability Analysis and Response	How to protect against exploitation of specific vulnerabilities.	Possible existence of new vulnerabilities.
Artifact Analysis and Response	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	None.	Practical examples and motivation; Knowledge.
Incident Tracing	Identify new or additional attack profile; Identify new sites impacted by incident.	Request for tracing as result of ongoing incident.
Intrusion Detection	New incident report.	New attack profile to check for.
Auditing and Penetration Testing	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.
Security Product Development	Availability of new tools for constituency use.	Need for products; Provide view of current practices.
Collaboration	Details of appropriate requirements for interactions with collaborating parties.	Need for new collaborations; Operational problems with existing collaborations.
Coordination	Provide up-to-date details of trusted parties.	Need for specific new trusted relationships; Operational problems with existing relationships.

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

Due to the limited resources available within 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 serv-

ices into one service; the substituting parts can then be labeled “functions” within the service according to the terminology of this handbook.

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 IR service) a penetration testing service. During penetration tests, administrators of the systems and networks involved are rarely made aware that the penetration test will take place. So, if during the test an insecure host is penetrated, the system administrator for the penetrated machine may notice the activity, perceive it as a break-in, and report it as such to the CSIRT. If the penetration service provides the IR service with advance notification of the test, the IR team may first verify with the penetration team if the activity is due to the test. If not alerted in advance, the IR team 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 the intrusion tracking group. As a result, precious resources of the CSIRT may be needlessly wasted. More importantly, the reputation of the CSIRT may be also 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, the left hand does not know what the right hand is doing.

2.2.3 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 service or training service), policies are mainly quantities internal to the team that dictate appropriate behaviors in some specific field. Examples include “categorization of information,” security policy, media policy, and code-of-conduct. The latter two examples may prompt the question “but these are hardly only internal, they have a lot to do with external communication.” True enough, but this external aspect is not something offered to the customer, it is not a service in itself, it merely impacts the manner in which the service is delivered.

A policy may be service-specific; an incident response service may require a specific policy on caller authentication (e.g., laid down in a procedure for verification of a caller before inci-

dent information can be discussed). Caller authentication may not be necessary within another service such as education 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.

2.2.3.1 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 media policy again as an example) 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 those attributes that every policy should have.

2.2.3.2 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, how to classify incoming information, and how to 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 (all drawn from the media policy arena).

Attribute	Description
Management Endorsement	Just as with the mission statement, without endorsement from senior management, a policy cannot be enforced.
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. Tip: If possible (according to your disclosure restrictions), ask an average person who is not in security and not in IT to read your policies. If they 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 includes a lot of procedures. Security policies in practice often tend to be not concise, confusingly mixing the management aspect (the policy) with the operational aspect (the procedures), resulting in a mixture that nobody really cares for.
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.
Useable	Avoid statements that sound nice but are of no use as they are open to interpretation, like “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 useable, 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.

Table 6: Basic Policy Attributes

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, such as media, media liaison, and other staff.
Responsibility	Duties and responsibilities of the identified parties should be defined, when appropriate (you can not 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 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.
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 IR service.
Maintenance	Describe responsibilities and guidelines for document maintenance and update.
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 what the policy is about, especially in the case of a new team member.

Table 7: Policy Content Features

2.2.3.3 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 just means finding out if all the great ideas inside the policy, written down so well, also survive when compared to 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:

- Where possible, ensure that the people responsible for the policy validation are not the same people who created the policy, thus hoping to avoid conflicts of interest and blind spots.
- Pay particular attention to validating the policy attributes and content features detailed in Tables 6 and 7 to ensure that policies are not so ambiguous that anyone can apply their own interpretation to them.

- Undertake consistency checking of the policy in relation to other policies, services, and procedures; and also within the policy itself.
Example: When espousing network security yet using the practice of transmitting passwords in clear text, one's security policy will be inconsistent due to this contradiction.
- Validate implementability and enforceability. Pilot-implementing the policy, then choosing some worst-case scenarios and checking on real-life behavior including enforceability is the best way to accomplish this.

2.2.3.4 Implementation, Maintenance, and Enforcement

Validation is followed by feedback to the policy makers, revising the policy, and finally (maybe after repeated validations) implementing the policy.

Once that is done, the policy will need to be maintained, i.e., 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, with regards to media policy, to check if the media is indeed informed within a pre-set number of hours following a media request for information. Clearly this real life behavior could not have been previously measured within the validation phase.

Both the checks originating from the validation process and the newly defined ones 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 2.2.4 “Quality Assurance.” There, it is also implied that every policy must have its regular maintainer who keeps track of the policy’s real life behavior in relation to quality of service and proposes changes to the policy if appropriate. Things change over time, and no policy should be implemented once and used “as is” forever on. The excuse “That’s the way it has always been done” is not acceptable.

2.2.4 Quality Assurance

Defining services (such as IR or vulnerability analysis and response), policies (such as security policy and code-of-conduct), 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, and liability and penalty clauses.

In the CSIRT arena, standard approaches are rarely used. A few teams attempt to at least prioritize incidents and work on what they regard as high priority incidents first. Another team reported their most commonly used measure of quality to be “absence of complaints reaching senior management.” However, these were exceptions, as very few teams undertake QA (Quality Assurance) either formally or informally. Lack of QA results in inconsistencies in service, 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: quality system definition, 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.

2.2.4.1 Definition of Quality System

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

The quality system should be defined using a top-down approach, starting with the mission statement 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 is the sum for all policies, services, service-functions and procedures: all of these elements will have their own subset of unique quality parameters, however, in some cases individual parameters may be common between e.g., services or service-functions. It is important to bear in mind however that "quality" is a dynamic quantity, 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 response services. Obviously it is then practical to define two different sets of quality parameters, one for the incident response, the other for the vulnerability response service. A typical parameter for the incident response set would be the maximum time that it takes to respond to a constituent's initial incident report. A parameter for the vulnerability response set would 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 response services also yielding good examples of quality parameters. For example, such a parameter is the maximum time that a vulnerability service

should take to provide an assessment of a vulnerability when an incident response service finds 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 (the term “service event” introduced below is best defined by example: e.g., an incident or a vulnerability report):

- response time for service events 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.

Example: Parameter: follow-up time on vulnerability reports. Value: for all non-urgent vulnerabilities, the CSIRT will follow-up with a constituent within 5 working days of the initial 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 quality 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 5 days.” These are communicated to the constituency, who need not be explicitly aware when the CSIRT is in crisis mode.

2.2.4.2 Checks: Measurement of Quality Parameters

It is insufficient to just define a quality system if you do not check in real life whether or not it lives up to expectations. Checking quality parameters (measuring real-life behavior) is thus an essential part of any QA system.

This demand explains why quality parameters should be clear-cut and preferably quantifiable: it's 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.

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 thing 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.

It is useful noting that the frequency used to check on your parameters is really a quality parameter in itself. Its value must be carefully optimized: 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 getting the work done.

2.2.4.3 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, such as: to operational management, to overall management, to the constituency, to the world; to name a few obvious categories. When workflow management software is in place there is a tendency (or desire) to make reports automatically available. This often results in the generation of reports with either no information or too much information.

Regular auditing of the QA check-system itself is necessary to ensure the quality. The system must be checked for both sloppiness and inadequacy. To help limit sloppiness, it is necessary to

- minimize the number of procedures necessary and make them crystal-clear.
- ensure that the CSIRT staff members understand why there are good and reasonable procedures, enhance their motivation.
- forget about the tiny details: It helps staff motivation if they are allowed to think for themselves (and besides, it is impossible to make rules for everything).
- do audits and feed the results into review cycles.

One common mistake is to make large complicated rules and to compensate for this by doing very rigid audits. 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: help assuring quality.

A quality check-system can become inadequate even if it was perfectly adequate when originally set-up. This of course is because quality parameters can change.

Example: If you define the initial response time to incident reports by customers as a parameter this may be fine until you introduce an automatic email response service which is very fast indeed but whose speed probably is not the quality parameter you set out to measure.

2.2.4.4 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 for when standard enforcement procedures fail, or when 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 lawsuits. These procedures and clauses are summarized using the word “balances”: no assurance without checks and balances.

In the demanding environment of a CSIRT, where staff stress levels are high and resources are stretched, it is important to ensure that staff members are able to accomplish their work to a high standard of quality without overburdening them with unnecessary hurdles. So there is the need to get the balance right 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.

Also it is advisable to give constituents methods to enforce quality, though this will normally be an indirect process. Not only does this “sell well,” but also the best quality judgment often comes from those who actually use (or suffer) the service. The convenient way of granting constituents influence is by implementing measures such as user groups and advisory boards. Admirable though these measures are, the most effective way probably is by implementing penalty clauses, meaning you have to pay, or refund the customer money if you perform 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 staff’s perspective, escalation is usually part of the daily routine. 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. This is business for 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 on 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 on the long run and also creates the possibility of making 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. Only too often procedures are made in a 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.

2.2.4.5 Constituents' View of Quality

The set(s) of quality parameters for internal use must be complete to assure that an appropriate quality level is maintained with respect to the mission statement. However, the set of quality parameters communicated to the constituent is some subset of those used internally. The subset can range from nothing to the full set being communicated.

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.” From the same commercial point of view however and sometimes also from a liability point of view, it is wise only to communicate those parameters that are easy 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 parameters that are impossible to quantify, however useful these may be to help assure overall quality. Constituents tend to dislike what they cannot grasp.

2.3 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: An organization with significant computer virus problems will build a CSIRT with at least a proper virus handling capability.

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

Example: A CSIRT 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 itself to the extra stringent security measures, and will have to extensively screen its personnel.

We will not try and generalize the topic of adaptation in this chapter; the subject is too specific to each team's situation. Clearly adaptation starts when defining the mission statement and services. Naturally it must be reflected in the quality assurance system. But where adaptation will be most evident is in a CSIRT's policies and procedures—and in the rather practical treatment in Chapter 4 of team operations, including fundamental policies, the topic implicitly surfaces, especially in the examples.

Example: A military CSIRT will have a rather restrained media policy.

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

However, two topics remain that deserve attention at this level, and they are detailed below. The first one is about the *general* ability to readily adapt to arising circumstances that every CSIRT must have, is it to do a proper job. The second topic is that of law, liability, and regulation.

2.3.1 The Need for Flexibility

CSIRTs need to be prepared for the dynamic environment of computer security incident response. 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 to the resulting impact on the CSIRT, are given in Table 8.

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 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 that it interacts with.	Situations will arise when the resources within an unprepared CSIRT may be insufficient to respond effectively to meet the conflicting demands placed upon it.

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

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 change over time. CSIRTs must ensure that they have flexible procedures and policies to enable the team to easily adapt to change, whether the change results from a variation in work load, technical focus, or legal issues.

Although these factors are usually outside the direct control of a CSIRT, planning can ensure that the team is prepared for these issues. To address these factors, the CSIRT should

- Be prepared to 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 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 legal counsel and law enforcement.
- Ensure service definitions, policies, and procedures are not so rigorous that they do not anticipate and allow for unexpected circumstances.

Most of these issues are dealt with 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.

2.3.2 Legal Issues

As we are not legal experts, we can only offer opinions of 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, as 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 seek legal advice when making major changes in policy or operating procedures.

Legal issues are a bit like quality assurance: permeating 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 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) 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 IR 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 and then focus on the important topics of liability and the main cause of liability, disclosure of information.

2.3.2.1 Legal Issues Management

Management of legal issues involving CSIRT teams means exercising a coherent view on the legal issues that the team faces. Legal advice should be given by a fixed set of people (mainly legal experts) experienced in this area and with an understanding of 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. Especially when 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.

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.

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 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 when for the best of reasons, you may comply with your own procedures, but it is very doubtful whether you are giving due care to meeting your service goals. Depending on local laws, the interpretation may determine that the latter is more important than following procedures to the letter, and may thus find you wanting.

Waivers and Disclaimers

Disclaimers are often found in many places: service descriptions, policies, Web site, 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. Since 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 wherever 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 hatches 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 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 IR work. It 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. To make such agreements more than just a psychological safeguard, the legal advisors should gauge them.

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 you only save your logs for a year and have stated so publicly, and there is no law against this, after two years nobody can complain if logs are not available any more.

The second point is where the proactive measures come in, and the legal advisors should lend a hand. 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 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 seizes CSIRT equipment).

2.3.2.2 Liability

A liability issue is everything that you say, do, or write; or that you omit to 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 really made a big mess of your operations resulting in damage to other parties, such as your constituents, in the process.

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.

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 9: Examples of Liability Issues Arising From Omission

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 he experiences an intrusion and seeks 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, he may well try to get back some of it through you and will not settle for an excuse not related to work.

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

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 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 his 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.

Table 11: Examples of Liability Issues Arising From Information Disclosure

How to limit your liability is again asking for the 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 or are legally unclear.
- Make disclaimers legally sound.
- 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.

2.3.2.3 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, 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 can not 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, his reputation, or his business.

Example: If a site is under an on-going 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.

2.3.3 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

3 Incident Response (IR) 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 IR service in detail. In this chapter we will describe the fundamental components that make up such services and the procedures that need to be in place to support such an operation.

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

- **specification—the logical dimension**
A description of the purpose and structure of the IR 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 IR service (or, for that matter, for any CSIRT service): interactions (Section 3.7 “Interactions”) and information handling (Section 3.8 “Information Handling”).

3.1 IR 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 that it interacts with; these definitions might be provided at different levels of abstraction. As discussed in Section 2.2.1.1 “Service Descriptions,” 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 IR 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 (e.g., to the CSIRT’s constituency), the results of the IETF working group “Guidelines and Recommendations for Incident Processing” (GRIP) should be consulted [RFC 2350]. 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 95].

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, the objectives for the IR service will be derived from the CSIRT mission statement, which in turn was derived from the mission of the security team of the parent organization. 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 IR service objectives based on different types of teams with differing missions.

CSIRT Type	Nature of Mission	Possible IR 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 co-ordination with other CSIRTs around the world; Through incident response 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 in regard 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.
IT Vendor	Improve the security of its products.	Provide technical support in response to vulnerabilities, coordinates with CSIRTs to analyze the basic source of incidents. Create and disclose public alerts about new patches and best current practice.
Corporation	Improve the security of the corporation's information infrastructure and minimize threat of damage resulting from intrusions.	Provide a center of excellence for incident response support to system and network administrators and system users within the corporation; Provide on-site technical support for incidents impacting company systems to isolate and recover from intruder threats.

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

3.1.2 Definition

Before you can describe how your IR 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 IR 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, teams (on the national, organizational, and service provider level) exist today that provide a minimal IR service consisting of simple instantiations of the triage, incident, and feedback functions (see Section 3.2 "IR Service Functions Overview") all rolled into one. 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 among the parties affected by an incident. In addition, some provide suggestions on the approach that a constituent might wish to adopt in response to the incident, provide an encrypted communications path, or (in the case of a national team) provide a language translation service.

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 full-blown IR service with all functions in place independently. The lack of in-depth technical expertise prevents the team from providing an in-depth incident function, i.e., not being able to fully grasp what specific incidents are technically about. In this case, the team may rely on information generated by more technically adept teams to use and disclose within their own constituency. By relying on the services of other teams, the function will degrade to the relaying of information.

With an understanding of the available resources, limiting factors, mechanisms that you may be able to leverage off within your existing organizational structure, and the purpose that you are trying to achieve, it should be possible to define the IR 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 IR functions.

It might be appropriate to produce two resulting service descriptions based on the same set of definitions. One description, written for external consumption, providing information such as to whom the service is available (within a view of the overall IR service). The other description, written for internal consumption, would include a view of the implementation of the service through the functions that constitute it and will include the external description as well, so the external description should really be a subset of the internal description. De-

pending 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 IR field.

3.1.3 Function Descriptions

The incident response service is usually composed of four functions: triage, incident, announcement, and feedback. A 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) IR 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 IR service description that will be made available to the parties that may access the IR service. However, various 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 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 incident function might be accessible during business hours, or on a 24x7 (twenty-four hours a day, seven 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 via completion of mandatory information requested via 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 impacted 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 Sections 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 internally within the IR team between the IR service and other CSIRT services with which it interacts, depends on what other services you provide. E.g., 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 IR 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 incident 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 incident function will leave few resources for other activities, which will be apparent to the constituency.

However, triage is a pre-requisite for the incident function to operate effectively. So limited triage might take place at a reduced level for all feedback and announcement needs. Until the team can revert to its usual operating state detailed triage effort will be focused on the incident function and explaining the current situation to other requesters.

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

3.2 IR Service Functions Overview

As stated above, the incident response service usually consists of the triage, incident, announcement and feedback functions (see Figure 4). These functions and their relationships are explained below and are covered in more detail in the next four sections.

It is important to realize that many CSIRTs exist that correspond to the functional specification portrayed in Figure 4, although they may differ greatly in their implementation. The differences occur due to factors such as funding, available expertise, or organizational structure. Some of these differences were discussed previously in Section 3.1.2 “Definition.”

Example: In a small team, the IR functions may not be individually distinct; a single person (with the necessary skill set) may provide them. A larger team may set up a help-desk composed of staff with a limited range of technical skills to handle the triage and feedback functions, and separately provide the incident function with staff with a higher technical skill set.

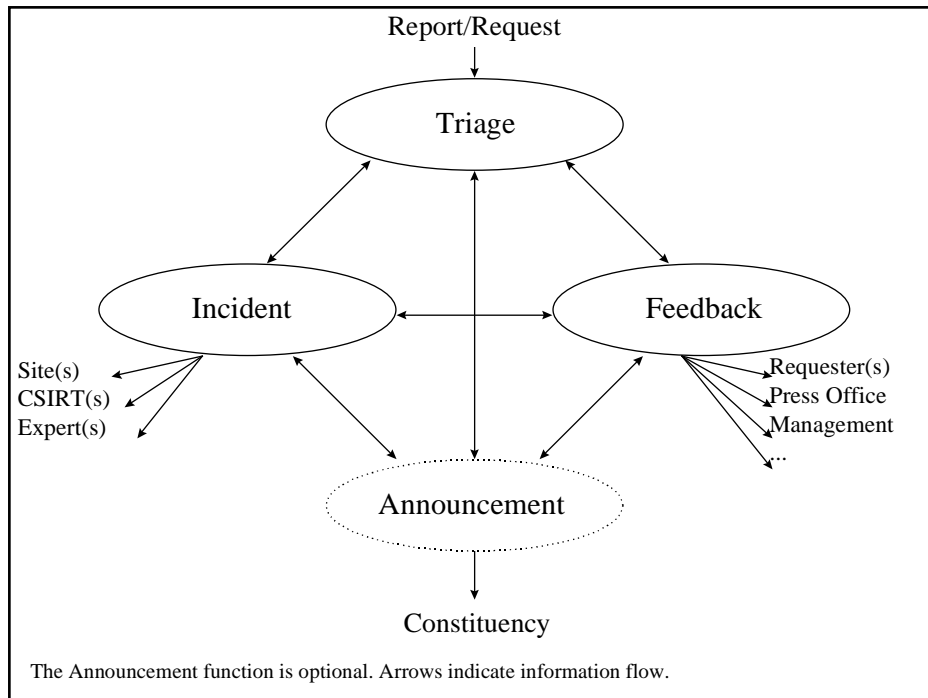


Figure 4: IR Service Functions

Triage Function

Provides a single point of contact and the focal point for accepting, collecting, sorting, ordering, and passing on incoming information for the IR service. It supports different input channels suitable to the needs of the team and constituency. An initial priority and possibly an associated tracking number is assigned to any apparent new event. The triage function might also undertake additional steps such as archiving, translation, or media changes.

Incident Function

Provides support and guidance related to suspected or confirmed computer security incidents.

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. For the purpose of this document, the scope of this function will be limited to its direct applicability with the IR service. However, within a CSIRT providing a broader range of services, 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

Provides support for giving feedback on issues not directly related to specific incidents. Feedback can be provided both on explicit request (e.g., by the media) or unsolicited, on a regular basis (such as annual reports) 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 IR 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 IR 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 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 the triage 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 IR 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 IR 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 hard-copy materials, it is common for the triage function to ensure that this information is entered online or a reference made online to the physical file 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)
- 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 scope of the IR 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 pre-registration of contacts. The following three sections will 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 numbers allocated 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. It also eliminates the need within the triage function to analyze information supplied with a tracking number. 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 IR service. Different prefixes might be used for the different services. As external communications have to be considered, too, 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# and CERT-INFO# for incident and feedback tracking numbers respectively.

3.3.1.1 Unique Intra-CSIRT Tracking Numbers

A fundamental requirement for tracking numbers is that they must be unique. Commonly, teams allocate numbers from a predefined range of integers as the basis for their numbering scheme. Within a team's own IR service and preferably across all of their CSIRT services (as tracking numbers can also be used for other services such as vulnerability handling and artifact analysis), 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, difficulties might arise because parties often forget to provide the prefix and refer just to the number. In other words, it should not be possible to have incident number 60 and feedback number 60. The number itself should be sufficient to refer to a unique event. If a team plans to re-use numbers, strong controls must be enforced to ensure that there is enough time between closing a particular event and re-using 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: The DFN-CERT uses numbers between 1 and 65,535. There are no plans to re-use any of these numbers. After 4 years of operation, approximately 600 numbers were used; that implies that at DFN-CERT's current rate of use, it will take 108 years until all the numbers are used. Even if the rate used increases dramatically, there will still be a significant number of years before old numbers will need to be reused or a different set of numbers is adopted.

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: Initially AusCERT used an incident numbering scheme of the form YYMMDDHHMM. This was generated from the date and time, that AusCERT opened the incident.

3.3.1.2 Unique Inter-CSIRT Tracking Numbers

Tracking numbers need to be unique not only within a single CSIRT, but also across other CSIRTs. As multiple CSIRTs may be involved in responding to an incident, they will each use their own individual identifiers to refer to that incident. Potentially there is the possibility that two teams will use the same identifier for different incidents.

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#123 and

DFN-CERT#123 are two separate and unique tracking numbers that may refer to two totally unrelated incidents.

If a team's tools support recognition of various tracking number formats used by different teams, it will further facilitate the triage function. However, all teams are encouraged to reference each of the tracking numbers of other involved teams during their communication with each other to allow efficient identification and processing on both sides.

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 (particularly in the case of incidents) reported. For these reasons, use of some random number generating scheme (while retaining uniqueness) is required.

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 as new 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.

3.3.2 Use of Standard Reporting Forms

The use of standard reporting forms can facilitate the provision of complete and appropriate information being supplied to the team by parties reporting to it. This facilitates the timely identification of new reports to associated activity, routing of information to the right function, and also improves completeness and comprehensibility of initial communications which

makes further processing easier. For most services, useful forms can be designed and used (e.g., vulnerability reporting forms within a vulnerability handling service).

Within the IR 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 made 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 incident function itself, incident reporting forms commonly request (for example, the CERT/CC form [CERT/CC 97a])

- 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
- logs detailing the activity (with associated time-zone information)
- tracking numbers that may have already been assigned (say by a local security team or another CSIRT)

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

```
Mar 2 96 10:34:12 myhost tcpd[52345]. connect REFUSED from cumber.some.where
```

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

Sometimes teams have trouble in convincing people of the need to make a report in the first place. As some prospective reporters feel that a form is cumbersome and not very effective, they are even more reluctant to report an incident if a reporting form is required. A team might choose the risk of losing some initial information in preference to not obtaining a report at all. So 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. 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.

3.3.3 Pre-Registration 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 IR 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. Useful items to pre-register include

- 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 pre-register other information such as 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, AusCERT initially established a constituency registration process as detailed in [Smith 94]. This process includes establishing trusted points of contact, and information disclosure restrictions. AusCERT later changed this process when it became a fee-for-service team.

Example: CERT-NL serves a constituency defined by contract (between an Internet service provider and its customer sites) and therefore can support a registration process for site security contacts and PGP (Pretty Good Privacy) keys. Furthermore because its constituency is of the academic type, CERT-NL is able to uphold a default information disclosure policy.

3.4 Incident Function

The goal of this function is to provide response to computer security incident reports. 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
- **Notification:**
Passing 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 IR 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).

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

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

Table 13: Possible Instantiations of Incident Function Attributes

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 given in this section. As described in the Section 3.3 “Triage Function,” part of the life cycle of an incident may take place within the triage function, where a new incident can be identified, a tracking number assigned to it or where further tracking of an open incident might occur. A new incident can also be identified within the incident function as a result of incorrectly triaged information, information provided to the team under an incorrect tracking number, or new information being discovered.

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 and the incident is closed. This 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 informing all parties impacted 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, if there is nothing more the team can do. As a result, the criteria for closing an incident can vary from team to team.

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

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

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.

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 or have been informed of the closure. This will help to set the appropriate expectation and avoid confusion in the case 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 regular incident correspondence. CERT/CC commonly uses words to the effect that if no further feedback is provided by the correspondent by a specified date then their thread of the incident will be considered closed by CERT/CC.

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. However, if the activity is not considered to be a continuation of the original incident, it might be appropriate to generate a new incident for the activity and issue a new tracking number. Similarly, 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 as 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 impacted will need to be capable of supporting such events. Such technical problems can usually be easily solved, but the human issues are not so trivial to solve and need to be taken into account. People have a habit of us-

ing incident numbers originally assigned to them. 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.

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 profound technical analysis of log-files, malicious code, and incident texture.

Example: Consider an analogy with a hospital emergency unit. The triage function decides which incoming patient goes where. The analysis aspects come next such as blood tests, scans, ECGs, and X-rays. The results of these tests help determine the next actions such as medicine and surgery.

Different types of analysis can be CSIRT services in their own right, separate from the IR service. One could, for example, offer an artifact response service, additional to (or even totally separate from) the IR service. 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 pop up during IR and since artifact analysis often is part of IR to some extent, reference will be made to these topics, 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:
 - log-file analysis
 - analysis of any artifacts left by intruder activities
 - 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 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 response 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., numbers 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.

Obtaining the “bigger picture” is often difficult as different people may be assigned to tackle different types of analysis. Various people within the team will have different pieces of information, resulting from their specific type of analysis. To retain the bigger picture from the information available to the whole team, it is important to institute a process to collate the information to yield the bigger picture. This can be done through team members interacting in regular meetings or ensuring incident supervisors obtain the information from the team.

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 gleaned 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 as unlike personnel, knowledge bases neither quit nor have holidays.

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. Then, the next time this weird directory name is found, it would prompt the team to search for the associated Trojan horse *immediately*.

It is useful and advisable to make the “bigger picture” (appropriately sanitized) available to other teams and possibly law enforcement. 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 developed 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”).

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 mainly on the factors given in Table 14.

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.
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.
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 passing on relevant information to constituents and 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, then in-depth analysis is necessary to understand any new type of 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.

Table 14: Analysis Depth Factors

There is a whole range of actions that a CSIRT can take if it has the time to both analyze events thoroughly and 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 the SATAN/SANTA program⁵ or ISS program⁶ 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 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

⁵ <ftp://ftp.win.tue.nl/pub/security/satan.tar.Z>

⁶ <ftp://coast.cs.purdue.edu/pub/tools/unix/iss/>

need to be analyzed by the team very carefully to avoid any liability for overlooked security weaknesses.

3.4.2.3 Log-File Analysis

Every decent hardware platform, 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 pre-defined (usually undesirable) event takes place, such as a packet flood. Logs are files where events (both harmful and innocent) are recorded. Alarms ring when a specific log entry meets pre-defined “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 (telnet, rlogin, X, etc);
- to what destinations was email 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 IR 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 IR teams 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 use 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.

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. 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 as the log was produced. Such actions are mainly of interest if legal action is possible.

Table 15: Notable Characteristics of Log-Files

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

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.

Example: Don't hand over to the media a log-file containing specific details about a constituent of yours.

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: Ethernet sniffer logs can commonly contain any information, ranging from not sensitive at all to explicit username/password combinations.

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 as 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 an IR team, but you don't want it to leak out. You may ask the parties involved to change 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.

Disclosing Log Extracts

If incident follow-up is undertaken and other constituents and teams are to be informed about the activity and what part of the incident 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.

Example: You don't send one Internet service provider specifics about break-ins at a competitor provider's site, even if both are involved in the same incident.

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, and source code to various programs. Generically we name these *remnant files*. We address scripts, source code, and programs here, naming these samples of potentially malicious code 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. Trojan-horse programs are popular among intruders. These are programs that seem to do everything that 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 do something else (e.g., updating the intruder on what is happening).

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 their original 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 idea (unless you favor reinstalling your systems from scratch after even a minor intrusion). Programs such as Tripwire⁷ for UNIX and some Windows 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 an IR team should analyze malicious code as part of their IR service (or as a separate service) is an important question. Various CSIRTs have differing views. The following examples are from different extremes.

Example: CERT-NL will not (as a rule) analyze malicious code itself but leave this task to its constituents. Only in rare cases will CERT-NL assume this task itself.

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.

No matter who performs the task, proper analysis of malicious code should, to some extent be addressed. 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 so if the flaw that enabled the intrusion has not been removed. Artifact analysis can help do just that, for "inside the artifacts lurks the intruder."

Assuming the IR team takes the responsibility of analyzing malicious code as part of the IR 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 want to return to normal operation as soon as possible, and you don't want to further endan-

⁷ <ftp://coast.cs.purdue.edu/pub/COAST/Tripwire>

ger 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 be performed by a member of the IR team, not by the constituent just sending some files. This means the constituent might grant the team member temporary access to the system involved, or undertake the task themselves using instructions provided by a IR team member.

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 un-isolated environments. This resulted in News "control messages" escaping from their test systems and did what they are supposed to do inside NNTP, the News protocol, 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. Thousands of such messages were received. Had the teams used appropriate isolated laboratory environments, this would not have occurred.

Involve expert groups?

Given the size of the average IR team, 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 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 CERT-NL constituents. The experts benefit by receiving new information first-hand. CERT-NL benefits by obtaining feedback from the experts.

Example: Teams like CERT/CC, AusCERT, DFN-CERT, and others sometimes cooperate when performing analysis (mostly vulnerability analysis; artifact analysis is still a fairly new area of cooperation), with one team taking the lead and the others contributing as "experts."

When to stop?

Criteria should be defined in advance on 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. 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, 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 which 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 software environment is firmly tied to the artifact analysis, so firmly that the 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, IR 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. This requirement indicates how involved 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, an academic environment where skilled technical staff members are available and easily accessible, an 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 96] 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 IR team, be part of the CSIRT, or external to the team (e.g., colleague or vendor teams). In the ideal case, the vulnerability will then be promptly patched and the world 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 hopping on from there to 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 provide access into the center of their web, the telephone company may help to tracking down the intruder's telephone number. (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), 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.

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 89, Shimomura 95].

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

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 IR teams on a need-to-know basis by 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 on the way that the intruder seems to go about his work. This will help the other teams recognize this kind of intrusion when it occurs, and you can expect the other teams to return the favor and inform you.

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 that have caused several CSIRTs a tremendous amount of effort tracking and notifying all references to 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 sub-directory names invented.

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 a result of intruders working together and sharing information. By following one intruder's trail, you might start finding their colleagues' trails too. These trails will probably show both clear differences and remarkable similarities.

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: 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.

3.4.3 Tracking of 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 will provide information required for responding effectively to the incident. It will also provide a historical record of reported activity and 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 incident function or by other functions or services.

The level of detail recorded may vary from team to team based on their needs, the level of IR service provided and the analysis depth undertaken. For every incident, the information detailed in Table 16 should be tracked.

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 CSIRT's incident tracking numbers	Tracking numbers assigned by other teams involved. This facilitates the appropriate coordination of this incident with 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	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 impact 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	Specify the 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.
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 impact the handling of this incident.
Textual description	A free form description that accommodates other information not covered in any other tracking field.

Table 16: Incident Tracking Information

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 re-opening. If the tools support the re-opening of incidents, this is a must. In some cases CSIRTs have reported incidents reopening a year after the original report! Such cases are mainly due to sites failing to regularly review logs for incident activity. Depending on the nature of the team there may be no need to store the whole set of collected incident information any longer, or it might be helpful for historical purposes. This topic will be discussed in more detail in Section 3.8 "Information Handling."

3.5 Announcement Function

As previously stated in Section 3.2 “IR Service Functions Overview,” the announcement function generates information tailored for the constituency in various formats. The purpose of announcements vary from disclosing details of ongoing threats, 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 purpose of this document, the scope of this function will be limited to its direct applicability with the IR service. However, within a CSIRT providing a broader range of services, 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 88] provide information on preventing threats that are generally discovered as the result of incident reports and cannot generally be created without use of the team’s vulnerability service.

Since the formation of the first CSIRT, announcements to teams’ constituencies have been part of a CSIRT’s daily business. As previously discussed, this function is optional as it is not critical to the provision of a basic IR service. The main objective is to disclose information to the constituency to assist them in protecting their systems or looking for possible signs of attack by providing notification of potential, current or recent threats; and suggesting 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

Usually takes the form of a short message, issued when detailed information is unavailable. The purpose is to inform of something that is likely to be important in the near future. Announcing a heads-up has two benefits. First, the recipients may already know more about the issue detailed in the heads-up than the CSIRT. This gives the constituency the opportunity to provide feedback to the team. Second, 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.

Alert

Alerts are short-term notices about critical developments containing information about recent attacks, succeeded break-ins or new vulnerabilities. Examples are more recent CERT Summaries like the one on the “named” problem [CERT/CC 98b] and the more recent CERT Incident Notes and Vulnerability Notes [CERT/CC 98c], [CERT/CC 98d].

Advisory

Mid-term and long-term information about problems and solutions suitable to raise awareness and help avoid incidents. Examples are CERT Advisories [CERT/CC 88].

For Your Information

Mid-term and long-term information, similar to advisories, but shorter and less technical to address a wider audience, including readers new to the topic or area, or interested bystanders (management levels, media). An example is CIAC C-Notes (originally called CIAC Notes) [CIAC 94].

Guideline

A sequence of steps suitable to 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). Examples include the Site Security Handbook [RFC 2196] and CERT Security Improvement Modules [CERT/CC 97c].

Technical Procedure

Guidelines with more technical details addressing an expert audience. Examples of this are the CERT Tech Tips such as the “Problems With The FTP PORT Command” [CERT/CC 98e] Tech Tip.

3.5.2 A Priori Considerations

Having defined a set of announcement types, is only the first step towards a comprehensive announcement process. Several other factors need to be considered and addressed before the first announcement issued. These factors (discussed in this section) range from the criteria that trigger an announcement to how it will be distributed.

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 upon the activities reported to the team or research done by the team, or the source is 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 (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 technical advisory aimed at experienced system administrators. Whereas an equally 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 impact the perceived importance of each individual announcement. Care should be taken to pre-assign objective priorities to each announcement based not only on announcement type, but also on its content (i.e., a handful of broad topics such as denial-of-service attacks or viruses).

3.5.2.4 Clearance of Information in Announcements

According to the team'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.

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 channel safe enough?
- audience addressed: Is the channel adequate to reach this audience?
- speed: Is the channel fast enough?
- cost: Is the expected result worth the money?

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. 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 an internal tracking number allocated. Together, announcement type, content type and audience determine the following important parameters:

- announcement priority
- 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 made at this point include the proposed time schedule, responsibility for the task, and other aspects such as collaboration with other parties (to provide content or improve the quality of the announcement).

3.5.3.2 Prioritization

This phase may be revisited regularly for all announcements under development. Yet unissued announcements are prioritized based on the pre-defined criteria and other (at that time) relevant criteria. Certain types of announcements (by their very nature) might have the high-

est priority if they are time-critical, even when other important announcements are in the queue. Others that simply provide general statistical information may be of the lowest priority. 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 threats addressed and 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. 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: A draft announcement may be sent out to all FIRST teams to seek their review and comment, but is not for further distribution. To protect it against sniffers, drafts and comments are encrypted.

3.5.3.4 Final Preparation

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). However some technical issues still need to be addressed such as generating cryptographic checksums of the announcement itself or items that it references. 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 their response. For a fellow CSIRT this might mean preparation of an announcement of their own. 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 final drafts of their advisories encrypted 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 leaked while others still believed that it would be confidential for some time longer.

Team members providing triage, feedback or other IR functions need to be prepared to handle 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 be protected against tampering, for example, by applying a digital signature.

Example: Every CERT Summary distributed by the CERT/CC has a number of the form CS-YY.XX (where YY is the year in which the summary was issued and XX the number sequentially allocated from 1 for each Summary issued in that year) with authenticity and integrity provided by a digital signature generated with PGP.

3.5.3.5 Distribution

This activity is related to the effort involved to distribute 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, 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 news group comp.security.announce which 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 news group) on its information server⁸.

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 incident response. Ignoring such requests and issues will affect the team's reputation and the attitude of the constituency toward the team. Hence, it is in the interest of all CSIRTs to have a function to provide feedback at some level to such requests. From a management perspective, the type of incoming requests received by the team will 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 provide a better service and at the same time clarify the expectations of the constituency instead of ignoring obvious problems and misconceptions. No request should go unanswered, no matter what the request is.

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 to be arrogant or worse. To avoid this perception, the team should at least provide a statement of the purpose of the team and why no further feedback is possible. Keep in mind, that some of these requests can be the result of "investigative" journalism, which comes in many shapes and sizes.

⁸ <ftp://ftp.cert.org>

Commonly incoming requests 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 help the constituency avoid incidents and improve their 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 or to a specific incident. Whenever possible, the CSIRT should be prepared to deal with 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 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 should not be ignored. Also, not explicitly requested feedback issues like annual reports can be placed in this category.

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

These requests have nothing to do with the IR service provided by the team, but even then, a simple acknowledgment with a reference to some FAQ or policy statement 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 penpal.

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 that are similar to those of incidents. However, it is uncommon for requests to remain open 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. Most teams choose to develop previously prepared documents such as a team Frequently Asked Questions (FAQ) document, which provides details of the IR service provided by the team, and how to access this document and other general 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.

For media requests, depending on the policies of the parent organization, the team might use an existing organizational media office or interact with the media through a team member or associate experienced in dealing with the media. Once the team has established a policy of 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 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 for the team members that describes the procedures related to handling the various types of request is very beneficial, enabling a consistent reaction to all requests. Such an FAQ 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.

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 to the "right" parties (i.e., points of contact) (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").

Example: Say that an incident is in progress. A person 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 news groups to a large number of readers.

3.7.1 Points of Contact

During the course of any incident, contacts are established as necessary. To establish the “right” 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 the information and/or the person authorized to make any necessary decisions.

Therefore, establishing and maintaining good contacts must be an ongoing effort with the intention of building a web-of-trust to suit the needs of the incident response process, with the CSIRT maintaining the web.

For our purposes, contacts can be considered 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 may need to correspond with when handling a specific incident. They include contacts within and external to an organization experiencing an incident. Examples of such contacts are

- upper management
- other departments
- technical 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 pre-determined initial point of contact (POC) that the IR team reports an incident to. 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 provide background information for the team, to help it to fulfill its service, might support the team's operation or can be used for obtaining input from domain experts. The following list provides a starting point for the type of contact that should be considered when generating a contact database. Examples are

- (constituency) site security contacts
- other constituency site contacts (like management)
- sites external to constituency
- Internet service providers
- other CSIRTs
- law enforcement
- 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.

3.7.1.3 Finding Contacts

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

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

To keep the confidence of the constituency, great care must be taken to use the correct contacts. If publicly available contact information can be forged, manipulated, or corrupted (which is a possible threat, but the risk might vary 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, a more daunting challenge than finding contacts is maintaining them. Contact information becomes (partially) obsolete when people leave an organization, are promoted, or just move to another desk 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. Regular (e.g., annual) check-ups, in addition to asking the contacts to relate changes, can help address this problem.

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 `ssc@somesite.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@somesite.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.

3.7.2 Authentication

An important aspect when interacting with others is authenticity. This term usually applies to ensuring 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 within the organization. By “correct,” we mean that the person is authorized to receive, accept, or act on the information. Without such procedures in place, the teams and their constituencies are susceptible to social engineering (discussed below).

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 him 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."

3.7.2.1 Social Engineering

Social engineering is when someone presents a fake identity to trick a person into doing something that they would not normally do if the real identity were known. 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 the 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 understand 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 example given above.

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.

Depending on the technical communication facilities support can be available to prove or verify authenticity. Most well known in relation to today's networks are digital signatures, like 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. Every recipient can check this signature with PGP. This check depends 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 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, these keys are 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 its 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 with 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, they should be very carefully protected. If an attacker could manipulate the database, seemingly "authenticated" data could be entered and the team members would trust it.

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

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 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 using anonymous information is that it makes no difference whether a bomb warning comes from an anonymous caller or not, ultimately, to be safe, you will go and check for the bomb.

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, but also information transmitted via more traditional means like post or couriers which are also vulnerable to attack (or loss).

In the same way as 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).

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 (FIRST has almost 70 members today) 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 the message. The procedures for use and maintenance the keys are distributed among FIRST members [FIRST 98].

In case of telecommunication networks additional black boxes can be applied, as 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 STU (Secure Telecommunication Unit) III devices, which can protect telecommunications. This only applies to the U.S. and Canada. Such devices are controlled equipment that have special handling/reporting procedures and requirements for their usage.

3.7.4 Special Considerations

The following text will present considerations specific for given environments. The objective here is to explain the practical considerations for interactions that have already been intro-

duced. 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 beneficial approaches and pitfalls to avoid. Although the examples provided include a wide range of possible partners, others might exist. But we believe we've covered the most important ones to consider. Any others that you may identify can likely be treated similarly to one of the categories below or will be similar to 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. For interactions one additional consideration 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.

As 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.

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 resource (e.g., vacations or conferences), prioritization will become even more important to distribute the activities among the available staff. Documented and public 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 constitu-

ency 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 IR Teams

Incidents that require no external interactions with other parties are rare in today's networked environment, as they only arise 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 they have been in the CSIRT business for a long time or have particular technical expertise.

Support might be provided in one of the forms described in Table 17.

Support Type	Description
Education	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 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	As 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 17: Possible Inter-Team Support Types

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 that teams are able or willing to exchange information and to cooperate on confidential issues depends on any existing 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.

Issue	Description
Confidentiality/Secrecy	As 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.
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	As the information may be distributed to the public at some future point in time, disclosure restrictions should be stated. 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 means 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	As 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.

Table 18: Considerations for Information Sharing

If two teams want to initiate a cooperative relationship, it is difficult to establish the necessary foundation of trust. One of the most important steps towards such a relationship is 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.

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. As 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 the 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 operational procedures.

Mandatory Information

The issue of dealing with incoming information was described above. 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; possible reasons include 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.

Not all inter-team relationships are at the peer-to-peer level. Some teams elect to participate in a voluntary hierarchy; less frequently teams exist within a mandatory hierarchy. Even if teams interact as peers for one activity, does not preclude them interacting in a hierarchy on other occasions.

Example: In 1997 a coordination center for CSIRTs within Europe was established. EuroCERT provides more information about its task on the World Wide Web⁹. EuroCERT acts as the main point of contact for incidents involving European CSIRTs and is used as an interface between European teams and other teams across the world.

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 to take place. Someone needs to take the lead otherwise a 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 decided on a case-by-case basis. Usually 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.

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

⁹ <http://www.eurocert.net/>

to CERT-NL. This is even true if the site involved is not within the formal constituency of CERT-NL, the customers of the Internet service provider SURFnet.

Whenever 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 maybe 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, there is no way that a doctor or nurse can ignore you if you ask for help (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: Every now and then, the CERT/CC receives anonymous calls announcing a new "Internet Worm." To date, these calls have proved false. On the other hand, if DFN-CERT would call the CERT/CC and provide evidence of a potential worm, this report would receive considerable attention.

When setting up a team and allocating resources and responsibilities, it is important to understand that requests that originate from outside of the declared constituency 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. 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 does not like resources spent on "outside" 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 40% of all European nations have a funded (not volunteer based) CSIRT. Usually, such CSIRTs were established for research networks, so depending on their policies, may or may not they handle incidents involving commercial sites in their countries.

Therefore each team should set clear expectations and establish understandable and enforceable policies to deal with external requests. Whenever there is another team responsible for the reporting site, that team should be notified about the report. If the reporter requests complete confidence they should be encouraged to contact the responsible CSIRT directly. As the existence their report together with the request for confidence by itself is valuable information to the responsible team, a team might choose to inform the responsible team about the report and request without revealing any details on the origin of the request. With this knowledge the responsible team might try to understand why the original reporter opted for confidence and it may result in the team improving their service or changing some of their procedures.

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 the team's services, from incident response to consultancy 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 consider serving incidents involving those units with a high priority and immediately escalate the incident to the CSIRT management's attention.

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 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 legally required to with a legal order, a team 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 law enforcement 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.

One other issue in cooperating with law enforcement agencies is the exchange of statistics and addressing the need of raising the awareness within the larger community. As CSIRTs will have first hand knowledge not only about computer crimes but incidents not considered as crimes, they can substantially enhance the statistics of law enforcement agencies to build the bigger picture.

3.7.4.6 Media

As 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. Consequently, the media often try 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, these individuals 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 if specific information relates to a site, a product, a new vulnerability, an ongoing attack, or a password.

Firstly, information must be collected and incorporated. Every piece of information must be stored and protected throughout the time it is held by the team. Tagging the information according to its type and sensitivity will facilitate its further handling. Before the information is processed further it must be prioritized to ensure that the most important information is worked on first. Finally, the information itself or an aggregation of multiple information pieces is 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 team handles will be sent to it directly, there is also a need to collect information, such as proactively searching for information on the Web or retrieving information from other sources.

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:

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. Exchange with other parties: As 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.")

As the information available is continually changing information collection and other related policies and procedures must be reviewed and verified frequently to take advantage of all possible information sources.

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 re-

lated to events, vulnerabilities and potential interesting discussion threads, the reporting users must be provided with appropriate support, such as reporting forms. This in turn requires a point of contact for reports be set up allowing for more types of information to be reported to the team.

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

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 impact the further processing of the information provided.

Example: If DFN-CERT reports a SATAN scan over large IP address ranges to CERT-NL, it will get a high priority 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 type of caller makes the situation more difficult.

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

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 impact the subsequent processing of the information.

Example: Concerned constituents who have received hoax virus reports from other parties often send the report 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.¹⁰ 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 impacts how the information is further handled e.g., storage, dissemination distribution and disposal. This is the approach taken by UNI-CERT in its guidelines [UNI-CERT 96]. Without differentiation, all information must be protected to the highest level and similarly disposed of.

Even if no explicit categorization is used, perceptions exist on the kind and importance of each piece of information. As these perceptions may differ amongst individuals, clear and concise procedures (as explained in Section 4.2.2 “Information Categorization Policy”) must be available to guide categorization.

Many CSIRTs handle contact information differently than other information, and as a result they are subject to specific policies and procedures. Contacts (people, organizations) are usually sheltered from exposure, even to trusted fellow teams. Therefore specific statements for sanitizing might be included and contact information might be placed in a category of its own.

Categorization is often based on the information itself. Sometimes the categorization is obtained following a dialog interactively with the information provider. At other times the information provider also specifies a category for the information.

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 classified 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 97a]), 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.

¹⁰ 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.

3.8.4 Information Storage

Whenever information is stored (whether it is recorded, written or stored in a computer system) security is of major importance. Without security, a team cannot pretend 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 has a value 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 quantity (e.g., one or two email messages) of information is disclosed due to inappropriate storage and protection. However, exposure of a small quantity of collected information (e.g., the unsanitized summary of a single incident) will greatly increase impact to the CSIRT's reputation.

CSIRTs are attractive sites for intruders. Clearly, putting a CSIRT out of business by discrediting it one possible motive for an intruder to gain unauthorized access to a CSIRT's data. However, another motive to consider is the information an intruder can learn from access to the data. An intruder might easily be able to determine to what extent their activities have been identified and reported to a CSIRT, identify information about vulnerable sites, or 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
- incidents
- vulnerabilities and patches
- exploits
- artifacts

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 disposal of in a consistent way.

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: To identify if a given site was compromised, a captured password file might be provided to the system administrator for verification. If incomplete information exists about the origin of a password file, it can be provided to a likely site so they can check to see if it really belongs to them. The CSIRT sends the site a copy of the file from which all encrypted passwords are removed. 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 (in this case, the encrypted passwords).

The storage of user and organizational 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 affects every party for which information is stored.

Example: If there is a legal requirement to provide specific information about one 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 media. Knowledge of such instances might result in reluctance of constituents to report future problems to the CSIRT.

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. This is particularly difficult in the case of backups, because the whole purpose of a backup scheme is ensure 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: Two different backup schemes are used: 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 is not easy.

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 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.

Continuous re-prioritization of incidents must occur. Whenever new information on a given incident comes in, its priority might change. As a change of priority also affects the reporter and impacted sites, these parties will also need to be informed accordingly. This is most important whenever incidents are downgraded to a lower priority.

As 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 can not be handled, the reporter must be informed. Without such statements, 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

Depending on the influence of a target, a value is assigned. Targets within the constituency boundary can be viewed as more important than targets outside the constituency, as the constituency "pays" the team. 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 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. This is hard even after the fact and is even more difficult to predict 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 team 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. Money might be chosen for this purpose. However, it is very difficult to calculate damage financially for some incidents. For example, it is difficult to estimate the amount of money that an organization might lose due to public knowledge about an intrusion. As a result, this approach is of limited use in prioritizing incidents.

Example: Hospitals and emergency teams have similar prioritization schemes:

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

3.8.6.3 By Incident Type

Known incident types are ordered, depending on their overall (potential) technical impact, such as denial of service or 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 another.

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 5 hosts at sites where the staff is experienced responding to incidents. One is a denial of service attack at a hospital affecting 2 hosts that hold a medical records database and laboratory test result data. The other two incidents involve hundreds of other hosts as the attacker is running attack scripts on other sites from compromised hosts.

The issue here 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 as they have experienced personnel available? Are other resources available to 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?

3.8.6.4 Feedback Request Prioritization

Generally requests for feedback can be handled differently from incidents. 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 skills). One of form of feedback request prioritization that is occasionally used is that 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 similar, escalation is concerned further raising the importance of an activity regardless of it’s priority. Escalation invariably requires at least one level of management to become involved for decision-making purposes or other activities that require their authority. 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 should be defined in advance in preparation for use. There is a continuous need to review the criteria and to adapt to changing needs and new developments, such as new attack styles and incident types.

By its very nature, incident escalation is driven by similar issues as those involved in the incident prioritization. However escalation criteria can be applied to the incident response service as a whole as well as to a given incident.

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 the incident handler being unable to address one or more aspects of the incident appropriately and they either need additional support, management oversight or to offload other work in order to appropriately handle the escalated incident. As an incident evolves and new information has come to light it may be apparent that the person to whom it is assigned does not have the technical expertise required to handle it appropriately.

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 a thousand 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.

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 news group. The report identifies connections that were made from the compromised system at their site to 1000 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 site is unable to identify the legitimate connections from the unauthorized ones. Over fifty of the remote systems listed fall within your constituency. The incident should immediately be escalated to management due to the possibility of media attention related to the activity.

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

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 being either technically or procedurally 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 it, the incident may be escalated allowing additional techniques to be used to attempt 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.

3.8.7.2 Multiple Incident Escalation

From an incident response service perspective, escalation criteria must also take into account additional factors including the overall workload a team is experiencing, the need to meet its mission, retaining the bigger picture and additional resources available to the team.

There are times when a team has more incidents than it can possibly respond or, or it is unable to meet its published response guidelines. These situations sometimes arise gradually as the rate of incident reports increase. 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) and reducing the level of service provided.

Teams are often faced with 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 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 impacts. In particular, it will adversely impact 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.

One major benefit a team coordinating the response to incidents is that it is able to develop, see and interpret the “bigger picture,” as discussed in Section 3.4.2.1. The 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. So losing your grasp on the bigger picture is especially serious during escalation, a time when the bigger picture is critical to both the team and its constituency.

While all available incident information is of interest when gathering the “bigger picture,” not every incident is of equal importance. When escalation criteria are implemented it is common for less analysis to be undertaken and so less information will be available to form the “bigger picture.” Wherever possible retain the necessary level of analysis on incidents critical to maintaining the bigger picture.

At times, such as during escalation, when a team is not able to cope with its workload, additional resources might be available for the team to draw on. These resources should be planned for in advance as when a crisis strikes you need to target all your available resources at the workload and not into arranging for and training additional staff. 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 group(s) chosen, the help of these people might be easier or more difficult to arrange.

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 the team, 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. As 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 practical issues are discussed here.

Different types of disclosure will need to be handled by different teams. Here are some examples:

- to other teams about a new vulnerability
- to other teams with regard to an incident
- to sites with regard to an incident
- to management for budget purposes
- to management for reporting
- to risk management group for improvements
- to funding body or shareholders for justification
- to law enforcement
- to governmental organizations
- to everybody who has a vested interest

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 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 German research network, the DFN-CERT will inform the whole constituency rather than only the known victims. Usually the source of such an attack isn't 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 isn't 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 preventative and reactive measures) to the constituency and let the sites handle it themselves.

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.

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 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, the Australian CSIRT, initially implemented a registration process where sites were asked if AusCERT could pass their contact information to other CSIRTs whenever the sites were involved in an incident. If a site does not wish to be contacted in regard to a specific incident, the site can inform AusCERT, and the contact information will not be passed on to other CSIRTs.

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 teams provide forms for submitting information to them. Usually forms make it easier for both the reporter and the team, 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 provide without the form.

Example: The CERT/CC makes its incident and vulnerability reporting forms available on its anonymous FTP server¹¹ [CERT/CC 97a, CERT/CC 96].

Sometimes teams or organizations place specific requirements on their constituency to fill out forms or to provide a defined set of information. Depending on the 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

- build up the bigger picture
- provide the funding body with additional background information
- provide a better service to the constituency
- raise awareness with pragmatic projections

¹¹ <ftp://ftp.cert.org/pub>

By definition, to make use of the collected information is vital to fulfill the mission. It might also become part of the interaction with others; for example, the disclosure of CSIRT statistics to FIRST is discussed as a possible requirement with each prospective member.

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

4 Team Operations

In Chapter 2, we discussed the main functions that an IR 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 IR 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, fire alarms and burglar-alarms and procedures. Therefore this chapter discusses operations.

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 the topics here are not exclusive to IR services. Therefore it is not surprising that some aspects of these issues have already been covered in Chapter 3. Where appropriate we will refer back to that section instead of repeating ourselves. However, in this section we give a more practical approach than was possible on 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 IR 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), and backup and all-hands-on-deck arrangements.

When considering work shifts, remember that after 2 hours of routine work, you definitely need a break; but that after only 1 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, 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 have the technology available to them to initiate communications to the constituency and other parties as required. Implementation used 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, you are not likely to have a working cellular phone. Be aware that constituents will be displeased if they have to wait for more than 4 consecutive rings when they try to access your service by telephone; and that they will be very displeased if nobody gets back to them within 15 minutes if theirs is an urgent case. A voice box might be useful to provide a first acknowledgment and further information 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 the CSIRT team member can be reached at.

4.1.3 Electronic Mail (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 IR service will have additional demands, like good filtering capabilities, advanced search facilities and automatic response tools.

Usually IR teams 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 plenty 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, this might change. 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, users 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.

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 have become obsolete (which it hasn't yet). 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 IR teams. Integration with email tools, the Web, and the telephone system (and pagers) is necessary to collect all incoming information and to interconnect events with each other. Technology is improving rapidly in this area.

4.1.5 World Wide Web Information System

The World Wide Web (WWW) is currently the hottest medium in use for retrieving information. Certainly no commercial team could do without it, and other teams should not lightly discard the reasons for the Web's popularity. 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 too, and that Web-based information can link to it. Clearly Web servers and any public information server for a IR team (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 on the outside of the firewall 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 the inside and download it to the Web server on a regular basis, like every night. Additional checks based on cryptographic checksums (like 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 because of security, you will require ownership of IP address space dedicated to the team. By using a Classless Inter-Domain Routing (CIDR) block of the Internet IP address space, it is of no consequence if only you and you alone use these numbers and not e.g., also some other part of your parent

organization. Alternatively, a team might choose to use some private address space (e.g., 10.0.0.0) and either address translation or a firewall for all external connections.

The 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 `my-csirt.some-org.tld` or `my-csirt.tld` (tld stands for top-level domain, like `.com`, `.edu`, `.org`, `.nl`, `.de`, or `.uk`).

4.1.7 Network and Host Security

An IR 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 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. Most of the time it is unnecessary to connect the testbed to other networks at all. Only when data transfer is necessary, a temporary connection can be established, but be careful that it is truly temporary. 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.

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 possibly other defensive tools should help identifying and preventing intrusion attempts. But even the security of home systems, 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 carefully considered, as they impact 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. But as it is most likely that service and quality specifics will affect the content of fundamental policies, the discussion below is generic, with 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 the intent of the organization’s mission statement and the organization’s character. A code-of-conduct applies to all staff members at all levels within the organization; it is an attitude, and attitudes should be classless. It provides basic direction that impacts any interaction and how to react in certain situations. It 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. Having worked long enough in an organization, the code-of-conduct should become natural professional behavior; the code-of-conduct should fit the organization’s stance.

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 definitely should not be the case. 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 constituency and fellow teams? This will also help form the sort of basic understanding needed for collaboration between teams.

A very simple code-of-conduct example (which has to be conducted in accordance with the CSIRT policies and procedures) follows:

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 [CERT/CC 98f].

4.2.2 Information Categorization Policy

CSIRTs must have a policy on information categorization. Without it 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.

A slightly more elaborate scheme could define the categories "internal" for use only within the team and on a need-to-know basis also for interactions with fellow teams, "external partner" for interaction with the constituency and fellow teams, and finally "external public" for public information. This is exactly the approach taken by CERT-NL as detailed in their operational framework [CERT-NL 92].

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, 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 classification. A pragmatic way of setting an initial scheme in place might be to ask the team members separately to classify some documents, then adopt a scheme based on their names and classification.

Military teams 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.

The category selected will impact the way the information is handled (e.g., storage, disclosure, disposal). As a result, policies and procedures must be developed for each category. Then, regardless of the real content of the information, a consistent set of policies and procedures apply 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 will differ.

Example: The default for categorization for contact information may be “internal” and differ from the default of “public” for publicly released of advisories issued by other CSIRTs.

Sometimes it is not clear what category information should be placed in as it might be considered as a candidate for more than one category. Within the CSIRT environment the category chosen is normally that which ensures that the information has greatest protection. If at later time new details become available to indicate the information has been incorrectly categorized, it can easily be re-categorized.

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, as there will be a reluctance 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 and when they can say to whom 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 where 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 the 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 their 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. Under what circumstances must a team pass sensitive (even contact) information to law enforcement or a court? CSIRTs do not have a similar legal status regarding client confidentiality as doctors or lawyers do.

Example: Consider the example where CERT-NL provides the CERT/CC with information about a security 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 to forward to the U.S. site and will indicate if 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 con-

tact 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.

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 a site that their machines may have been compromised by a news-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. However, if an incident involves break-ins through Telnet, 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 to what information is 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 of the information will decide how disclosure will take place and what pieces of the information will be disclosed.

Example: Consider a large-scale incident, with intrusions involving hundreds of hosts all over the world. As a result, 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 to them or their constituency. You might

warn law enforcement by telling them size and spread of the event, plus generic exploit details. You may tell the media about size and spread of the activity, and a warning and some comforting words. However, to trusted experts you might give all the gory details (by sanitizing the information to make it anonymous) so that they might learn more about exploits, trends, and signatures. To victims, you pass the relevant log entries (sanitized to be anonymous) to enable them to check their own logs, together with guidelines on how to protect against future attacks of this kind.

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 this information will only be used on behalf of the teams' constituencies, and it is not to be spread beyond.

One approach helpful for others is to place a label on disseminated information clearly stating the expected use of it (for example: "For internal use within the CERT/CC 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 words, may also cause the constituent to file a complaint or lawsuit. This kind of behavior is less likely to strike a team who has no authority over, or contract with, its constituency, such as the CERT/CC.

4.2.4 Media Policy

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, like CSIRTs and related teams do, it is advisable to have the media spokesperson external to the team. Then the individual spokesperson has little or no access to sensitive data, as they only know as much as they need to know to fulfill their function to fill in 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 law suits. 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 [McGillen 93].

Establishing 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. The current Internet rage is helpful. Many publications have good people on these jobs nowadays; 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.
- Provide 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 finds you prepared to answer questions that might otherwise unnerve 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 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. Additionally, this list can be used to send a heads-up or invite them for a detailed briefing to alert them to an upcoming event.

4.2.5 Security Policy

Every self-respecting organization nowadays 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 98, 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, additionally they are of great interest for the professional intruder as they might find information on companies who 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 policy of this type 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 has 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. The error noted, managers and staff member alike should put aside their sentiments for the moment and work *together* on containment of 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 staff member and manager (plus 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 a hot week, pressure and workload are high, and the week is coming to an end. A staff member incorrectly cut-and-pastes the 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 also 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, more coffee).

If mistakes by any staff member start becoming regular occurrences, then additional steps outside of the human error policy will be necessary.¹²

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 aspects of operations, including three aspects that 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 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

¹² In risk management, there are known principles to deal with such situations. This includes "separation-of-duty" and "four-eyes-principle."

3.8.6 “Prioritization Criteria.” What it means for an IR service is that you let a sniffer log that is 2 weeks old wait a bit if at the same time all your attention is drawn onto an acute case of intrusion. Without a pre-defined 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 effects the team’s service. When you have a lot of work at hand, it is helpful to make notes of 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 like 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 and using these in the course of work is the oldest form of workflow management, but still workable, even if software packages exist. Workflow management is treated below in more detail.

Academic CSIRTs are particularly vulnerable to incidental lack of time, which is caused by an 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 pre-assigned 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 spreading 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 actually coming into action and performing specific actions. Simply by spreading the business hours, this can be avoided. If personnel are on a business trip, they might be available to help out or their specific expertise may be 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 just being one of them. Off-site coverage is discussed separately below as 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 transitions between shifts pose 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 out-of-hours and business-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 in a telephone call. Additional explanations are necessary in these cases. In the second case, 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 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.

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.

4.3.1.2 Medium-Term Issues

For the medium term, the useful thing to do to help continuity is getting 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. Both 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 IR business (and whenever there is a lack of funding). Good work and holiday conditions will help to ease the burden on the individual. Job rotation will help too; the latter will also help against staff boredom, which is also a form of staff burnout. Boredom is not unusual in the IR business, 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 also have positive impacts for the team as its 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 changing trends, somebody leaving, or falling ill. Section 4.5 "Staff Issues" covers this issue 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.

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 of it. 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.

IR continuity problems arise as IR teams have to deal with a lot of problems over longer periods of time, with continually 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 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 should also be available. This also allows the hand-off of ongoing incidents by the team members new to 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 this incident with this number 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, all electronic files are numbered and stored too, the email and files usually 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. This does not even include the mention of tracking records of the incident, assuming that one of these four archives is used for that. Several teams actually use a fifth archive for that (some sort of database).

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 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.

Getting rid of paperwork is not that difficult: Scanning techniques are quite sophisticated nowadays and relatively inexpensive. Incorporating these into everyday IR practice would be a good thing to do, though 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 archived for requisite length of time.

Perhaps the best current practice for integrating email, files, and access to the Web is indeed using a Web browser. Converting email archives to the Web is possible most of the time (certainly within UNIX environments). Accessing files from a browser is easy. Search functions and indexing are also easily implemented. The Web solution is still one that you have to devise yourself.

One further degree of complexity now has to be added: how to properly keep the history of an incident. In the above light 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, all by hand, possibly helped by some home-brew tool. 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 within 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!). Using the IR terminology, a new incident is stored in the database; and from then on, all related events are logged. Every incident has a status field ranging from “new” to “closed.” Lack of status change may trigger alarms. This is just the core functionality; many additional possibilities exist.

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 still rudimentary.

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 policies, 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 while off-site coverage is handled in the next section.

4.3.3.1 Hotline Coverage

There are different choices on how to implement the coverage of a hotline. The most important 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 that 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 a cellular phone. Last but not least, the person on duty might sleep in the office.

If hotline calls are relayed through other staff members or external parties, they can have a list of home telephone numbers of each team member, or a hotline number can be provided to call or page the staff member.

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 easy, as other team members might be able to help; but what happens out-of-hours? Thought should be spent on this issue. For most of the teams, it might be a good approach to consider having at least one other team member available as a backup on short notice (or a backup might be chosen in a crisis situation by finding out who is available). As this will impact 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, 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 visa-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 one 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 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 and preferably should not be 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. However, 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 IR work they are requested to conduct off-site. The priorities associated with the tasks must be clear and determined in advance. These priorities determine what tasks take precedence and what can be left until the next day back in 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 clearly 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 your security management:

1. confidentiality: to get what you are allowed to get and nothing more
2. availability: to get what you want when you want it
3. integrity: to be sure that information stays the way it was intended
4. authenticity: to know for sure where the information is from
5. exclusivity: to assure that only the intended recipients can use the information
6. privacy: to guarantee that the interests of persons and organizations are protected
7. obligation: to guarantee that the due diligence requirements were fulfilled

4.4.1.1 Use of Encryption and Signing Applications

The use of encryption tools is unavoidable for any CSIRT. Within the team, they offer good possibilities for securing data on computer systems and during the 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 DES and PGP 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 transferred as encrypted using email to and from constituents to keep private the sensitive information about victims and the systems involved. With regard to internal encryption, one can choose proprietary standards and several good possibilities exist but will not be discussed in more detail here. When dealing with the outside world, you have to opt for (defacto) standards such as DES (now superseded by Triple-DES) and PGP. In some communities, PEM (like PGP, mainly used for encrypting and signing email) is still in use. S/MIME, also for email, may also become a defacto standard, judged by the support it receives from Microsoft,

Netscape, and the rest. 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).

Because of the already mentioned export controls, it is sometimes not possible to use all of the available applications in every country. For example, at present, you cannot find a secure (i.e., using strong cryptography) S/MIME application outside the U.S. Every legally available version in Europe for encryption of messages will use only 40 bit keys (some vendors now can export 56 bit keys under specific circumstances). Although internally more bits are used, the remaining ones are fixed and well-known to government departments. This makes it especially difficult for international operating companies that want to use the same application everywhere. Even U.S. companies and individuals might use weak cryptography, if they obtained the program (especially Web browsers) from a public server in the U.S. All users and especially CSIRTs (as they specialize in security issues) are advised to adhere to the familiar phrase: Weak cryptography is no cryptography. So be aware and get the best cryptography that is available to you.

The good news is that serviceable programs such as PGP and DES 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. PGP Version 5.x has brought more user friendliness and better integration with email clients. Other activities also support the development of a new standard for PGP within the IETF [IETF 97].

4.4.1.2 Key Management and Certification

Use of cryptography introduces a key management and certification problem. PGP, PEM and S/MIME use asymmetric encryption (also known as private key encryption) for providing strong authentication; avoiding the weaknesses of symmetric (also known as single or secret) key encryption schemes as the secret key must be known to all communication partners. Hence it is impossible to provide authenticity of the origin and destination. 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 your password.

Within 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, 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, for 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 pass-phrase,

which is evidently not very secure. For this reason some people carry their private key on a floppy, though one might wonder about the security that provides.

Beware, though; using strong cryptography without common sense is no cryptography at all. If you use a 3-letter pass-phrase 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 pass-phrases to unlock the private keys on computer disks or within programs and scripts.

The 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 PGP-key-signing-parties are so common and necessary nowadays, at which people try to check each other's public keys and check passports to prove the identity of others. PEM introduced specific entities, called Certification Authorities (CA) and Policy Certification Authority (PCA) [RFC 1422], to carry out such tests after which signatures are added without the need for end users to do it on their own. S/MIME initially relied mainly on a similar approach. Trusted third parties (TTPs) like Verisign Inc. will sell you a key pair (public and private key), though one might wonder about the security that provides. From a user's perspective, it is absolutely necessary to be the only person with access to its private key. If you buy a key from a TTP, this TTP will also sign your key, thus making it more trustworthy for the community at large. None of these systems currently provide a digital identity for network citizens worldwide to reliably compare personal ID systems such as passports. Caution should be exercised when relying on these entities as they rely on proprietary policies of PCAs or TTPs.

With PGP where users can sign 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 verify the fingerprint, users have to rely on the web-of-trust, which means another user has certified that he verified the binding between a key and its user.

Example: If Moira signs Don's public key, and Peter wants to send secure email to Don, Peter will see Moira's signature on Don's public key. (That is a bad example since Peter also knows Don, but the same concept would apply for Moira's new colleague, for example, whom Peter hasn't met before.) As Peter 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), Peter might then also trust Don's key without checking this with Don in person. That is the idea of the user certification of PGP, which builds a web-of-trust within smaller user communities.

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. 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 95]. 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, but tools such as swatch¹³ and logsurfer¹⁴ allow for online recognition of suspicious log-file entries).

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, surplus workstations, and hot spares). To protect against interrupted power supplies, backup arrangements should be made.

Do not forget to have extra backup tapes in another building; think of fire, for instance. But the other place may be less secure than your own, so 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 can not block each attack.

Any testbed e.g., for viruses, artifacts, programs with unknown behavior etc., must be separated from the operational network of a CSIRT, to ensure the availability and integrity of the mission critical computer, communication and network systems. Contamination or denial-of-service events will badly impact the ability of the team to perform its function and will ruin the standing of the team in the public. This is even more true if a virus escapes or if an attack involves other systems of the Internet for example.

Example: 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. This could have been prevented had the people performing the tests used an isolated testbed or ensured that they had secured their testbed properly.

¹³ <http://www.ja.net/CERT/Software/SWATCH/>

¹⁴ <ftp://ftp.cert.dfn.de/pub/tools/audit/logsurfer>

4.4.1.4 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 the local systems holding the 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 solutions. The first one is dialing into the network. This is in itself relatively secure, especially when the access procedure uses strong authentication like 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 telephones (in the U.S., STU III). In addition to this, physical protection against loss or theft of a mobile computer, 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. The neatest solution is application-level encryption, but this is often not feasible or not good enough. Until recently, U.S. export controls allowed only for 40-bit keys for encryption, even if the application or protocol used was built for 128 or more bits.

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) and AltaVista tunnels 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.

But to protect 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 impact 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 also. It might be easier to restrict the necessity to a minimum or disallow the handling of specific categories which are especially sensitive.

4.4.1.5 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 or printers, or even hotlines inside the "safe" environment. Telephone conversations should not have the possibility of being overheard by other persons like guests.

Concern should be raised over wiring schemes in the building, location of hubs etc., which might be easy eavesdropping targets. Distrust all other public communication mechanisms, especially mobile ones, which are possible to eavesdrop (although this is not a distinctive physical security problem). Consider using encryption for connections in question. Be aware that beside technical means, information can also be revealed to visitors and guests, not only as they are possible seekers of information but also in the normal course of small talk or just being in the room when incident-related information is discussed. 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%.

Consider 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 these situations into account.

4.4.1.6 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. Clearly some definition should exist on when to enter the “disaster mode” (and on 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. This should include gathering sensitive and critical information, including hardcopy documents and electronic media.

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.

4.4.1.7 Handling Internal Security Incidents

Organizations like to keep quiet about incidents internally. If nothing is said within the security policy nothing is said on this topic, the “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 people may be aware of the activity and may disclose the information publicly. Certainly the perpetrator knows of the attack and intruders like to brag and publicize their activities, often supplying proof to support their claims. So attacks against the CSIRT cannot be ignored. This is another

instance where CSIRTs need to practice what they preach. CSIRTs must prepare for and address such incidents, not just for the obvious reason of containment of the overall incident, but also because if they try to hide it and someone subsequently exposes the activity, then the team's reputation may be damaged irreparably.

This holds even more 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 criteria 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 to ensure that an individual is a competent and effective team member, as they are constantly communicating with their team, 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. Hence attention to an individual's interpersonal skills are extremely important.

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
- willingness to own up to one's own mistakes
- 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 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, FTP, TELNET)
- basic security principles
- risks and threats to computer and networks
- security vulnerabilities/weaknesses and related attacks (IP spoofing, Internet sniffer and computer viruses)
- network security issues (firewalls or virtual private networks)

- encryption technologies, digital signatures, cryptographic hash algorithms
- 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) and computer systems (UNIX, Windows NT, 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 up and maintain strong skills, and support the continuous improvement to reflect changes in constituency, technology, 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 range of 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. But this will lead to other problems whenever one person leaves the team. Although 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 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, social skills and team fit
- candidate technical presentation
- reference checks including criminal records

Depending on specific organizational needs of the parent organization or the constituency, more specific requirements such as security clearances and/or background checks may also be necessary.

The overall hiring process should be designed to ensure that the candidate has the suitable interpersonal skills for the position and has or can be trained in the necessary technical skills. As many 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 to the interview process is kept to a minimum, yet is used to the maximum effect [Crabb 96, Fithen 96].

A pre-interview document check and telephone screening with the candidate can 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 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.

The requirement to have a candidate give a technical presentation provides the CSIRT with an opportunity to understand other technical and interpersonal qualities of the candidate. The team can understand how much common sense the candidate has and how the candidate

cope under stress. They can quantify other attributes such as general presentation skills, attention to detail, technical accuracy, and ability to answer questions on the fly.

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 and 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 CSIRT members (such as system administrators). Exit procedures might include

- change of 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, different exit procedures might apply since there are underlying reasons for the decision that affect the trust placed in the employee.

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 staff members for personal development and overall team benefit; and keeping the overall CSIRT skillset 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 at the same time, that 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 important to maintain a steady awareness as to why it is important to follow the established policies and procedures, as well as to exercise situations where personnel must apply their own common sense if a gap within 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 incident response 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 of the form of background reading and then learning by experience. This holds true for incident handling, but there is no formal education, nearly no literature, and written material comes in the form of workshop reports or presentation slides. More important for preparation is the review and study of internal documents, like the policies and procedures, and case studies or past incident summaries collected for this purpose. As the written material through which people can learn to handle incidents is limited, on-the-job-training becomes a necessity.

Even the most experienced staff members associate some level of stress with 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. Initially it is strongly encouraged 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 onto small-scale incidents. In each area, the approach could take the form of the new staff member simply 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. So 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 shows the new member how policies and procedures effect the handling process [Longstaff 93a, Smith 96].

Training on-the-job can also be used for existing team members who need to be trained to maintain their knowledge base. This is vital for the team, as the technical world is changing rapidly. In addition, attending conferences, work in appropriate international task forces and working groups provide knowledge not only to the team member involved, but to the team as a whole.

4.5.5 Retaining Staff

As discussed in the introduction of this document, experienced CSIRT staff are in short supply and expensive to hire. So having invested in the time and resource 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 are not exceptions) where the constant pressures and stress from daily (and often nightly if a 24-hour service is offered) incident response tasks become a burden and intrude into the private life. Staff 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% of any individual's effort dedicated to IR service

- attendance at technical conferences/workshops/tutorials (such as the FIRST Conference¹⁵)
- 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 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 field of incident response is unpredictable, 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 for reaching 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 within escalation policies and procedures.

These additional staffing resources might be drawn from

- other areas of the overall security team
- other groups within the CSIRT parent organization
- other groups within the CSIRT's constituency
- other CSIRT organizations
- external experts

When considering staff to serve in this role, the same hiring principles should apply for them as for any CSIRT member. To ensure that such arrangements are effective, procedures and arrangements should be established in advance to allow for them to be enacted as quickly as possible:

- agreed-on criteria for calling in their participation
- non-disclosure agreements

¹⁵ <http://www.first.org/conference/>

- up-to-date contact information
- prior agreements from management
- procedures to establish secure communications
- initial and regular training

It is essential to provide extension staff the opportunity to go through some on-the-job-training before she/he is allowed to participate in the actual incident handling process. This will give all personnel the chance to socialize with each other and to get familiar with the way policies and procedures are executed through the day.

5 Closing Remarks

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 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 the security in general, the needs of each CSIRT are unique and the CSIRT environment is dynamic. There is no chance of long-term stability as technology, constituency base, and the intruder community can change any time. To ensure successful operation, a CSIRT must have the ability to adapt to 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 computer security incident response, we have found it rewarding work (despite being hard, sometimes frustrating and demanding work). The rewards come from believing in the work, the chance to interact with other dedicated members of teams from around the world, the willingness of the CSIRT community to share lessons learned and support each other, and 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 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 computer security incident response 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 continued IR development and maturity. If not, we hope that it can at least be a starting point for further refinements, improvements, and initiatives to develop better documents, policies, or even standards. We will be happy to be involved with or contribute to other efforts of this nature.

We cannot overemphasize the need and importance for the exchange of ideas, experiences, procedures, or documents. As every CSIRT environment is a little different from others, everyone has something to share, even if this fact isn't immediately evident to those involved. Instead of waiting for others to come forward, you should examine what you can share and then find the right way to actually do it!

If you have comments on this document, if you want to share your opinions, or if you have suggested additions 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 the following address:

`csirt-handbook@cert.org`

About the Authors

The authors of this handbook have extensive experience in the formation, documentation, and operation of their own team's incident response 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 impacts.

Moira J. West-Brown <mjw@cert.org>

Moira J. West-Brown is a senior member of the technical staff within the CERT Coordination Center (CERT/CC) based at the Software Engineering Institute (SEI). She leads a group responsible for facilitating and assisting the formation of new CSIRTs around the globe. The group has assisted in the formation of a wide range of teams including those for national, government, Internet service provider, and academic environments.

Working in the computer security field for more than 7 years, she 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 rate of increase in 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 is an active figure in the international CSIRT community. She has 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 is currently the Steering Committee Chair.

West-Brown holds a first-class bachelor's degree in Computational Science from the University of Hull, U.K..

Klaus-Peter Kossakowski <kpk@work.de>

Klaus-Peter Kossakowski is a senior consultant and project manager at SECUNET, an IT security provider; and he is a visiting scientist within the CERT Coordination Center based at the Software Engineering Institute (SEI). Kossakowski's work currently involves incident response services, intrusion detection, network security, and security improvement.

Kossakowski has worked in the security field for more than 10 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 Coordination Center. He successfully led the team from a research effort to a functional and well-respected entity in the CSIRT community.

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 "CERTs in Europe." This task force outlined the concept and service definition of a European CSIRT Coordination Center. Resulting from 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 in this role he actively supports international CSIRT cooperation and the move of FIRST toward a new organizational structure.

Kossakowski is in the process of completing his Doctorate in Information Technology-Incident Response Capabilities. 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 company M&I/STELVIO, offering senior consultancy services in the areas of Internet, intranet, and security.

He has worked in the security area for more than 7 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 on he managed subcontract-

tors 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 1992-1998.

Stikvoort is an active participant internationally, particularly in regard to security issues, in RIPE, TERENA, IETF, and especially the FIRST community. Together with Klaus-Peter Kossakowski, he initiated the closer cooperation of European CSIRTs and contributed to the efforts leading to a more structured European incident coordination. He was chairman of the TERENA “TAG” group that selected the party currently delivering the EuroCERT service, and he was the first chairman of the EuroCERT Monitoring Committee. Recently, he has been actively involved in helping FIRST to evolve into a new organizational structure. Stikvoort is 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 its Dutch Chapter and the “Nederlands Genootschap voor Informatica” (NGI), and he participates in several national security groups.

Bibliography

- [Alfano 96]** Alfano, Joseph A. "Developing a Malicious Code Analysis Capability to Support Incident Handling." See [CSIIHW 8/96].
- [Aslam 95]** Aslam, Taimur. "A Taxonomy of Security Faults in the UNIX Operating System." Master's Thesis, Purdue University, 1995.
- [Brand 90]** 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.
- [Carpenter 98]** Carpenter, Jeffrey J. & Dunphy, Brian P. "Moving Towards the Exchange of Incident Statistical Data." See [CSIIHW 10/98].
- [CERT/CC 88]** "CERT/CC Advisories 1988-98." CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa., 1998. <<http://www.cert.org/advisories/>>; Tuesday, December 15, 1998; 4:13 P.M. EST.
- [CERT/CC 96]** "CERT/CC Product Vulnerability Reporting Form Version 1.0." CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. <ftp://ftp.cert.org/pub/vul_reporting_form>; Tuesday, December 15, 1998; 3:54 P.M. EST.
- [CERT/CC 97a]** "CERT/CC Incident Reporting Form, Version 4.3.3." CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa., December 1997 <ftp://ftp.cert.org/pub/incident_reporting_form>; Tuesday, December 15, 1998; 3:58 P.M. EST.
- [CERT/CC 97b]** "The CERT Coordination Center FAQ, Revision 10.8." CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa., April 1998. <http://www.cert.org/faq/cert_faq.html>; Tuesday, December 15, 1998; 4:02 P.M. EST.

- [CERT/CC 97c]** “CERT Security Improvement Modules.” CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Tuesday, December 15, 1998; 4:14 P.M. EST.
<<http://www.cert.org/security-improvement/modules.html>>
- [CERT/CC 98a]** “Incident Reporting Guidelines.” CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Tuesday, December 15, 1998; 4:16 P.M. EST.
<http://www.cert.org/tech_tips/incident_reporting.html>
- [CERT/CC 98b]** “CERT Summary CS-98.05 - SPECIAL EDITION.” CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Tuesday, December 15, 1998; 4:18 P.M. EST. <<http://www.cert.org/summaries/CS-98.05.html>>
- [CERT/CC 98c]** “CERT/CC Incident Notes.” CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Tuesday, December 15, 1998; 4:21 P.M. EST.
<<http://www.cert.org/summaries/CS-98.05.html>>
- [CERT/CC 98d]** “CERT/CC Vulnerability Notes.” CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Tuesday, December 15, 1998; 4:22 P.M. EST.
<http://www.cert.org/vul_notes/>
- [CERT/CC 98e]** “Problems With The FTP PORT Command – Tech Tip, Version 1.1.” CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Tuesday, December 15, 1998; 4:27 P.M. EST.
<ftp://ftp.cert.org/pub/tech_tips/FTP_PORT_attacks>
- [CERT/CC 98f]** CSIRT Code-of-Conduct. CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pa. Materials from the Managing Computer Security Incident Response Teams (CSIRTs) course, November 1998.
- [CERT-NL 92]** CERT-NL. “CERT-NL Operational Framework, Version 2.1.” Utrecht, Netherlands, June 23, 1992.
- [Chapman 95]** Chapman, D. Brent & Zwicky, Elizabeth. *Building Internet Firewalls*, 1st ed. Sebastopol, Calif.: O’Reilly & Associates, 1995.

- [CIAC 94]** Lawrence Livermore National Laboratories. "CIAC Bulletin: Computer Incident Advisory Capability." Livermore, Calif. Wednesday, December 16, 1998; 4:30 P.M. EST.
<<http://ciac.llnl.gov/cgi-bin/index/notes>>
- [Crabb 96]** Crabb, Michele. "How To Find and Hire Good Technical People." *Proceedings of SANS 1996 Conference*, Washington, D.C., May 12-18, 1996.
- [CSIHW 1/89]** Invitational Workshop on Computer Security Incident Response. Carnegie Mellon University, Software Engineering Institute. Pittsburgh, Pa., August 1989.
- [CSIHW 2/90]** Workshop on Computer Security Incident Handling. Carnegie Mellon University, Software Engineering Institute. Pittsburgh, Pa., June 1990.
- [CSIHW 3/91]** 3rd Workshop on Computer Security Incident Handling. Carnegie Mellon University, Software Engineering Institute, Herndon, Va., August 1991.
- [CSIHW 4/92]** 4th Workshop on Computer Security Incident Handling. Forum of Incident Response and Security Teams. Denver, Colo., August 1992.
- [CSIHW 5/93]** 5th Workshop on Computer Security Incident Handling. Forum of Incident Response and Security Teams. St. Louis, Mich., August 1993.
- [CSIHW 6/94]** 6th Workshop on Computer Security Incident Handling. Forum of Incident Response and Security Teams. Boston, Mass., July 1994.
- [CSIHW 7/95]** 7th Workshop on Computer Security Incident Handling. Forum of Incident Response and Security Teams, Karlsruhe, Germany, September 1995.
- [CSIHW 8/96]** 8th Workshop on Computer Security Incident Handling. Forum of Incident Response and Security Teams. San Jose, Calif., July 1996.
- [CSIHW 9/97]** 9th Workshop on Computer Security Incident Handling. Forum of Incident Response and Security Teams. Bristol, U.K., June 1997.

- [CSIIHW 10/98]** 10th Annual FIRST Conference on Computer Security Incident Handling and Response. Forum of Incident Response and Security Teams, Monterrey, Mexico, June 1998.
- [Dalton 90]** Dalton, Jerry. "Building a Constituency: An Ongoing Process." See [CSIIHW 2/90].
- [Devargas 95]** Devargas, Mario. *The Total Quality Management Approach to IT Security*. Oxford: NCC Blackwell, 1995.
- [FIRST 97]** Forum of Incident Response and Security Teams. "Forum of Incident Response and Security Teams (FIRST) Operational Framework." Wednesday, December 16, 1998; 4:35 P.M. EST.
<http://www.first.org/about/op_frame.html>
- [FIRST 98]** Nijssen, Teun & Ley, Wolfgang; Forum of Incident Response and Security Teams. "FIRST PGP FAQ Version 1.3." Wednesday, December 16, 1998; 4:37 P.M. EST.
<<http://www.first.org/docs/pgpfaq/>>
- [Fithen 96]** Fithen, Katherine T. "Hiring IRT Staff Interview Process." See [CSIIHW 8/96].
- [Garfinkel 96]** Garfinkel, Simson & Spafford, Eugene. *Practical UNIX & Internet Security*, 2nd ed. Sebastopol, CA: O'Reilly & Associates, 1996.
- [Gordon 95]** Gordon, Sarah. "Social Engineering: Techniques and Prevention.," 445-451. *Proceedings of the 12th World Conference on Computer Security, Audit & Control*, Westminster, U.K., October 1995.
- [Greening 96]** Greening, Tony. "Ask and Ye Shall Receive: A Study in "Social Engineering." *ACM SIG Security, Audit & Control Review* 14, 2 (1996): 8-14.
- [Halil 97]** Halil, Eric. "Coordinating Multi-Vendor Vulnerabilities: Why Is It So Difficult?" See [CSIIHW 9/97].
- [Icove 95]** Icove, David; Seger, Karl; & VonStorch, William. "Computer Crime: A Crimefighter's Handbook." Sebastopol, CA: O'Reilly & Associates, 1995.

- [IETF 97]** Internet Engineering Group Task Force. "An Open Specification for Pretty Good Privacy (openpgp), Charter 1997-1998." Wednesday, December 16, 1998; 4:48 P.M. EST.
<<http://www.ietf.org/html.charters/openpgp-charter.html>>
- [Kaufman 95]** Kaufman, Charlie; Perlman, Radia; & Spencer, Mike. *Network Security: Private Communication in a Public World*. Englewood Cliffs, N.J.: Prentice Hall, 1995.
- [Kossakowski 94a]** Kossakowski, Klaus-Peter. "The DFN-CERT Project." Wednesday, December 16, 1998; 4:50 P.M. EST. See [CSIIHW 6/94].
<<ftp://ftp.cert.dfn.de/pub/csir/dfncert/papers/6csihw.dfn-cert.ps.gz>>
- [Kossakowski 94b]** Kossakowski, Klaus-Peter. "The Funding Process: A Challenging Task." Wednesday, December 16, 1998; 4:52 P.M. EST.
See [CSIIHW 6/94].
<<http://www.cert.dfn.de/eng/team/kpk/6csihw2.html>>
- [Kossakowski 95a]** Kossakowski, Klaus-Peter. "Computer Emergency Response Teams and Their Role in Internet Security." *Proceedings of the HP International User Group Conference 1995*, Stuttgart, Germany, May 1995.
- [Kossakowski 95b]** Kossakowski, Klaus-Peter. "The Role of Site Security Contacts." See [CSIIHW 7/95].
- [Kossakowski 96a]** Kossakowski, Klaus-Peter. "Incident Trends: Observations Made By CERTs." *Proceedings of the Open System Security Europe, London*, U.K., February 1996.
- [Kossakowski 96b]** Kossakowski, Klaus-Peter. "Providing Incident Response Services." Tutorial during *Open System Security Europe*, London, U.K., February 1996.
- [Kossakowski 96c]** Kossakowski, Klaus-Peter. "Coordination of Incident Response Teams." *Proceedings of 28th I-4 Meeting*, Oslo, Norway, June 1996.
- [Kossakowski 96d]** Kossakowski, Klaus-Peter. "Commercialization of Incident Response." See [CSIIHW 8/96].

- [Kossakowski 97]** Kossakowski, Klaus-Peter. "From Incident Response to Incident Control Management." See [CSIH 9/97].
- [Kossakowski 98]** Kossakowski, Klaus-Peter. "Information Technology Incident Response Capabilities." Work in Progress. [Doktor Thesis at the University of Hamburg, Germany]
- [Longstaff 93a]** Longstaff, Thomas A. "Incident Role Playing: An Exercise to Develop New Insights Into the Process of Investigating a Computer Security Incident." See [CSIH 5/93].
- [Longstaff 93b]** Longstaff, Thomas A. *Results of a Workshop on Research in Incident Handling* (CMU-SEI-93-SR-20). Pittsburgh, Pa.: CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University, September 1993.
<<http://www.sei.cmu.edu/publications/documents/93.reports/93.sr.020.html>>
- [McGillen 93]** McGillen, Terry. "CERT Incident Communications." See [CSIH 5/93].
- [McGillen 97]** McGillen, Terry & Fithen, Katherine T. "Public Communications in the World of Incident Response." See [CSIH 9/97].
- [McMillan 96]** McMillan, Robert D. "Vulnerability / Advisory Processes." See [CSIH 8/96].
- [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 95]** Naval Research Laboratory, IS Security Group. *IS Security Incident Response Manual* (Code 1220.2). Washington, D.C: Naval Research Laboratory, 1995.
- [NRL 97]** Naval Research Laboratory, IS Security Group. *IS Security Incident Response Plan*. Washington, D.C: Naval Research Laboratory, January 1997.
- [Olson 94]** Olson, Jon. "Development of Security Policies." *Computers & Security* 13, 8 (1994): 628-636.

[Pethia 90a]	Pethia, Richard D. "Forming and Managing a Response Team." See [CSIHW 2/90].
[Pethia 90b]	Pethia, Richard D. "Developing the Response Team Network." See [CSIHW 2/90].
[Pethia 90c]	Computer Emergency Response: An International Problem / Richard D. Pethia; K. R. van Wyk. CERT Coordination Center. Pittsburgh, PA, 1990.
[Rezmierski 98]	Rezmierski, Virginia; Deering, Stephen; & Fazio, Amy. "What Did That IT Incident Actually Cost? The IT Incident Cost Analysis Model and Findings." See [CSIHW 10/98].
[RFC 1281]	Pethia, Richard D.; Crocker, Steve; & Fraser, Barbara. <i>Guidelines for the Secure Operations of the Internet</i> . Request For Comments 1281 (November 1991).
[RFC 1422]	Kent, S.T.; and J. Linn. <i>Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-based Key Management</i> . Request For Comments 1422 (February 1993).
[RFC 1984]	IAB and IESG. <i>IAB and IESG Statement on Cryptographic Technology and the Internet</i> . Request For Comments 1984 (August 1996).
[RFC 2196]	Barbara Fraser, ed. <i>Site Security Handbook</i> . Request For Comments 2196 (September 1997).
[RFC 2350]	Brownlee, N. & Guttman, E. <i>Expectations for Computer Security Incident Response</i> . Request For Comments 2350, Best Current Practice, June 1998.
[Schneier 95]	Schneier, Bruce. <i>Applied Cryptography: Protocols, Algorithms, and Source Code in C</i> . Chichester, U.K.: John Wiley & Sons, 1995.
[Sebring 93]	Sebring, Jeffrey. <i>Incident Aftermath and Press Relations: A MITRE Perspective</i> . See [CSIHW 5/93].
[Shimomura 95]	Shimomura, Tsumotu & Markoff, John. <i>Takedown</i> . Secker & Warburg, 1995, 324pp, ISBN 0-436-20287-5.

- [Smith 94]** Smith, Danny. *Forming an Incident Response Team*. University of Queensland: Brisbane, Qld., Australia, July 1994.
- [Smith 96]** Smith, Danny & West-Brown, Moira. "Incident Handling – Experience Through Role-Playing." Tutorial at [CSIIHW 8/96].
- [Sparks 97]** Sparks, Sandy; Fithen, Katherine; Swanson, Marianne; & Zechman, Pat. "Establishing an Incident Response Team." Tutorial at [CSIIHW 9/97].
- [Stikvoort 96]** Stikvoort, Don & Kossakowski, Klaus-Peter. "Incident Response Teams: the European Perspective." See [CSIIHW 8/96].
- [Stoll 89]** Stoll, Clifford. *The Cuckoo's Egg*. Doubleday, 1989, 326pp, ISBN 0-370-31433-6.
- [TERENA 95]** Kossakowski, Klaus-Peter, ed. "Final Report of the TERENA Task Force 'CERTs in Europe,'" October 1995.
- [UNI-CERT 96]** UNI-CERT. "UNI-CERT Operational Framework." [Private communication, 1996].
- [West-Brown 95]** West-Brown, Moira J. "Incident Trends." *Proceedings of the UNIX Network Security Conference*, Washington D.C., November 1995.
- [Wood 98]** Wood, Charles Cresson. *Information Security Policies Made Easy*, 6th ed. Sausalito, Calif.: Baseline Software Inc., 1998. ISBN# 1-881585-04-2.

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 this handbook.

Acronyms and Abbreviations

24x7	Twenty-four hours a day, seven days a week
AFS	Andrew file system
BCERT	Boeing CERT
CA	Certification Authority
CERT/CC	CERT Coordination Center
CERT-NL	Computer Emergency Response Team Netherlands
CIDR	Classless Inter-Domain Routing
CSIR	computer security incident response
CSIRT	computer security incident response team
DES	Digital Encryption Standard
DFN-CERT	Deutsches Forschungsnetz Computer Emergency Response Team
DNS	Domain Name System
DOE	Department of Energy
FIRST	Forum of Incident Response and Security Teams

FTP	file transfer protocol
FYI	for your information
Global Integrity REACT	Global Integrity's Rapid Emergency Action Crisis Team
GRIP	"Guidelines and Recommendations for Incident Processing"
HTTP	Hyper-Text Transmission Protocol
IBM-ERS	IBM Emergency Response Service
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
INND	Internet news daemon
IP	Internet protocol
IRT	incident response team
ISP	Internet service provider
ISS	Internet security scanner
MCERT	Motorola Computer Emergency Response Team
MD5	Message Digest 5
MIME	Multipurpose Internet Messaging Extension
NTP	Network Time Protocol
PCA	Policy Certification Authority
PEM	Privacy Enhanced Mail
PGP	Pretty Good Privacy
POC	point of contact

RFC	request for comments
SATAN/SANTA	System Administrator Tool for Analyzing Networks
S/MIME	Secure Multipurpose Internet Mail Exchange
SMTP	Simple Mail Transport Protocol
SSC	site security contact
SSH	Secure Shell
STU III	Secure Telecommunication Unit III
SUNSeT	Stanford University Network Security Team
TCP	Transmission Control Protocol
TERENA	Trans-European Research and Education Networking Association
Triple-DES	Triple Data Encryption Standard
TTP	trusted third party
UDP	User Datagram Protocol
UNI-CERT	Unisource Business Networks Computer Emergency Response Team
WWW	World Wide Web

Glossary Terms

Artifact (a.k.a. Critter)

Instances of malicious code. 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. Due to their characteristics it is usually differentiated between the authenticity (sometimes also referred to as integrity) of a message or file and the authenticity of a transaction.

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”)

Within the computer security arena, these events are often simply referred to as incidents.

Computer Security Incident Response (CSIR)

By providing the basic set of services (triage, incident, 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.

A team providing these service is called a Computer Security Incident Response Team (CSIRT). Within the computer security arena, these teams are often simply referred to as incident response teams (IRTs).

Depending on factors such as expertise and resources, the level and range of service provided might be different for various teams. Therefore each team will have to define the type of incidents that fall into the scope of their work and the level of service that they will provide under what circumstances.

Constituency

A specific group of people and/or organizations that have access to specific services offered by a CSIRT.

Intruder

An intruder is 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 didn’t constitute an intrusion. But because the intention of the person responsible for the attack the term 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.

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

Instead of collecting information by technical means intruders might also apply methods of social engineering like impersonating individuals on the telephone, or using other persuasive means to encourage someone to disclose information. As 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 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

Existence of a software weakness, design, or implementation error that can lead to an unexpected, undesirable event compromising the security of the system, network, application, or protocol involved.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

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	3. REPORT TYPE AND DATES COVERED Final
3. TITLE AND SUBTITLE Handbook for Computer Security Incident Response Teams (CSIRTs)			5. FUNDING NUMBERS C — F19628-95-C-0003
6. AUTHOR(S) Moir West-Brown Don Stikvoort Klaus-Peter Kossakowski			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			7. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-98-HB-001
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/DIB 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES			
12.A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12.B DISTRIBUTION CODE
13. ABSTRACT (MAXIMUM 200 WORDS) This document provides guidance on the generic issues to consider when 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 response (CSIR) service, which is the core service of a CSIRT. The document discusses 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 often sensitive information. In addition, operational and technical issues are addressed, 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 consists of managers responsible for the creation or operation of a CSIRT or a CSIR 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 response, CSIRT, incident response service, team operations, information handling, continuity assurance, security management			15. NUMBER OF PAGES 190
			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

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102