NIST Special Publication 800-150 (Draft)

Guide to Cyber Threat Information Sharing (Draft)

Chris Johnson
Lee Badger
David Waltermire

C O M P U T E R S E C U R I T Y

National Institute of Standards and Technology
U.S. Department of Commerce

NIST Special Publication 800-150 (Draft)

Guide to Cyber Threat Information Sharing (Draft)

Chris Johnson
Lee Badger
David Waltermire
Computer Security Division
Information Technology Laboratory

October 2014

U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie E. May, Acting Under Secretary of Commerce for Standards and Technology and Acting Director

Authority

This publication has been developed by NIST to further its statutory responsibilities under the Federal Information Security Management Act (FISMA), Public Law (P.L.) 107-347. NIST is responsible for developing information security standards and guidelines, including minimum requirements for Federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate Federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130, Section 8b(3), *Securing Agency Information Systems*, as analyzed in Circular A-130, Appendix IV: *Analysis of Key Sections*. Supplemental information is provided in Circular A-130, Appendix III, *Security of Federal Automated Information Resources*.

Nothing in this publication should be taken to contradict the standards and guidelines made mandatory and binding on Federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other Federal official. This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States. Attribution would, however, be appreciated by NIST.

National Institute of Standards and Technology Special Publication 800-150
Natl. Inst. Stand. Technol. Spec. Publ. 800-150, 73 pages (October 2014)
CODEN: NSPUE2

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by Federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, Federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. All NIST Computer Security Division publications, other than the ones noted above, are available at http://csrc.nist.gov/publications.

Public comment period: *October 29, 2014* through *November 28, 2014*

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in Federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, and its collaborative activities with industry, government, and academic organizations.

Abstract

In today's active threat environment, incident detection and response is an ongoing challenge for many organizations. This publication assists organizations in establishing computer security incident response capabilities that leverage the collective knowledge, experience, and abilities of their partners by actively sharing threat intelligence and ongoing coordination. This publication provides guidelines for coordinated incident handling, including producing and consuming data, participating in information sharing communities, and protecting incident-related data.

Keywords

computer security incident; coordinated incident handling; incident handling; incident response; information security; information sharing

Acknowledgements

The authors, Chris Johnson, Lee Badger, and David Waltermire of the National Institute of Standards and Technology (NIST), wish to thank their colleagues who contributed to this publication.

Table of Contents

List of Appendices

List of Figures

List of Tables

Executive Summary

As the magnitude and complexity of cyberspace increases, so too does the threat[1] landscape. Cyber attacks have increased in both frequency and sophistication resulting in significant challenges to organizations that must defend their infrastructure from attacks by capable adversaries. These adversaries range from individual attackers to well-resourced groups operating as part of a criminal enterprise or on behalf of a nation-state. These adversaries are persistent, motivated, and agile; and employ a variety of tactics, techniques, and procedures (TTPs) to compromise systems, disrupt services, commit financial fraud, expose sensitive information, and steal intellectual property. To enhance incident response actions and bolster cyber defenses, organizations must harness the collective wisdom of peer organizations through information sharing and coordinated incident response. This publication expands upon the guidance introduced in Section 4, Coordination and Information Sharing of NIST Special Publication (SP) 800-61, *Computer Security Incident Handling Guide* and explores information sharing, coordination, and collaboration as part of the incident response life cycle.

This publication assists organizations in establishing, participating in, and maintaining information sharing relationships throughout the incident response life cycle. The publication explores the benefits and challenges of coordination and sharing, presents the strengths and weaknesses of various information sharing architectures, clarifies the importance of trust, and introduces specific data handling considerations. The goal of the publication is to provide guidance that improves the efficiency and effectiveness of defensive cyber operations and incident response activities, by introducing safe and effective information sharing practices, examining the value of standard data formats and transport protocols to foster greater interoperability, and providing guidance on the planning, implementation, and maintenance of information sharing programs.

Implementing the following recommendations enables organizations to make more efficient and effective use of information sharing and collaboration capabilities throughout the incident response life cycle.

Organizations should perform an inventory that catalogues the information an organization currently possesses, the information that it is capable of producing, and document the circumstances under which this information may be shared.

By conducting an information inventory, an organization gains a better understanding of where its critical information resides, who owns it, how must it be protected, and when it can be shared. When deciding what incident-related information to share with other organizations, the following factors should be considered:

- Risk of disclosure
- Operational urgency and need for sharing
- Benefits gained by sharing
- Sensitivity of the information

[1] NIST Special Publication 800-30, Revision 1, *Guide for Conducting Risk Assessments* defines a *threat* as "any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destruction, disclosure, or modification of information, and/or denial of service".

- Trustworthiness of the recipients
- Methods and ability to safeguard the information

Organizations should exchange threat intelligence, tools, and techniques with sharing partners.

Organizations should move from informal, ad hoc, reactive cybersecurity approaches where the organization operates in isolation to formal, repeatable, adaptive, proactive, risk-informed practices where the organization coordinates and collaborates with partners. The Cybersecurity Framework[2] describes an approach that enables organizations – regardless of size, degree of cybersecurity risk, or cybersecurity sophistication – to apply the principles and best practices of risk management to improving the security and resilience their infrastructure. Through sharing, an organization benefits from the collective resources, capabilities, and knowledge of its sharing peers. When sharing threat intelligence, organizations have the opportunity to learn from each other; gain a more complete understanding of an adversary's tactics, techniques, and procedures; craft effective strategies to protect systems; and take action, either independently or collectively (i.e., as a sharing community) to address known threats.

Organizations should employ open, standard data formats and transport protocols to facilitate the efficient and effective exchange of information.

The use of standard data formats and protocols enables the automated exchange of information at machine-speed and allows different types of information from diverse sources to be readily correlated and analyzed. Standards can provide common identifiers that allow different organizations to unambiguously identify concepts, artifacts, or objects of interest (e.g., vulnerabilities, malware); define a common vocabulary to establish a shared understanding, or describe structures for encapsulating information for exchange. The use of standard formats and protocols fosters interoperability and allows disparate products, data repositories, and tools to rapidly exchange data and enables organizations to identify and mitigate threats in cyber-relevant time[3]. Organizations should choose formats that are widely adopted, readily extensible (i.e., new data elements or features can be incorporated with minimal engineering and design effort), scalable, and secure. Standardized formats are often highly expressive and support a wide-range of data elements; organizations should focus on using a manageable subset of data elements that provide maximum interoperability and the greatest value.

Organizations should enhance their cybersecurity posture and maturity by augmenting local data collection, analysis, and management functions using information from external sources.

By enhancing its local data collection and analysis capabilities, an organization can gain a more complete understanding of its systems and networks, and is able to make better use of the information that is available from external sharing partners. Correlating this data with information received from external sources and sensors can enhance data collected within an organization. Through the aggregation and analysis of information from internal and external sources the organization can build richer context about activities on its networks, identify campaigns, or better detect blended threats (i.e. threats that use multiple methods of attack). This enrichment process allows ambiguous data to be transformed into actionable information.

[2] See the *Framework for Improving Critical Infrastructure Cybersecurity* for additional information, http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf

[3] The term *cyber-relevant time* is a relative value that is based on the attack speed of the adversary. If an attack is unfolding then the network defender must implement response actions at the same speed or faster. This concept is discussed in greater detail in "Active Cyber Defense: A Vision for Real-Time Cyber Defense", MJ Herring, KD Willett, Journal of Information Warfare, Volume 13, Issue 2, April 2014.

Organizations should define an approach for adaptive cybersecurity that addresses the full cyber-attack life cycle.

Organizations should engage the adversary throughout the cyber attack life cycle and develop and deploy defensive measures that detect, limit, or prevent reconnaissance, delivery of malicious payloads, and the execution of exploits that allow an adversary to establish or maintain a persistent presence on an organization's systems or networks. Organizations should acquire cyber threat intelligence from both internal and external sources and use it to disrupt the adversary's cyber attack life cycle.

Organizations should ensure that the resources required for ongoing participation in a sharing community are available.

Participation in an information sharing community may require an organization to commit personnel; deliver training; and provide hardware, software, services and other infrastructure needed to support ongoing data collection, storage, analysis, and dissemination. Organizations must have a sustainable approach that provides the resources needed for ongoing participation to achieve sustained benefits from information sharing activities.

Organizations should protect sensitive information by maintaining an ongoing awareness of information security, vulnerabilities, and threats.

Organizations should implement the security controls necessary to protect its sensitive information, enforce its information sharing rules, and ensure that information received from external sources is protected in accordance with applicable data sharing agreements. Organizations should maintain an ongoing awareness of information security, existing vulnerabilities, and threats in the operational environment to support organizational risk management decisions.[4]

Organizations should establish the foundational infrastructure necessary to maintain its cybersecurity posture and clearly identify the roles and responsibilities for installing, operating, and maintaining these capabilities.

Organizations should have basic asset, vulnerability, and configuration management capabilities in place to ensure that the organization can actively monitor and manage the hardware and software residing on its networks and ensure that vulnerabilities are patched in a timely manner.

[4] NIST SP 800-37, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach* and SP 800-39, *Managing Information Security Risk: Organization, Mission, and Information System View* address the concept of information security risk management from the organization-level, mission/business process-level and the information system-level. NIST SP 800-137, *Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations* is intended to assist organizations in the development and implementation of an ISCM program.

1. Introduction

Cyber attacks are increasing as evidenced by reports from governments describing the security breaches to their computer systems. Further evidence comes from major corporations that have reported similar successful incursions. In addition, it is likely that many intrusions are undetected, go unreported, or have never been made public. As a consequence, criminal groups cause substantial losses to individuals and businesses and adversaries acquire valuable intellectual property and government secrets. All of these actions have a negative effect on the economic well-being and national security of the United States.

Among the challenges business and governments face is the need for a high degree of interconnectivity. The issue is such interconnectivity can allow attacks to spread quickly. To defend against cyber attacks, it is important for a defender to have timely access to relevant, actionable threat intelligence and the ability to act on that intelligence. This threat intelligence includes indicators (i.e., an artifact or observable that suggests that an attack is imminent, that an attack is underway, or that a compromise may have already occurred); the TTPs of an adversary; and recommended actions to counter an attack. Attackers often use similar strategies, tools, and methods against multiple organizations; therefore, it is important for organizations to share information with their peers.

When an organization identifies and successfully responds to a cyber attack, it acquires information that can be used by other organizations that face the same or similar threats. When information is shared, threatened organizations have access to threat intelligence provided by peer organizations and are able to rapidly deploy effective countermeasures and detect intrusion attempts. As a result, the impact of a successful cyber attack can be reduced.

1.1 Authority

The National Institute of Standards and Technology (NIST) developed this document to further its statutory responsibilities under the Federal Information Security Management Act (FISMA) of 2002, Public Law 107-347.

NIST is responsible for developing information security standards and guidelines, including minimum requirements for federal information systems, but such standards and guidelines shall not apply to national security systems without the express approval of appropriate federal officials exercising policy authority over such systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130, Section 8b(3), *Securing Agency Information Systems*, as analyzed in Circular A-130, Appendix IV: *Analysis of Key Sections*. Supplemental information is provided in Circular A-130, Appendix III: *Security of Federal Automated Information Resources*.

Nothing in this publication should be taken to contradict standards and guidelines made mandatory and binding on federal agencies by the Secretary of Commerce under statutory authority. Nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other federal official.

This guideline has been prepared for use by federal agencies. It may also be used by nongovernmental organizations on a voluntary basis and is not subject to copyright, though attribution is desired.

1.2 Purpose and Scope

This publication provides guidance that is intended to help organizations share information related to computer security incidents, communicate and coordinate with external groups, and manage the impact of the incidents on their organizations as well as the wider community. This document explores information

sharing architectures, examines how the maturity of an organization's cybersecurity capabilities affects its participation in a sharing community, and presents specific considerations for participation in an information sharing community. The guidance in this publication applies primarily to organizations that are familiar with the incident response life cycle presented in NIST SP 800-61, have some basic incident response capabilities in place, and are interested in exchanging information with other organizations.

1.3 Audience

This document is for computer security incident response teams (CSIRTs), system and network administrators, security staff, technical support staff, chief information security officers (CISOs), chief information officers (CIOs), computer security program managers, and others who are responsible for preparing for, or responding to, security incidents.

1.4 Document Structure

The remainder of this document is organized into the following sections and appendices:

- **Section 2** discusses the benefits of information sharing and incident coordination as well as the challenges facing organizations as they implement these types of programs. In addition, this section describes the fundamental concepts associated with incident coordination and information sharing including: (i) the cyber attack life cycle; (ii) threat intelligence; (iii) information sharing architectures; and (iv) formal and informal sharing communities.
- **Section 3** identifies the characteristics of organizations that have mature cybersecurity capabilities. The maturity of the organizations shapes their ability to effectively participate in incident coordination and threat sharing organizations. Individual organizations can perform a self-assessment, identify gaps, and define a plan to improve their organization's cybersecurity capabilities.
- **Section 4** identifies the key activities involved in implementing an incident coordination and information sharing capability. These activities are grouped by: (i) establishing sharing relationships; (ii) participating in sharing relationships; and (iii) maintaining sharing relationships. The section also provides guidance on how to protect shared information throughout the information life cycle.
- **Section 5** presents the general recommendations made in the publication.
- **Appendix A** contains computer security incident response scenarios that show how sharing threat intelligence and coordinating a response to incidents increases the efficiency and effectiveness of the organizations involved and enhances their network defense by leveraging the cyber experience and capabilities of their partners.
- **Appendix B** contains an alphabetical list of terms and their associated definitions
- **Appendix C** provides an alphabetical list of acronyms used and their expansion
- **Appendix D** lists resources that may be helpful in establishing and maintaining an incident response capability.
- **Appendix E** is the document change log.

2. Incident Coordination and Information Sharing Overview

In today's active threat environment, effective incident detection and response is an ongoing challenge for organizations. Information sharing and coordination provides a means of increasing the effectiveness of an organization's cybersecurity capabilities. Through collaborative incident response, organizations forge sharing partnerships that provide access to threat intelligence and tools that might otherwise be unavailable. Using these shared resources, organizations are able to enhance their network security posture by leveraging the knowledge, experience and capabilities of their partners. Allowing one organization's detection to become another's prevention is a powerful paradigm that can advance the overall security of organizations that actively share and coordinate. Threat information exchanged within communities organized around industry sector (or some other shared characteristic) can be particularly beneficial because the member organizations often face adversaries that use common TTPs that target the same types of systems and information.

Attacks may be part of coordinated campaigns targeting related industries or organizations by adversaries using sophisticated tools and techniques that are difficult for a single organization to detect or defend against. An organization whose threat horizon is limited to the activities that occur on their own systems and networks may be unaware of targeted attacks against their industry sector, technology stack, or the specific information that they possess. These attacks, when successful, are often quickly commoditized and directed against other organizations. An organization can gain greater awareness of the larger threat landscape by establishing the communication channels, data sharing agreements, and automation necessary to share information in advance of an incident. These preparations enable the organization to act decisively throughout the cyber attack life cycle.

Network defense is an intrinsically collaborative undertaking that is most effective when organizations coordinate and work together to face well-organized, capable adversaries. Coordination consists of multiple organizations communicating, cooperating, and exchanging information before, during, or after an incident in order to achieve common goals. Organizations can use shared information such as indicators, tactics, and tools to develop proactive defense strategies that focus on predicting an adversary's next move.

Organizations seeking to participate in sharing relationships need to be able to manage both the information they publish and the information they receive through all stages of the information life cycle. The life cycle of information, as described in OMB Circular No. A-130[5], consists of the following six phases:

- **Creation or Collection:** generating or acquiring information
- **Processing:** aggregating, transforming, correlating, and classifying information
- **Dissemination:** publishing and distributing information to authorized recipients
- **Use:** applying information to support organizational decision-making
- **Storage:** short and long-term retention of information in file systems, content management systems, databases, or other repositories
- **Disposition:** implementing and enforcing policies for the retention and disposal of information

[5] OMB Circular A-130, Transmittal Memorandum #4, *Management of Federal Information Resources*, http://www.whitehouse.gov/sites/default/files/omb/assets/omb/circulars/a130/a130trans4.pdf

The processes, guidelines, and agreements put in place for information sharing and coordination should address each of the information life cycle phases. The life cycle is an ongoing process that directly supports the generation, enrichment, maturation, and exchange of information between organizations.

2.1 Benefits of Information Sharing and Coordination

Incident response activities often include communication and interactions between a variety of organizations. By working together, these organizations can build and sustain the trusted relationships that are the foundation of secure and responsible information sharing and coordination. The benefits of collaboration throughout the incident response lifecycle include:

- **Shared Situational Awareness.** Information sharing and coordination enables organizations to leverage the collective knowledge, experiences, and analytic capabilities of their sharing partners, thereby enhancing the defensive capabilities of both organizations. Each member of a cybersecurity community of interest can profit from the knowledge and experience of other community members. Even a single contribution—a new tool or a description of an intrusion artifact—can increase the awareness and security of the entire community.

- **Enhanced Threat Understanding.** By developing and sharing threat intelligence, organizations gain a more complete understanding of the threat environment and are able to tailor and deploy security controls, countermeasures, detection methods, and corrective actions based on observed changes in the threat environment.

- **Knowledge Maturation.** When raw intelligence in the form of seemingly unrelated observations is shared and analyzed, it can be correlated with other data sets to build robust sets of indicators that are associated with a specific incident or threat and impart valuable insights into the relationships that exist between indicators.

- **Greater Defensive Agility.** As cybersecurity technologies advance, adversaries continually adapt their TTPs to counter the protective and detective measures implemented by network defenders. Organizations that possess the agility to rapidly detect and respond to changes in the adversary's TTPs can shift from reactive to proactive cybersecurity strategies.

- **Improved Decision Making.** Organizations that are able to consume and act on shared information are generally able to make decisions with greater speed and confidence. When adversaries are better understood, it is sometimes possible to anticipate their actions and deploy defensive measures before they act.

- **Efficient Handling of Information Requests.** Information sharing and coordination is an essential activity when reporting or investigating cybersecurity incidents that are criminal in nature. Organizations that have the processes, tools, and trained personnel in place to exchange information are better prepared to handle such information requests that arise and understand ensure that the computers and artifacts involved in the incident are treated as evidence and should be handled in a manner that preserves the chain of custody.

- **Rapid Notifications.** In the event an incident results in the release of information about another party (the victim), organizations are typically required to notify their affected customers or business partners. Government agencies and some industry sectors are subject to regulations that levy specific requirements for reporting of cybersecurity incidents. Organizations that understand their notification requirements and have notification procedures, contact information, and communications channels in place are able to rapidly disseminate breach notifications to affected customers or business partners. Appropriate sharing capabilities may be used, at least in part, to support these requirements.

2.2 Challenges to Coordination and Sharing

While there are clear benefits to sharing information, there are also a number of challenges to effective sharing and collaboration that must be considered.

- **Legal and Organizational Restrictions.** An organization's executive and legal teams may restrict the types of information that the organization can share. Restrictions may include limits on the types of information and the level of technical detail provided. Such restrictions are appropriate when they address legitimate business, legal, or privacy concerns; but the imposition of unwarranted or arbitrary restrictions may diminish the quality and timeliness of shared information.

- **Risk of Disclosure.** Knowledge of an adversary's TTPs is advantageous to a network defender but sharing of this information may put the contributor at risk by exposing the protective or detective capabilities of the organization and result in threat shifting by the adversary[6]. Additionally, disclosure of sensitive information, such as Personally Identifiable Information (PII), intellectual property, trade secrets, or other proprietary information can result in financial loss, violation of NDA's or other sharing agreements, legal action, and loss of reputation. Organizations should manage these risks using an appropriate risk management strategy.

- **Preserving Privacy.** Organizations may openly participate in information sharing communities, but still require that their contributions remain anonymous. This lack of disclosure may limit the usefulness of information to others since they cannot query the source of the information or understand the information's original context and provenance.

- **Producing Information**. Organizations seeking to produce information must have the necessary infrastructure, tools, and training to do so. While basic incident data (e.g., indicators, vulnerabilities) is relatively easy to produce; information such as an adversary's motives and TTPs generally requires greater effort.

- **Consuming Information**. Organizations must also have the infrastructure needed to access external sources and incorporate the information provided it into local decision-making processes. Information received from external sources has value only to the extent that an organization is equipped to act on the information.

- **Interoperability**. Standardized data formats and transport protocols help facilitate the interoperability needed for the secure, automated exchange of incident data between organizations, repositories, and tools, but agreement on formats and protocols requires careful analysis of costs and benefits.

- **Classification of Information.** Information received from government sources may be marked as classified information, making it difficult for an organization to use. It is also expensive and time-consuming for organizations to request and maintain the clearances needed for ongoing access to

[6] According to NIST SP 800-30, *Guide for Conducting Risk Assessments*, threat shifting is the response of adversaries to perceived safeguards and/or countermeasures (i.e., security controls), in which adversaries change some characteristic of their intent/targeting in order to avoid and/or overcome those safeguards/countermeasures. Threat shifting can occur in one or more domains including: (i) the time domain (e.g., a delay in an attack or illegal entry to conduct additional surveillance); (ii) the target domain (e.g., selecting a different target that is not as well protected); (iii) the resource domain (e.g., adding resources to the attack in order to reduce uncertainty or overcome safeguards and/or countermeasures); or (iv) the attack planning/attack method domain (e.g., changing the attack weapon or attack path).

classified information sources. In addition, many organizations employ non-U.S. citizens who are not eligible to hold security clearances and are not permitted access to classified information[7].

- **Establishing Trust.** Trust relationships form the basis for information sharing, but can be time consuming to establish and maintain. Ongoing communication, through regular in-person meetings, phone calls, or social media can help accelerate the process of building trust.

2.3 Cyber Attack Life Cycle

The attacks perpetrated by adversaries are growing in scale, scope, complexity, and frequency. Reactive defense strategies are not suitable for dealing with the advanced persistent threats that leverage sophisticated tools, zero-day exploits, and advanced malware to compromise systems and networks. While vulnerability and configuration management continue to be an important part of an organization's defensive strategy, these practices cannot fully address the threat posed by persistent adversaries who use advanced intrusion techniques. Although it is not feasible to fully predict adversary behavior, a cyber attack life cycle model can provide a simple, but useful abstraction for analyzing potential threats. Each phase in the cyber life cycle is an opportunity for a network defender to take action against an adversary. By using a cyber attack life cycle, in concert with both internal and external threat intelligence, network defenders can craft proactive incident response strategies that focus on disrupting the adversary earlier in the life cycle (i.e., before an exploit has occurred).

A number of the cyber attack life cycles exist, including Lockheed Martin's "Cyber Kill Chain®"[8] (shown in Figure 2-1) and the attack phase steps presented in NIST SP 800-115[9]. Figure 2.1 depicts 6 phases of a cyber attack:

- **Phase 1—Reconnaissance:** Adversary identifies and selects a target(s).
- **Phase 2—Weaponize:** Adversary packages an exploit into a payload designed to execute on the targeted computer/network.
- **Phase 3—Deliver:** Adversary delivers the payload to the target system(s).
- **Phase 4—Exploit:** Adversary code is executed on the target system(s).
- **Phase 5—Install:** Adversary installs remote access software that provides a persistent presence within the targeted environment or system.
- **Phase 5—Command and Control:** Adversary employs remote access mechanisms to establish a command and control channel with the compromised device.
- **Phase 6—Act on Objectives:** Adversary pursues intended objectives (e.g., data exfiltration, lateral movement to other targets)

[7] Executive Order 12968, *Access to Classified Information*, http://www.gpo.gov/fdsys/pkg/FR-1995-08-07/pdf/95-19654.pdf
[8] "Cyber Kill Chain" is a registered trademark of Lockheed Martin.
[9] The attack phase steps presented in NIST SP 800-115, *Technical Guide to Information Security Testing and Assessment: A Security Life Cycle Approach* are presented in the context of a penetration testing activity, but the activities described are similar to those that would be performed by an actual adversary. This publication is available at http://csrc.nist.gov/publications/nistpubs/800-115/SP800-115.pdf

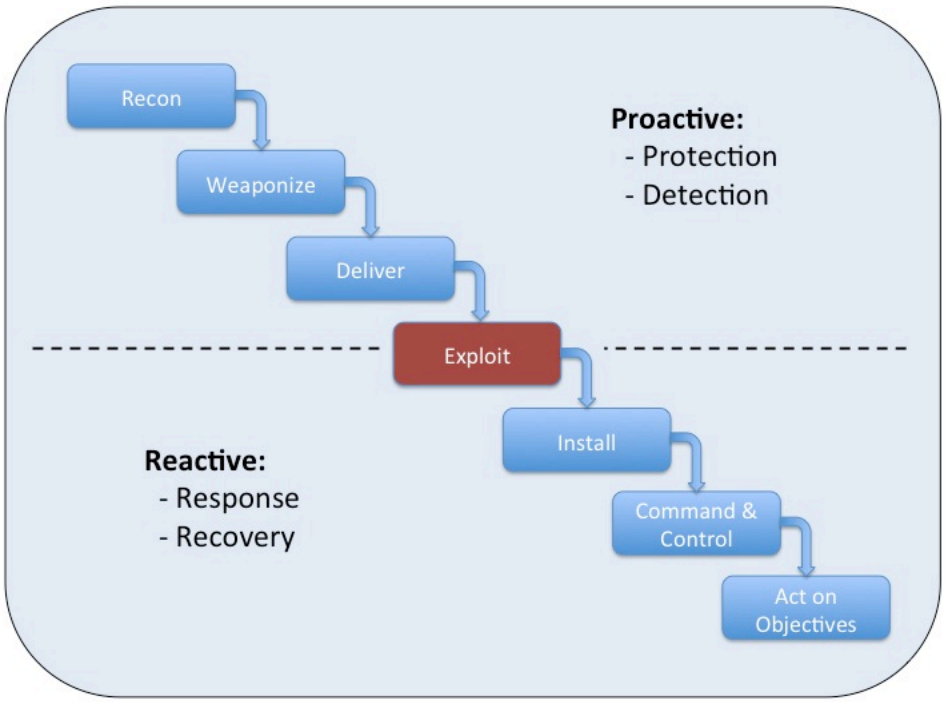

Figure 2-1: Cyber Kill Chain[10]

As depicted in Figure 2-1, proactive network defense (i.e., "above the line") consists of deploying protective and detective measures that disrupt an adversary before an exploit has been successfully executed. By recognizing and engaging the adversary during the reconnaissance, weaponization, and delivery phases of the cyber attack life cycle, network defenders are able to deploy mitigations or take some other course of action to ensure that mission-critical assets are protected prior to an adversary successfully executing an exploit. Reactive network defense (i.e., "below the line") relies on the organizations ability to detect the presence of an adversary on their networks and systems and craft an effective response and recovery strategy. Regardless of where interdiction occurs within the kill chain, the network defender must perform a retrospective analysis of the threat across the cyber attack life cycle to ensure that the response was effective. This analysis should include identifying indicators, determining where in the cyber attack life cycle these indicators were observed, and correlating these indicators with other threat intelligence. By understanding how an adversary operates over the cyber attack life cycle a network defender may be able to devise more effective defensive strategies. Examples of such defensive strategies and techniques, and where they can be applied within the cyber kill chain are described below:

- Reconnaissance. Perform monitoring and analysis of NetFlow, darknet, and passive DNS data to detect and investigate common network reconnaissance patterns such as port scans or probes. Employ anti-reconnaissance measures such as redirecting an attacker to a network black hole or by blocking specific IP addresses or domains.

- Weaponize. Develop, deploy, and refine high-fidelity signatures based on analysis of artifacts observed in malware payloads. Signature-based detection methods are generally fragile; adversaries can evade detection through minor modification to an exploit. By performing a more in-depth analysis of captured malware artifacts, more accurate and lasting detection signatures can be created,

[10] *Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains*, http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

and additional techniques can be selected and used, to identify new malware and variants of existing malware.

- Deliver. Understand the tools and techniques that an adversary uses to deliver malicious payloads and develop and deploy detective and protective measures that disrupt the adversaries deliver channels. These measures could be a technical (e.g., blacklisting of a site associated with a "watering hole" attack) or procedural (e.g., just-in-time awareness training for emerging threats).

- Exploit. Counter zero-day attempts by deploying defenses that help prevent attackers from injecting code into a running program, exploiting buffer overflow conditions, injecting operating system commands, or using access control weaknesses to gain greater system access. Organizations can also employ advanced threat modeling to characterize their attack surface and use fuzz testing to expose vulnerabilities in likely attack vectors.

- Install. Expose and actively respond to recently-installed malware by employing host and network-based intrusion detection signatures and tools such as file integrity checking, rootkit detection, and configuration change monitoring.

- Command and Control. Establish baselines of normal network and device activity and configure internal networks to detect anomalous inbound and outbound network traffic and changes in user and device behaviors. Monitoring against a baseline provides a means of detecting beaconing (i.e., outbound traffic on regular intervals) that may be associated with interactions with a command and control server.

- Act on Objectives. Deploy advanced data loss prevention solutions to detect abnormal data access, evasion techniques, and data exfiltration attempts to prevent unauthorized transmission or copying of sensitive data.

To mount an active defense, an organization should seek to understand an adversary's TTP within the cyber attack life cycle and possess and make use of detailed threat intelligence that is relevant, timely, and accurate. Information sharing among comparable organizations is an effective method for developing this level of intelligence. By observing an adversary's targets, activities, and behaviors over an extended time period a set of known TTPs can be developed for that adversary. Sharing this information with other defenders may enable those defenders to acquire valuable insights into an adversary's strategies and overall plans, thereby increasing the defender's ability to anticipate an intruder's behavior and develop a more vibrant and effective defense.

2.4 Threat Intelligence

Threat intelligence is a vital part of network defense and incident response. Organizations gather intelligence about the active threats to their environment and implement targeted defensive measures, both tactical and strategic. Threat intelligence includes information about threats, TTPs, and devices that adversaries employ; the systems and information that they target; and any other threat-related information that provides greater situational awareness to the network defender and incident responder. Effective threat intelligence exhibits the following characteristics:

- **Timely.** Intelligence should be rapidly delivered (i.e., ideally at wire speed), with minimal latency and provide sufficient opportunity for the recipient to anticipate the threat and prepare a suitable response. The timeliness of intelligence is context-dependent (i.e., cyber-relevant) and needs to take into account the volatility of the threat, the speed of attack, and the capabilities and TTPs of the adversary. Some decision cycles may require that tactical intelligence be delivered within seconds or

minutes to counter a fast-moving adversary, other threats may be more slow-moving and deliberate and can be effectively addressed using intelligence that is hours, days, or even months old.

- **Relevant.** Threat intelligence should have applicability within the recipient's operating environment, address threats that the organization is likely to face, attacks they are likely to see, and describe adversaries that the recipient is likely to encounter. Recipients of threat intelligence should perform a risk analysis to determine the degree of risk associated with a particular threat.

- **Accurate.** The threat intelligence should be correct, complete, and unambiguous. Inaccurate or incomplete information may prevent critical action, incite unnecessary action, result in an inappropriate response, or instill a false sense of security on the part of the recipient.

- **Specific.** Threat intelligence should depict the incident or adversary at a level of detail that addresses the salient facts about the threat, allows the recipient to understand how the threat may affect them, and allows them to evaluate possible courses of action.

- **Actionable.** Threat intelligence should ideally identify actions the recipient can take to counter the threat or provide sufficient information and context to allow the recipient to develop a suitable response to the threat.

Organizations should not only share information about successful intrusions, but also information about intrusion attempts — regardless of whether the intrusion actually succeeded. Sources of information include darknet servers (i.e., servers configured to capture traffic destined for unused address space or unallocated IP addresses), firewall, and IDS/IPS logs. Reports of attempted intrusions are often deemed less sensitive because sharing partners cannot readily draw conclusions about organization vulnerabilities or security resources from the information provided. Since information about attempted intrusions generally requires less sanitization and analysis, it can often be shared and acted on by the recipient more quickly.

There are many sources for cyber threat intelligence; organizations can collect and develop intelligence internally or acquire it externally through sharing communities, open sources, business partners; industry sector peers, product vendors, commercial cyber threat intelligence services, customers, law enforcement agencies, or other incident response teams.

Any insights regarding the motives and goals of the adversary are extremely valuable and should be documented. Personal relationships with trusted individuals or organizations are excellent sources of information, with the caveat that informal relationships may not be an enduring source of threat intelligence because individuals may move to other organizations or take on a new role within their current organization that no longer affords them access to the information that was previously shared. Internal threat intelligence sources include intrusion detection or protection systems, security information and event management products, antivirus software and file integrity checking software alerts; and operating system, network, service, and application logs[11]. The internal threat intelligence and related artifacts that are gathered should be retained and shared with partners as permitted by organizational policy.

Threat intelligence can also be acquired through sharing communities organized around industry sectors such as financial, electricity, or health. Organizations that operate within a specific sector should consider joining an established sector sharing community or, if none exist, consider forming one with other sector

[11] See NIST SP 800-61, *Computer Security Incident Handling Guide*, Section 3.2.3, for additional information on common sources of precursors and indicators.

peers. Organizations that operate in the same sector often have similar missions, operational environments, and data and often face the same threats and adversaries. In addition to industry sector groups, there are other communities that serve local, regional, and federal law enforcement; state and local governments; emergency responders, and other affiliations (see Appendix D for information on some incident response organizations).

There are many Internet-accessible open source threat intelligence outlets that publish indicators of compromise, blacklists, malware and virus information, spammer lists, and other information regarding emerging threats. Information originating from these sources may need to be manually collected and analyzed; a process that is time-consuming, labor-intensive, and potentially error-prone. Organizations that are unable or unwilling to take on such an effort may want to consider the use of a commercial cyber threat service provider that offers similar threat intelligence and other value-added capabilities for a fee.

2.5 Information Sharing Architectures

Most sharing communities exchange information using some variant of the following basic information-sharing architectures: (i) centralized; and (ii) peer-to-peer shown in Figure 2-2. The characteristic, benefits and challenges of each of these approaches are further explored in Sections 2.5.1 and 2.5.2.

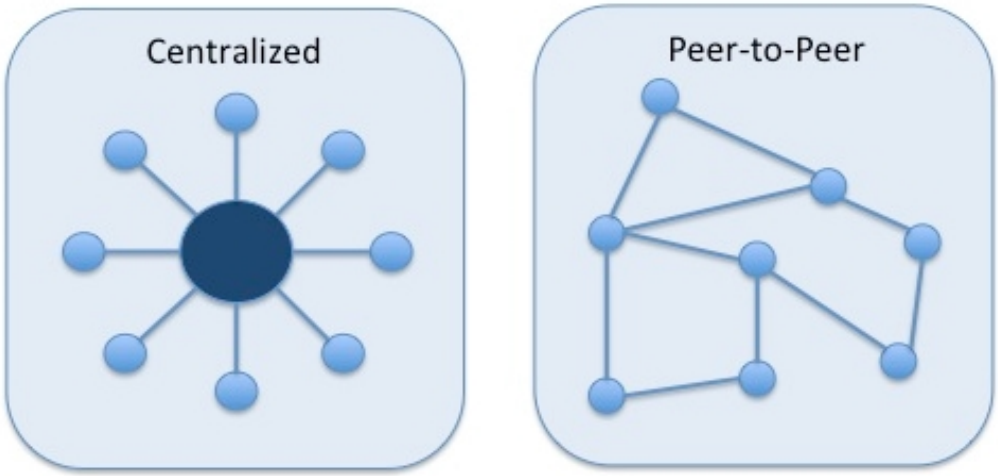

Figure 2-2: Information Sharing Architectures

The information sharing requirements for a community help determine the architecture that is most suitable. Some communities may benefit from a centralized approach; others may choose to exchange information directly among peers; still others may employ an approach that incorporates features and characteristics of both. When selecting an architecture for a sharing community, the following key factors should be considered:

- The characteristics, trustworthiness, capabilities, and composition of the participants
- The level of commitment of government, member organizations, and sponsors to support the community
- The type and sensitivity of information that will be shared
- The required frequency, volume, and speed of information distribution

2.5.1 Centralized Architecture

The centralized architecture is commonly described as a "hub-and-spoke", where a central "hub" serves as a repository or clearinghouse for information that it receives from the "spokes" (i.e., participating members) or other sources. Information provided to the hub by participating members is either directly forwarded to other community members (i.e., without any additional processing) or the hub may enhance the information in some way and then distribute it to designated community members. The enhancements performed by the hub may include aggregation and correlation of information from multiple sources, sanitization, de-attribution, enrichment of information by providing additional context, or trending and analysis that identifies common trends, threats, and malicious activity within the larger community.

Sharing communities based on this architecture usually establish formal data sharing agreements that stipulate what information can be shared, who it can be shared with, whether attribution is allowed, and the level of detail permitted. Information received by the central repository may be quite detailed, voluminous, and contain data elements that would enable attribution. The repository's summarization, sanitization and distribution processes should handle data in accordance with the data sharing agreements and provide abstracted, unattributed summary information to the sharing community as required. Central repositories that receive frequent, high volume submissions may choose to automate aspects of the summarization and sanitization process.

The benefits conferred by a hub-and-spoke architecture are largely determined by the services performed by the hub. The services provided by the central hub vary by community; some hubs may simply broker the information exchange, others may perform additional processing to enrich the information. In a hub-and-spoke community the central hub services can include consuming, aggregating, correlating, analyzing, validating, sanitizing, distributing, and archiving information from a multitude of sources.

Hubs that use open, standard data formats and transport protocols alleviate the need for participants to adopt multiple formats and protocols to exchange information with other community members. Additionally, participants have fewer connections to manage – once a connection to the hub exists, community members are connected to each other through the hub infrastructure.

The cost of the hub infrastructure is typically covered through membership or service fees paid by community members. If these fees are too high, they may present a barrier to entry and preclude organizations from participating in the community. A potential drawback to this architecture is that the information exchange system is entirely dependent on the hub's infrastructure, making it vulnerable to system failures, delays (e.g., due to network congestion, processing backlog, or other resource contention), or compromise at the hub. Though the time sensitivity of information varies, when the hub is not functioning or performance is degraded, all members of the sharing community are affected. A final consideration is that the hub, as a repository of threat intelligence, becomes an attractive target for attack.

Federal Government Response Teams

The hierarchical hub-and-spoke architecture (i.e., where security incidents are reported to centralized hierarchies within the government) is widely used within the Federal government. Figure 2-3 depicts a notional hub-and-spoke reporting structure for incident response teams operating across the Federal government and within specific departments and agencies. In this example, response teams participate as both a hub (to subordinate organizations) and a spoke (to a parent organization), depending upon where the team resides within the reporting hierarchy. In the Federal government, information flows from the agencies to the United States Computer Emergency Readiness Team (US-CERT) and/or the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT). In the DOD, information flows from

720 the combatant commands, services, agencies, and field activities to United States Cyber Command
721 (USCYBERCOM). USCYBERCOM coordinates with the US-CERT and ICS-CERT on cybersecurity
722 incidents, intelligence, and reporting involving the DoD[12].

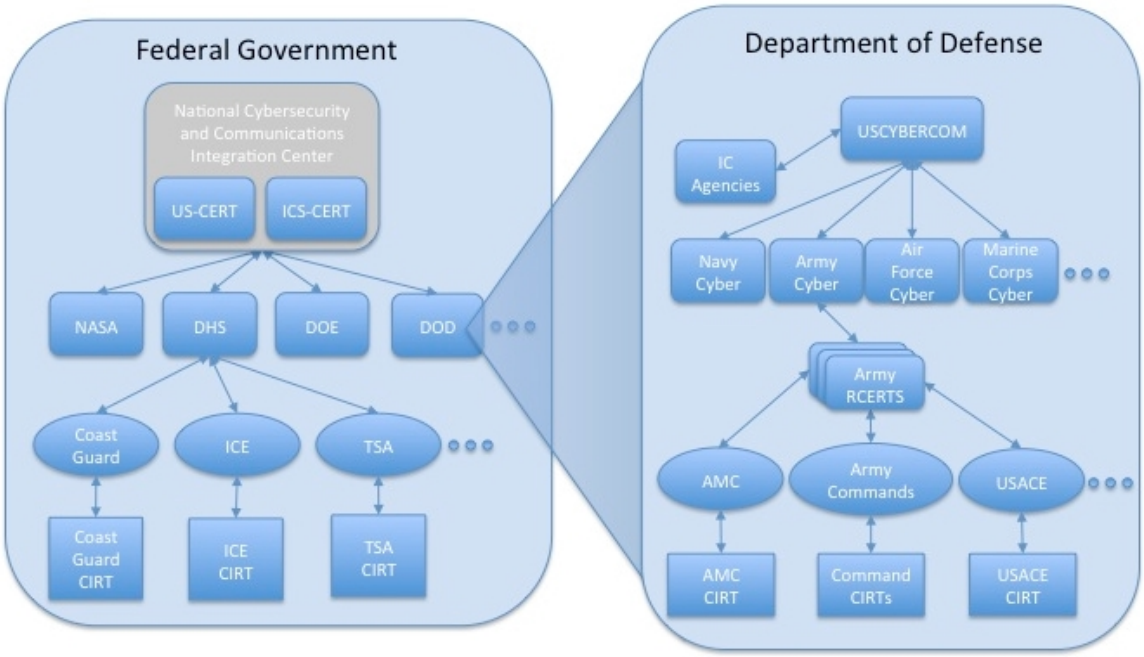

723
724
725 **Figure 2-3: Notional Federal Government Hub-and-Spoke Hierarchical Incident Reporting**

726 ***Information Sharing and Analysis Centers***

727 Another example of the hub-and-spoke model is the Information Sharing and Analysis Center (ISAC)
728 activities. Presidential Decision Directive-63 (PDD-63), published in 1998, describes ISACs as centers
729 for collecting, analyzing, sanitizing, and distributing information from the private sector to industry and
730 government. ISACs may also disseminate data from the government to the private sector. The private
731 sector participants determine the design and functions supported within the ISAC, with advice and
732 assistance from the Federal Government. Participation in an industry ISAC is voluntary. The National
733 Council of ISACs identifies 17 member ISACs[13].

[12] Chairman of the Joint Chiefs of Staff Manual (CJCSM) 6510.01B, Cyber Incident Handling Program, 10 July 2012
[13] http://www.isaccouncil.org/memberisacs.html

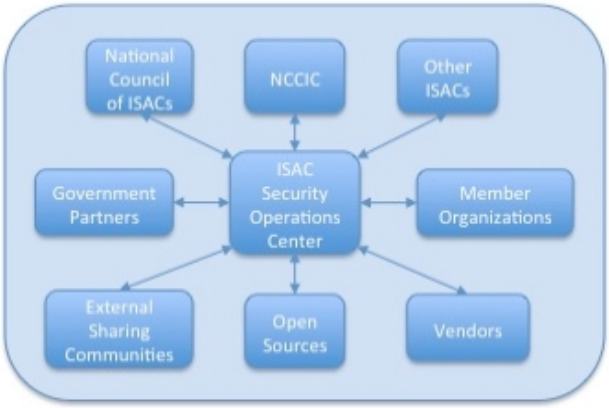

734

735 **Figure 2-4: Notional ISAC Hub-and-Spoke Incident Reporting Model**

736 In the notional ISAC model, illustrated in Figure 2-4, an ISAC Security Operations Center shares
737 incident, vulnerability and threat information with a variety of sources, including member organizations,
738 government partners, external sharing communities, vendors, and other ISACs. For example, a public or
739 private electrical utility company can join the Electrical Sector ISAC (ES-ISAC) and share information on
740 incidents and intelligence with that specific ISAC. The ES-ISAC would then share that information with
741 North American Electric Reliability Corporation (NERC), other ISACs, and the Federal government.

742 ### 2.5.2 Peer-to-Peer Architecture

743 Rather than routing data through a central hub, peer-to-peer participants share directly with each other.
744 Since no hub is present, each organization is responsible for consuming, aggregating, correlating,
745 analyzing, validating, sanitizing, protecting, and exchanging information with their peers. The
746 information that is exchanged between peers is limited to the data acquired, analyzed, and disseminated
747 by the participants. The dynamics of information exchange (e.g., security, speed, and frequency) will vary
748 according to the requirements and abilities of the communicating peers.

749 In a peer-to-peer relationship, trust is directly established with individual peers rather than brokered
750 through a central repository. Based on the level of trust established and the type of information being
751 exchanged, an organization may choose to share with a specific community member, a designated group
752 of recipients, or with all peers. Peer-to-peer trust is based on the belief that peers support a common
753 mission, respect the established sharing rules, and demonstrate a willingness participate in reciprocal
754 sharing.

755 The peer-to-peer architecture offers many benefits: (i) Peer-to-peer participants share directly with each
756 other (i.e., no intermediary such as the hub); this provides great agility and allows information to be
757 rapidly distributed as the receiver gets the information directly from the source. (ii) Peer-to-peer
758 architectures generally demonstrate greater resiliency since information is available through multiple
759 communication channels and there is no central hub that represents a potential single point of failure or
760 high-value target of attack.

761 The peer-to-peer architecture has some drawbacks including: (i) Peer-to-peer implementations that do not
762 employ standard methods of information exchange are difficult to scale since peers must support multiple
763 formats and protocols (ii) As the number of peer-to-peer sharing partners grows, the operating costs of
764 managing numerous connections, data (e.g., consuming, aggregating, correlating, analyzing, validating,
765 sanitizing, protecting, and exchanging), and trust relationships can grow exponentially.

Information exchanges between an organization and its Internet service provider (ISP), hosting provider, business partner, industry sector peers, law enforcement agencies, and other incident response teams and personnel often consist of peer-to-peer interactions. Such sharing, though not orchestrated through a sharing community, is nonetheless an important component of an effective incident response capability.

2.5.3 Hybrid Implementations

The two architectures previously described, are sometimes often in hybrid implementations that combine characteristics of both hub-and-spoke and peer-to-peer. Both centralized and decentralized P2P implementations exist. In a centralized peer-to-peer implementation, a central server(s) may be used for resource discovery, to broker requests, or as a trusted 3rd party for authentication. In a purely decentralized implementation, participants manage all aspects of their interactions with community peers.

An organization, for example, might exchange low-level intrusion indicators using a peer-to-peer architecture but send high-level incident reports to a central hub. Another scenario involves sending the same information directly to individual group members, as well as to the central hub. Such an approach enables both an effective tactical response (i.e., rapid action on time-sensitive data through direct, joint sharing) and makes use of the hub's ability to gather, combine, and analyze data received from multiple members to craft longer term strategies and courses of action. While the use of a hybrid approach may be advantageous in some cases, it can also increase costs and be more difficult to implement and operate.

2.6 Formal vs. Informal Communities

Information sharing communities exhibit varying degrees of formality. Some of the characteristics of formal and informal communities are presented below.

Informal sharing communities are generally self-organizing groups that operate through voluntary cooperation. Membership is mutable (i.e., no formal fixed membership), sometimes anonymous, and the members maintain full autonomy with minimal central coordination. These communities use informal data sharing agreements (i.e., rules of conduct rather than legally binding instruments) that establish the basic parameters for sharing information with the community.

Participants in an informal community publish information to a repository on a voluntary, ad hoc basis and are responsible for ensuring that content submitted to the repository is suitable for sharing. The repository operators maintain the repository but generally make no assertions regarding the quality and accuracy of the data contained within the repository; trust in the information is based on the reputation of the submitter. Organizations that wish to consume information subscribe to specific data sources hosted by the repository (e.g., email, RSS feed).

Formal sharing communities are often organized around a common characteristic (e.g. industry sector) and have official membership requirements that may define:

- Eligibility for institutions (e.g., specific industry sector)
- Eligibility for individuals (e.g., must have enterprise-wide security responsibilities)
- Nomination or sponsorship requirements (i.e., brokered trust)
- Probationary membership period
- Required organizational cybersecurity capabilities

Membership in such communities is generally fixed with minimal volatility in the membership rosters. Information exchange within the community is governed through SLAs, NDAs, and other agreements. Some communities collect an annual membership fee to cover the services and administrative costs of the community. These fees vary by community and the fee structure is sometimes tiered, providing for different levels of membership based on the organization type or size.

2.7 Recommendations

The key recommendations presented in this section are summarized below:

- Leverage the knowledge, experience, and capabilities of sharing partners to exchange threat intelligence, mitigation strategies, and tools, to enhance the cybersecurity posture of participating organizations and reduce the overall cost of cyber attacks.
- Establish and maintain information sharing relationships to enhance the organization's situational awareness and to foster a proactive approach to incident response.
- Use a cyber attack life cycle as a framework for observing and understanding an adversary's actions and for defining an active defense strategy that makes effective use of information available through both internal and external sources throughout the life cycle.
- Share information about intrusion attempts (regardless of whether the intrusion actually succeeded) rather than information about a specific intrusion. Intrusion attempt information is less sensitive and requires less sanitization and analysis; therefore it can be shared more quickly.
- Different sharing architectures exist for the sharing of information (e.g., centralized, peer-to-peer), as a participant in an information sharing community, understand both the benefits and drawbacks of these architectures.
- Seek out threat intelligence sources that provide information that is timely, relevant, accurate, specific, and actionable.

3. Understanding Current Cybersecurity Capabilities

Organizations should regularly assess the maturity of their cybersecurity capabilities and identify opportunities to enhance their overall security posture through information sharing and coordination. The purpose of this section is to describe the characteristics of a mature cybersecurity capability and a process by which an organization might become both a consumer and producer of actionable threat intelligence.

3.1 Characteristics of Mature Cybersecurity Capabilities

The maturity of an organization's cybersecurity practices is determined by its ability to establish and maintain an operational culture and the infrastructure necessary to actively manage cybersecurity risk. An organization must understand the cybersecurity threats to its systems, assets, data, and capabilities and prioritize its efforts, consistent with its risk management strategy and business needs. An organization should develop and implement protective measures that mitigate the impact of a potential cybersecurity incident, deploy capabilities that enable the timely detection and response to cybersecurity incidents, and be able to rapidly restore capabilities or services that were impaired due to a cybersecurity incident.

An organization should move from informal, ad hoc, reactive cybersecurity approaches where the organization operates in isolation to formal, repeatable, adaptive, proactive, risk-informed practices where the organization coordinates and collaborates with partners; such an approach is described in the Cybersecurity Framework.[14] The Cybersecurity Framework describes a process by which an organization can efficiently manage cybersecurity risk by selecting security controls that are consistent with the organization's risk management processes, legal/regulatory requirements, business/mission objectives, and organizational constraints. Security operations personnel should use information that originates from both internal and external sources to develop and deploy effective protective measures, detect network reconnaissance and attacks, identify threats, vulnerabilities, and indicators of compromise; and respond and recover from cyber attacks. Organizations that have high-performing security personnel in place are better poised to leverage sharing and coordination opportunities.

By participating in information sharing relationships an organization has access to a more extensive collection of cyber threat intelligence that can be used to help bolster its defenses. However, an organization that participates in sharing relationships does not thereby reduce or alleviate the need to deploy its own cybersecurity capabilities; it must still develop the local expertise and infrastructure to produce internal threat intelligence and to act on the information that it receives from external sources. Sharing and coordination is effective only if the recipient can act the information being shared; information is actionable when an organization possesses the core capabilities through which shared information can influence its detection, analysis, response, and recovery efforts. For example, shared threat intelligence that contains data elements, such as the IP addresses of a known or suspected adversary, is helpful only if the organization is monitoring IP addresses, has the ability to apply this information to a sensor device, and can identify what end points in the computer network were impacted. In another example, an organization may receive threat intelligence reporting that a compromise can be detected by observing the presence of a specific system artifact or a configuration setting holding a certain value. If the organization has no means of monitoring system artifacts or configuration settings, the shared information has no immediate value to the organization. Without core cybersecurity capabilities in

[14] The Cybersecurity Framework Tiers describe the degree to which an organization's cybersecurity risk management practices exhibit these characteristics (e.g., risk and threat aware, repeatable, and adaptive). See the *Framework for Improving Critical Infrastructure Cybersecurity* for additional information, http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf

place, sharing and coordination provides minimal benefit to an organization, since the information received is not actionable.

3.2 Consumer, Producer, and Capability Evolution

Often, entrants to a sharing community are primarily consumers of threat intelligence rather than producers of information. Sharing communities benefit from the dynamic and symmetric exchange of information, so an organization should seek to evolve from being a consumer only to become both a consumer and producer of threat intelligence. By producing threat intelligence, an organization gains greater expertise, helps other organizations more effectively respond to threats in their environment, and fosters trust with other community members.

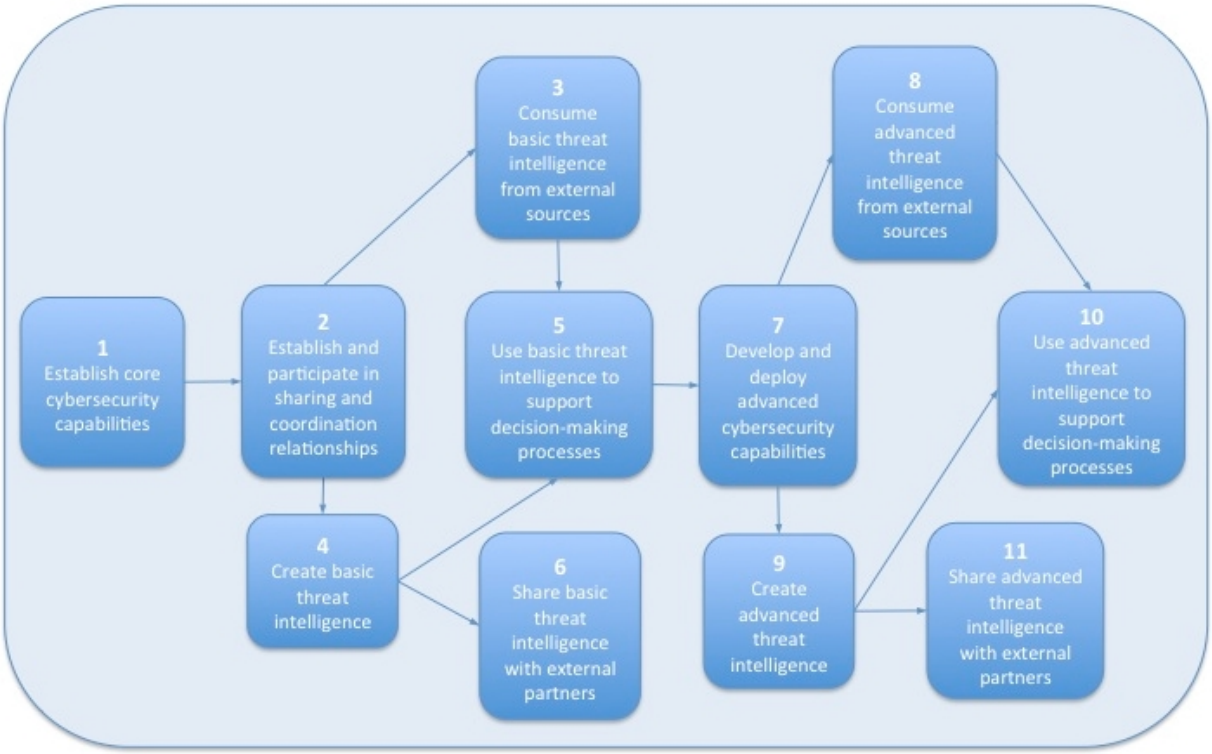

Figure 3-1: Notional Information Sharing Process

Figure 3-1 illustrates a process by which an organization can progress from an organization that initially possesses a set of core cybersecurity capabilities to become a more mature organization that consumes, creates, and shares cyber threat intelligence. The steps in this progression are described below:

1. **Establish core cybersecurity capabilities**[15]**.** An organization should deploy the infrastructure and processes necessary to support the core cybersecurity capabilities required to participate in information sharing and collaboration activities. These core capabilities include a monitoring infrastructure that is capable of supporting basic event and incident detection, analysis, and response

[15] The Computer Security Division's (CSD) Computer Security Resource Center (CSRC) facilitates broad sharing of information security tools and practices, provides a resource for information security standards and guidelines, and identifies key security web resources to support users in industry, government, and academia.

efforts. Examples are implementing boundary network monitoring capabilities such as an intrusion detection system (IDS) or a network-based antivirus (AV) appliance using vendor-provided signatures, and monitoring and responding to the alerts issued by these devices.

2. **Establish and participate in sharing and coordination relationships.** An organization should identify external sources of information that could be used to augment existing internal threat intelligence and enter into information sharing and coordination relationships. Section 4 of this document describes the process for establishing, participating, and maintaining information sharing relationships.

3. **Consume basic threat intelligence from external sources.** An organization should establish the infrastructure, processes, and training necessary to consume basic threat intelligence (e.g., simple indicators such as IP addresses, domains) from its sharing partners. External threat intelligence sources could include commercial, sector-based, or open source vulnerability, threat, and signature feeds.

4. **Create basic threat intelligence.** An organization should establish the infrastructure, processes, and training necessary to produce basic threat intelligence and disseminate it, as appropriate, to sharing partners.

5. **Use basic threat intelligence to support decision-making processes.** An organization should integrate the threat intelligence received from both internal and external sources into its current incident response processes and capabilities. For example, an organization might deploy enhanced IDS signatures, expand monitoring and assessment activities, or block IP addresses/ports based on the threat intelligence it possesses. The organization should use the threat intelligence to help prioritize response operations, enhance detection capabilities, and to develop and deploy effective courses of action.

6. **Share basic threat intelligence with external partners.** An organization should establish the infrastructure, processes, and training necessary to disseminate basic threat intelligence, as appropriate, to sharing partners.

7. **Develop and deploy advanced cybersecurity capabilities.** In some cases, external sources will possess threat intelligence that an organization has no means of consuming or acting on because of lack of infrastructure or expertise. In such cases, the threat intelligence is available only after the organization has expanded the scope of monitoring (e.g., monitor new sources or additional data elements or more frequently), performed skills development, or deployed more capable security tools. For example, the organization's host-based monitoring product may not be configured to (or able to) examine specific system artifacts and settings of interest. In addition, as an organization begins to engage more fully with its community peers, relationships grow and trust can be established which can help foster technical exchanges. Examples of advanced capabilities are establishing a forensics team that performs detailed network and computer forensics and malware analysis; deploying defensive capabilities such as honeypots, honeynets, and detonation chambers; or implementing advanced analytics and visualization functions that help expose an adversary's TTPs.

8. **Consume advanced threat intelligence from external sources.** An organization should establish the infrastructure, processes, and training necessary to consume advanced threat intelligence (e.g., TTPs, NetFlows) from its sharing partners.

9. **Create advanced threat intelligence.** An organization should establish the infrastructure, processes, and training necessary to produce advanced threat intelligence (e.g., TTPs, malware artifacts). As an organization develops new threat intelligence sources and new analysis techniques, they gain the

expertise needed to create and publish advanced threat intelligence and the ability to perform a more detailed and sophisticated analysis of incident data.

10. **Use advanced threat intelligence to support decision-making processes.** An organization should integrate the advanced threat intelligence received from both internal and external sources into its current incident response processes and capabilities. The use of advanced threat intelligence may allow the network defender to engage the adversary earlier in the attack life cycle and to deploy countermeasures or corrective actions that disrupt, delay, or prevent the adversary from achieving their goals.

11. **Share advanced threat intelligence with external partners.** Organizations that produce advanced threat intelligence possess information that may benefit others and should share it with others when possible. By acting as both a producer and publisher of information the organization is able to contribute new or enriched threat intelligence to the community.

3.3 Managed Security Services Providers Considerations

An organization's cybersecurity capabilities (core or advanced) may, in some cases, be implemented and maintained by a Managed Security Service Provider (MSSP). An organization may use a MSSP to provide capabilities that cannot be practically or cost-effectively developed in-house. MSSPs offer a variety of cybersecurity services and expertise that can be used to augment and enhance an organization's security capabilities.

There are many approaches to using MSSPs, and the degree to which an organization depends on an MSSP for their information sharing and incident coordination varies. Some organizations may choose to outsource all cybersecurity operations, while others only specific components or capabilities. Small to medium sized organizations may use an MSSP or a turnkey solution when the personnel and skills necessary to perform a task are not readily available within the organization, or in cases where the desired services can be provided by a MSSP at a lesser cost. When selecting a MSSP, the following factors should be considered:

- The MSSP should be engaged with information sharing communities and have ready access to actionable threat intelligence.

- The MSSP service level agreement (SLA) should clearly describe the responsibilities of the parties entering into the agreement and establish a dynamic, adaptive cybersecurity strategy that utilizes information received from both internal and external sources.

- An organization that relies on an MSSP to provide some portion of its cybersecurity operations needs to integrate the MSSP-provided capabilities with the organization's internal cybersecurity capabilities and support the exchange of threat intelligence between the organization and the MSSP.

3.4 Capabilities Self-Assessment

When considering incident coordination and sharing opportunities, an organization should determine if they have the capabilities necessary to effectively engage in these communities. The maturity of an organization's cybersecurity capabilities can be evaluated through an informal self-assessment. The self-assessment helps an organization better understand the maturity of its cybersecurity capabilities, which in turn helps determines its readiness to coordinate and share with external partners. For the purposes of the self-assessment process, maturity is defined at three levels: (i) underlying foundations and infrastructure; (ii) core cybersecurity capabilities; and (iii) advanced cybersecurity capabilities.

3.4.1 Underlying Foundation and Infrastructure Capabilities

Participation in an information sharing and incident coordination may require changes to an organization's policies and procedures, technology deployments, and personnel training. An organization must establish the groundwork and infrastructure necessary to maintain its cybersecurity posture and clearly identify the roles and responsibilities for installing, operating, and maintaining these capabilities. The underlying foundation and infrastructure, at a minimum, includes:

- **Organizational Structure for Incident Coordination.** An organization should have policies in place that: (i) define the management structures, roles, responsibilities, and authorities conferred to incident response team personnel; (ii) describe handoff and escalation procedures between team members and teams; (iii) identify the primary and backup communication mechanisms that allow incident response personnel to effectively coordinate with both internal and external stakeholders.

- **Asset, Vulnerability and Configuration Management.** An organization should have rudimentary asset, vulnerability, and configuration management capabilities in place to ensure that the organization can actively monitor and manage the hardware and software residing on its networks and ensure that vulnerabilities are patched in a timely manner.

- **Log and Alert Collection.** An infrastructure that supports the enterprise-wide collection of relevant log data and alerts generated by security products. The collection capability should provide wide coverage of the enterprise's computer network infrastructure; allow new log data sources to be incorporated with minimal effort; and allow the security analyst to change the type of data collected, the frequency of collection, or to discontinue the collection of certain data elements altogether.

- **Log and Alert Search and Retrieval**. Organizations should consider the use of a security information and event management solution that aggregates, analyzes, and correlates log and alert data and provides situational awareness for incident response personnel and network defenders and allows them to search and retrieve log and alert data and use the data to detect malicious activity, protect systems and data, and support incident response and recovery efforts.

- **Response Tools.** An organization should have the infrastructure and tools necessary to effectively contain, eradicate, and recover from a cyber incident. This includes tools and infrastructure for containment (e.g., sandbox network), digital system forensics, malware removal, and current system backups to support recovery efforts.

3.4.2 Core Cybersecurity Capabilities

Organizations that have the foundational infrastructure in place should monitor their infrastructure and establish a baseline for normal user, system, and network activities. By establishing a baseline, sensors can be configured to raise alerts when observed behaviors and activities significantly depart from the established baseline or exceed established thresholds for reporting.

Core cybersecurity capabilities include the ability to:

- **Deploy, configure, monitor, and update sensors.** An organization should have host-based sensors capable of collecting information regarding the status of processes, ports, files, services, hardware, software, and configuration settings on endpoint systems, and should have network-based sensors capable of active/passive monitoring of network activities to provide enhanced situational awareness. Operations personnel should review and respond to the alerts generated by these sensors and update the signature files and configuration of these devices to address false positives/negatives and to address emerging threats.

- **Manage log data.** An organization should generate, collect, aggregate, and manage relevant log, alert, and event information from across the enterprise. An organization may use a dedicated logging server, log management software, or a Security Information and Event Management product to allow the efficient collection, aggregation, analysis, and storage of log data.

- **Document, prioritize, and manage incidents**. An organization should have incident response procedures in place that document the incident handling process. These procedures should cover all phases of the incident response life cycle.

- **Perform basic network traffic forensics.** An organization should possess the tools (e.g., sniffer), log data and expertise necessary to correlate and analyze network events; identify common adversary techniques such as port scanning, probing, and IP address spoofing; and should possess a basic understanding of how adversaries use specific ports, protocols, and services to stage attacks.

- **Coordinate with system/information owners.** An organization should have processes and communication mechanisms in place that allow incident response personnel to effectively communicate with the owners of systems and information during an active incident. The owners may need to be consulted when response decisions may cause a service disruption or have some other operational impact.

3.4.3 Advanced Cybersecurity Capabilities

The distinctions between basic and advanced defensive capabilities are primarily based on the depth of analysis being performed and the role that information sharing and incident coordination plays in cybersecurity activities. Organizations practicing advanced cybersecurity capabilities are distinguished by their ability to:

- **Conduct "deep dive" digital forensics analysis of a compromise**. Advanced digital forensics includes the use of a full suite of tools, tactics, and procedures including:
  - Analysis of non-volatile data such as computer media, hard drives, USB sticks, and DVDs/CDs.
  - Analysis of volatile data including random access memory (RAM), running processes, open ports, open files, and network connections.
  - Export, analysis, and identification of malware and associated artifacts
  - Advanced packet capture analysis and network activity reconstruction
  - Dissecting network traffic and identify and export items of interest including command and control traffic and malware
  - Engaging in network traffic flow analysis (e.g. NetFlow)
- **Actively collect, produce, use, and share threat intelligence**. An organization should be actively engaged in the sharing of threat intelligence by:
  - Participating in coordination and sharing groups and forums
  - Acquiring and using threat and vulnerability information from external sources
  - Active coordination among computer network defenders, analysts, and operators
  - Using threat intelligence to drive sensor configuration and signature generation
  - Facilitating the production and sharing of threat intelligence within the organization and with external partners

- **Develop threat intelligence that reveals an adversary's TTPs, behaviors, and motives.** An advanced organization may seek to expose an adversary's TTPs through:
  - Malware capture, inspection, sanitization, and analysis
  - The use of a detonation chamber to explode files of interest (e.g., PDF, Word documents) for the purposes of malware and exploit detection, generally in temporary virtual environments
  - The deployment and monitoring of honeynets and honeypots
- **Use knowledge management practices to enrich data, mature knowledge, and inform cybersecurity decision-making.** An organization should develop and effectively use actionable information by:
  - Constantly refreshing and adapting defensive capabilities based on emerging threat intelligence
  - Using the knowledge of an adversary's TTPs to impede their progress, contain them, or prevent them from achieving their objectives
  - Using threat intelligence to inform the configuration of sensors, analysis platforms, and defensive measures

3.4.4 Information Sharing Capabilities

To consume and publish threat intelligence, an organization must demonstrate the ability to:

- **Coordinate the exchange of threat intelligence.** An organization should have the communication channels and business procedures in place that allow them to facilitate the exchange of information with both internal and external stakeholders.
- **Appropriately handle sensitive or classified information.** An organization should have the infrastructure and access control policies in place to preserve privacy and to ensure that sensitive information is afforded the required degree of protection.
- **Normalize or transform information**. An organization should have the ability to perform the data transformations necessary to make use of data received from external sources. These transformations may include, time synchronization, filtering, or rendering the information in alternate forms or formats.
- **Ingest information from external threat intelligence sources.** An organization should have the infrastructure and processes in place to ingest, store, and analyze the threat intelligence that it receives. Insufficient network, input/output, or processing capacity may result in information loss, data quality issues, and delays.
- **Produce and publish threat intelligence.** An organization should have the infrastructure and processes in place to produce and publish actionable threat intelligence.
- **Acquire actionable threat intelligence.** An organization must be able to acquire and use the threat intelligence from internal and external sources to:
  - Inform the development of signatures for intrusion sensors
  - Identify new artifacts and search terms during forensic analysis
  - Drive the configuration of honeypots and honeynets
  - Shape the tuning strategy for sensors and other monitoring instrumentation

3.5 Recommendations

The key recommendations presented in this section are summarized below:

- An organization should have, or develop, the underlying foundation and infrastructure in place to support information sharing and coordination activities
- An organization should seek out external information sources and enter into various information sharing and coordination relationships as their cybersecurity capabilities mature.
- An organization should consume information from external sources and apply the information to enhance their existing internal incident response capabilities
- An organization should expand their internal data collection, perform more sophisticated analysis, and begin to develop and publish their own indicators
- An organization may consider the use of an MSSP or outsourcing arrangement when the personnel and expertise necessary to perform a task are not readily available within the organization, or in cases where developing or maintaining a specific security capability in-house is not financial feasible
- An organization should perform routine self-assessments to identify opportunities for improved cybersecurity practices and more effective information sharing

4. Establishing, Maintaining, and Using Information Sharing Relationships

As defined in NIST SP 800-61, incident handling is structured as a four-phase life cycle: i) preparation; ii) detection and analysis; iii) containment, eradication, and recovery; and iv) post-incident activity, illustrated in Figure 4-1. Information sharing and coordination may occur in any or all of these phases. This section describes how an organization can establish, participate in, and maintain incident coordination and information sharing relationships throughout the incident response life cycle.

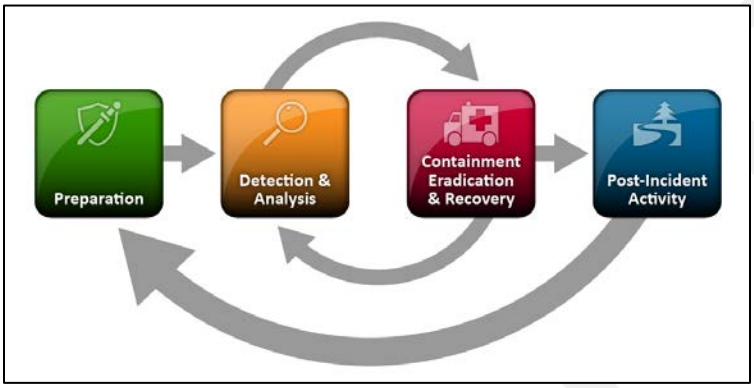

Figure 4-1: Incident Response Life Cycle

4.1 Establishing Sharing Relationships

When launching an information-sharing program, the following planning and preparation activities are necessary to help ensure the success of the initiative:

- Defining the goals, objectives, and scope of information sharing
- Conducting an information inventory
- Establishing information sharing rules
- Joining a sharing community
- Supporting an information sharing capability

These preparatory information-sharing activities are explored in greater detail in the following sub-sections.

4.1.1 Defining the Goals, Objectives, and Scope of Information Sharing

The first step in establishing an information sharing relationship is to set forth basic goals and objectives that describe what the organization hopes to accomplish. This need not be an onerous process; it is simply a matter of stating the desired outcomes of information sharing. In framing the information sharing initiative, the organization should also establish the general scope of the effort by identifying the resources (e.g., information, services, capabilities) that the organization could share, the resources that the

organization needs, the general conditions under which sharing is permitted, and potential sharing partners.

When establishing the initial parameters for information sharing, it is important to obtain approval from the management and legal teams (i.e., those with the authority to enter into commitments) and the support of key organizational stakeholders (i.e., those who will satisfy these commitments). Management commitment and authorization is generally easier to obtain when it can be demonstrated how information sharing helps to better protect the organization's critical assets, its reputation, and the well being of its customers, employees, and business partners. The leadership team plays an integral role and is responsible for providing continued oversight for the information coordination and sharing activities and for ensuring that resources are available to achieve specific objectives related to the organization's information sharing goals. The program's goals, objectives, and scope should be reevaluated and adjusted as needed, as mission or business requirements, priorities, technology, and regulations change.

Information sharing and coordination initiatives often require the participation of stakeholders from different internal organizational units. The stakeholders should possess a sound collective knowledge of cybersecurity operations; organizational business processes, procedures, and systems; and the ability to promote and support information sharing and collaboration within their functional units. The roles, responsibilities, and authorities (both scope and duration) of the stakeholders should be well understood, enabling decisive action before, during, and after an incident. Handoff and escalation procedures should be in place to allow the effective transfer of authority and flow of information to key decision makers throughout the incident response life cycle. The specific authorities given to team members should be enumerated; describing both the internal actions (e.g., empowered to add rules to an organization's firewall or temporarily disable specific systems or applications during an incident) and external collaboration (e.g., permission to share designated types of information with a specified sharing community, such as the US-CERT, law enforcement, legal teams, or the media) that team members are permitted to perform.

When possible, dedicated resources should be assigned to key leadership roles within the incident coordination and information sharing team, providing a trusted, consistent point of contact (POC) for internal and external sharing partners since high rates of personnel turnover can adversely affect the dynamics of sharing communities[16].

4.1.2 Conducting an Information Inventory

An organization initiating a sharing and collaboration effort should perform an inventory that identifies information that supports key business functions (e.g., financial, employee, or customer data that may contain PII; intellectual property) and security operations (e.g., security alerts, logs, analysis results, threat intelligence). Information should have an assigned owner who serves as the organizational point of contact for the information and is responsible for determining its sensitivity, the level of protection required, and for managing it throughout the information life cycle.

The inventory should identify the physical location (i.e., the geographic location of the server or storage media) and logical location (i.e., the network on which it resides) of the information. The inventory should identify how the information is stored; either as structured, machine-readable data (e.g., extensible markup language (XML), comma-separated values (CSV), JavaScript Object Notation (JSON)) or as unstructured data that has no pre-defined format (e.g., email message body, free text and images on web

[16] Merminod, V., Rowe, F., and Te'eni, D. *Knowledge Sharing and Knowledge Maturation in Circles of Trust: The Case of New Product Development*, 33rd International Conference on Information Systems, 2012

pages, business documents). The format of the information plays a significant role in determining the ease and efficiency of information exchange, analysis, and use. Information stored using open, machine-readable, standard formats can generally be more readily accessed, searched, and analyzed. As the number of sharing partners, frequency of sharing, and data volumes increase the need for standard data formats and interoperable protocols becomes more pronounced.

The inventory of information that supports security operations may include information derived from multiple sources within the organization including, IDSs, firewalls, antivirus software, and application logs. Specific data types and elements commonly of interest to incident handlers and network defenders include:

- IP addresses and domain names
- URLs involved with attacks
- Simple Mail Transport Protocol (SMTP) headers, email addresses, subject lines, and contents of emails used in phishing attacks
- Malware samples and artifacts
- Adversary Tactics, Techniques, and Procedures (and effectiveness)
- Response and mitigation strategies
- Exploit code
- Intrusion signatures or patterns
- Packet captures of attack traffic
- NetFlow data
- Malware analysis reports
- Campaign/actor analyses
- Disk and memory images

The information inventory is useful in a number of ways: (i) network defenders are able to develop prioritized monitoring and analysis strategies that focus on protecting the organization's most important information assets, (ii) an organization's resources can be more effectively allocated, (iii) ownership of information within the organization is formally established, (iv) information security analysts gain a better understanding of the likely value of the data source and the amount of effort required to acquire the information, **(**v) the organization is able to identify, understand, and document the information that is produced and consumed as part of business-specific workflows, (vi) the inventory can be used to develop guidelines, procedures, and mechanisms for information exchange.

As part of the inventory process, organizations consider how existing information sources might be used more effectively. For example, could information that the organization currently possesses be enhanced through additional analysis, through more frequent collection, or by aggregating and correlating information with other sources? Another consideration is to determine if incident response activities and defensive capabilities are adequately served by current sources of information. Any observed gaps should be documented and addressed through enhancements to local data collection capabilities, updates to

policy, or through external information sources as needed. The information inventory, once initially created, should be regularly updated to ensure that it is current, complete, accurate, and readily available.

4.1.3 Establishing Information Sharing Rules

Organizations should work with information owners, key management personnel, and the organization's legal team to establish and vet the rules governing the handling of sensitive information. This review should focus on identifying the general types of information that the organization may want to share with an incident response community and determining its sensitivity based on the risks of sharing the information inside and outside of the organization. Such risks may include, revealing the organization's network architecture and defensive capabilities to an adversary, exposing intellectual property, or the inadvertent release of PII.

4.1.3.1 Information Privacy

From a privacy perspective, one of the key challenges with sharing incident information is the potential for unauthorized disclosure of PII[17]. In the context of internal sharing, unauthorized disclosure could be disclosure to people who, by virtue of their job functions, would not typically have access to that PII in the normal course of business. They are performing a legitimate business function in terms of addressing the incident, but access to PII may not be truly necessary to adequately investigate the incident. For example, in conducting a forensics review of a hard drive, an analyst may review a file containing a list of employees that are under investigation for workplace hostility. The analyst does not have a need to know about the investigation, but may have a need to review the file for threat indicators associated with it. Generally, threat information that is shared externally is focused on actionable information for other organizations and should not contain PII.

Table 5.1 introduces various types of incident data, provides specific examples of each data type, and briefly discusses some of the sensitivity and privacy considerations when handling each type of data.

Type of Incident Data	Incident Data Elements	Sensitivity Considerations	Privacy Considerations[18]
Network Indicators	URLs, domains, IP addresses, script file names	Generally, information about the attackers is deemed less sensitive than information about the victim, so it can often be more readily shared. Before releasing information, the organization should consider the potential net intelligence-gain/loss. (e.g., a public	Attackers may possess personal information gleaned from open sources, acquired through social engineering techniques, or acquired from previous successful attacks (i.e., from a compromised system)

[17] OMB Memorandum 07-16 defines PII as information which can be used to distinguish or trace an individual's identity such as their name, social security number, or biometric records, alone, or when combined with other personal or identifying information which is linked or linkable to a specific individual, such as date and place of birth, or mother's maiden name. OMB Memorandum 10-22 further states that "the definition of PII is not anchored to any single category of information or technology. Rather, it requires a case-by-case assessment of the specific risk that an individual can be identified by examining the context of use and combination of data elements. In performing this assessment, it is important for agencies to recognize that non-PII can become PII, whenever additional information is made publicly available, in any medium and from any source that, when combined with other available information, could be used to identify an individual." NIST SP 800-122 includes a slightly different definition of PII that is focused only on the security objective of confidentiality and not privacy in the broad sense. Definitions of PII established by organizations outside of the federal government may vary based on the consideration of additional regulatory requirements. The guidance in this document applies regardless of the definition of PII by organizations.

[18] The PII confidentiality impact level as discussed in NIST SP 800-122 is a useful tool for gauging sensitivity of PII.

Type of Incident Data	Incident Data Elements	Sensitivity Considerations	Privacy Considerations[18]
		announcement that attacks are originating from a particular IP address will likely result in the adversary simply launching their attacks from an alternate IP address.)	
Packet capture	Network packet headers and payloads	Shared samples should filter on malicious traffic	Unencrypted or decrypted packets may contain PII such as logon credentials, financial information, health information, security investigation information, or information submitted via web forms
Phishing Email samples	Employee email	Email headers may contain infrastructure information such as internal IP address or hostnames	Consider anonymizing email samples and removing any sensitive information that is not relevant to incident responders
Webproxy logs	Logs of an organization's web activity, possibly including full URL's and parameters passed in requests	Log data may reveal business partner associations and contain logon credentials, portions of financial transactions, and other activities captured in URL parameters	Log data may contain PII regarding personal and business activity such as logon credentials, ID numbers used in URL parameters
Network traffic / "NetFlow"	NetFlow records provide a connection history between two IP addresses, including the time, duration, protocols used, number of packets exchanged, and number of bytes exchanged.	Generally less sensitive, though some organizations may not want to share full connection history and may "zero-out" low order bits in the IP addresses so that it is not possible to identify the network subnet.	NetFlow data may provide insight into employee behaviors or conditions that are not relevant to the investigation (e.g., access to websites about medical conditions)
Malware samples	Some artifacts associated with malware (e.g., log or staging files) may contain sensitive information from the victim's system.	Generally not considered sensitive, though proper handling, storage and encrypted transport should be used.	Context dependent based on a particular user's business and personal use of the resources that generate those artifacts

Table 5-1: Commonly Used Incident Data

The type of PII that may appear in incident data is situation-dependent, but the requirement to protect PII remains. To ensure adequate protection of PII in incident data, it is important to include the organization's privacy official in planning and development of an incident response program. Incident response policies and procedures should incorporate guidance from the organization's privacy official so that they address requirements for handling PII during incident response, including whether and how to share that information internally and externally. For example, incident response processes may include steps for identifying the incident data types that contain or are likely to contain PII similar to the table above and acceptable measures for addressing privacy risks associated with those data types.

When practicable, PII that is not relevant to investigating or addressing the incident should be redacted from incident data (e.g. working from a copy of incident data that has been scrubbed of known PII fields). Education and awareness activities are critical to ensuring incident response and sharing teams understand how to recognize and safeguard PII that is commonly encountered within the organization and are familiar with procedures for handling of PII.[19]

An organization may benefit from integrating security and privacy incident and breach response processes, as the processes are mutually supportive. Often times, incident response teams are in the position to first know when a security incident is also a privacy incident or breach. Privacy breaches carry an additional set of privacy requirements that must be addressed in close coordination with the organization's privacy official.[20]

4.1.3.2 Information Sensitivity

When participating in an information sharing community, it is sometimes necessary to share data collected from the business-critical computers and networks; data that could possibly contain sensitive information. It is therefore important that an organization document the circumstances under which information sharing is permitted by evaluating the risks of disclosure, the urgency of sharing, the trustworthiness of the information sharing community, and the methods available to safeguard shared information.

The information owner, management, and legal teams should adjudicate all sharing decisions using established procedures. The rules governing the sharing of information produced by the organization should be documented in local policies and procedures and expressed to external sharing partners through Memoranda of Understanding (MOUs), NDAs, Framework Agreements[21] or other agreements. Such agreements should established in advance of an actual incident and pre-vetted decision-making criteria should be in place, where possible, to control the risks of sharing while also enabling prompt coordination during an incident.

Many organizations handle information that is afforded specific protections under regulation or law. Examples of information requiring protection are privacy-related information such as PII, and information regulated under the Sarbanes-Oxley Act (SOX), the Payment Card Industry Data Security Standard (PCI DSS), the Health Information Portability and Accountability Act (HIPAA), the Federal Information Security Management Act of 2002 (FISMA), and the Gramm-Leach-Bliley Act (GLBA). An organization should consult its legal team and experts familiar with the various regulatory frameworks to identify protected classes of information within the organization.

The handling procedures established by an organization should specifically address the types of sensitive information that are likely to be encountered by incident response personnel and explicitly state the conditions (e.g., risk, urgency, trustworthiness of the information sharing community) under which management authorizes sharing of protected information, and the circumstances that require decisions be escalated to management. Information sharing rules are often context-dependent and require careful

[19] For additional guidance and examples of controls for protecting PII during incident response and sharing, see the following controls in NIST SP 800-53, Rev 4: IR-1, IR-2, IR-3, IR-4, IR-5, IR-6, IR-7, IR-8, IR-10, AR-3, AR-5, DM-1, DM-2, SE-2, TR-2, UL-2.

[20] See NIST SP 800-53, Revision 4, control SE-2, Privacy Incident Response

[21] An example of such an agreement is the Defense Industrial Base (DIB) Cyber Security/Information Assurance (CS/IA) Program standardized Framework Agreement which implements the requirements set forth in Title 32 Code of Federal Regulations, Part 236, Section 236.4 through 236.6. See Federal Register at http://www.gpo.gov/fdsys/pkg/FR-2013-10-22/pdf/2013-24256.pdf for additional information.

consideration of the nuances of the proposed sharing scenario to determine the extent or degree to which information should be shared. An organization's mission, legal requirements, regulatory environment, privacy concerns, and intellectual property considerations help shape these sharing policies. Through careful consideration of these factors, an organization must determine when the exchange of information is encouraged, limited, discouraged, or in some cases, forbidden. An organization may, for example, when a party in a lawsuit or other legal proceedings, chose not to share information that might under normal circumstances be readily shared. In some cases, information may be shared, but with specific restrictions (e.g., no attribution is permitted, specific data elements must be obfuscated before sharing).

These handling procedures seek to prevent the inappropriate release or mishandling of information, stipulate what information can be shared, when it can be shared, and how it must be protected. An organization's formal and informal information sharing agreements should stipulate protections consistent with approved information sharing rules. Should conditions change after a sharing agreement is in place, an organization should reserve the right to modify the agreement to accommodate emerging requirements. The documentation should be at a level of detail commensurate with organizational needs and updated at a frequency that does not impose an undue administrative burden. Incident responders, threat cell analysts, and operations personnel should, where possible, use automation to enforce information sharing rules to enable prompt, risk-managed, information coordination.

4.1.3.3 Marking

There are a variety of ways data can be marked-up or annotated in order to communicate how a message or document should be handled, or what specific elements might be considered sensitive and suitable for redacting, depending on an organization's needs.

Clear handling guidance should accompany any data that is intended for exchange. Examples of handling guidance or designations are:

- For Official Use Only
- Distribution limited to first responders
- Investigation underway, do not perform queries or active reconnaissance against these indicators

Data marking and handling procedures should be clearly documented and approved by management. The personnel responsible for handing data should be trained in these procedures. For some incidents or threat intelligence, the collection methods may be considered confidential or proprietary, but the actual indicators observed may be shareable. In such cases it is useful to organize reports with a so-called "tear-off" sheet of shareable items.

4.1.3.4 Procedures for Sharing and Tracking Incident Data

Over the course of time, an organization may face numerous attacks, participate in a large number of incident response efforts, and accumulate volumes of associated data. This data may be internally collected or may come from an external source. Tracking the source of data is important for both the protection of the information owners as well as for the enforcement of legal commitments such as NDAs. A balance must be struck between the need for rapid response and the obligations for protecting potentially sensitive data. When considering the capabilities of an organization's knowledgebase and data sharing processes:

- Develop a list of data types and content, such as indicators, that can be shared quickly with relatively minor review with established sharing partners.

- Develop a process for reviewing and protecting data that is likely to contain sensitive information.
- Store and track information regarding the sensitivity of data to be shared, including any relevant NDAs or other handling constraints.
- Track sources of data and with whom that data has been shared.

4.1.4 Joining a Sharing Community

Through the previous activities, an organization can better understand the information it currently collects and analyzes, the degree to which this information can be shared, and the additional information it needs to prevent incidents from occurring and to support the incident handling life cycle when they do occur. An organization can use this understanding to identify peers and other organizations with whom coordination and information sharing relationships would be beneficial. When evaluating potential sharing partners an organization should look to sources that complement the information collected internally (e.g., provides additional context), provide actionable information (e.g., indicators that an organization can readily use), and deliver information in a format and at a frequency that the organization is able to accept.

An organization may consider the use of open source information repositories, commercial services, government resources, and public/private sharing communities to enhance its IT, security, and incident handling processes. The public/private sharing communities often organize around some shared characteristic such as a geographic or political boundary, industry sector, business interest, threat space, or other common attribute. The coordination relationships may be team-to-team, team-to-coordinating team, or coordinating team-to-coordinating team. Potential sharing partners include: ISACs, CERTs, external CSIRTs, Product Security Incident Response Teams (PSIRTs), media outlets, security websites, social media, threat and vulnerability repositories, vendor alerts/advisories, commercial threat feeds, malware/antivirus vendors, supply chain partners, sector peers, customers, and known victims of cyber incidents.

When choosing a sharing community consideration should be given to the type of information that is shared within the community, the structure and dynamics or the community, and the cost of entry and sustainment. When evaluating the information that is shared within the community, consider the following questions:

- What information does the community provide/accept?
- Is the information relevant and does it complement locally-collected information? (i.e., provides meaningful insights into your organization's threat environment)
- Is the information actionable?
- Is the information timely, reliable, and of known quality?
- What is the frequency and volume of data disseminated?
- Does the organization have the capacity to ingest/analyze/store the information?

In addition to the information shared within the community, consideration should also be given to the dynamics of the community and its participants, including:

- What information-sharing model does the community use? (see section 2.5)

- What is the size and composition of the community? (e.g., number of participants, information producers, and information consumers)
- How active is the community? (e.g., number of content submissions/requests)
- How trustworthy are the community members?
- What are the technical skills and proficiencies of the community members?
- How are decisions made within the community?
- How is information communicated to its participants? (e.g., delivery mechanisms, formats, protocols)
- What is the cost of entry and sustainment? (e.g., commercial service offerings, resources)
- What type of sharing agreement does the community use? (e.g., formal vs. informal)
- Is the sharing agreement well aligned with organizational goals, objectives, and business rules?

When evaluating potential sharing partners, a great deal can be learned by observing the dynamics of the sharing community. Conversations with current or former community members may also provide valuable insights into community dynamics and the trustworthiness of its members. The trustworthiness of a community and its constituents is manifested in a multitude of ways, including the knowledge, skills, experience, integrity, reliability, communication abilities, and level of commitment of the community's members. NIST SP 800-39, *Managing Information Security Risk; Organization, Mission, and Information System View*, describes the following trust models that can be used to establish and maintain the level of trust needed to form partnerships, collaborate, share information, or receive services.

- **Validated Trust.** One organization obtains a body of evidence regarding the actions of another organization and uses that evidence to establish a level of trust with the other organization.
- **Direct Historical.** The track record exhibited by an organization in the past is used to establish a level of trust with other organizations.
- **Mediated Trust.** An organization establishes a level of trust with another organization based on assurances provided by some mutually trusted third party.
- **Mandated Trust.** An organization establishes a level of trust with another organization based on a specific mandate issued by a third party in a position of authority.
- **Hybrid Trust.** An organization uses one of the previously described models in conjunction with another model(s).

Mature sharing communities communicate regularly (e.g., using conference calls, email, portals with forums, social networking tools, and face-to-face meetings) to distribute and discuss current security threats, provide training and skills development, develop and share mitigation strategies, and define incident handling best practices. The level of maturity of the participating organizations often varies: some possess advanced monitoring, analytical, and forensic capabilities that allow them to produce information to share; other less mature organizations will participate primarily as information consumers.

One mechanism for building trust is to orient the information exchange around a shared mission or business objective—creating a setting where members often confront common threats. This focus on common threats fosters greater cohesion within the community and provides greater focus. Trust can be

further established and strengthened through face-to-face meetings between members and other events that help establish a level of personal rapport. Trust is also built as members share relevant technical insights, collaboratively build greater competency, work together to solve common problems, and lay a foundation to strengthen relationships through ongoing interactions with their peers.

The expectations and responsibilities of the participants in these sharing relationships may be expressed in a variety of ways, including data sharing agreements, association bylaws, or other agreements. Although some information sharing communities operate informally, based on personal reputation and verbal agreements, others are based on more formal expressions of policy such as NDAs, SLAs, or other agreements. Small informal circles of trust are generally tight-knit sharing communities where reputation-building occurs over time through personal relationships and the demonstrated technical prowess of its members. Regardless of the degree of formality, when entering into any type of information sharing agreement it is important to adhere to the organization's information sharing and handling rules and ensure that incident coordination personnel have clear guidance regarding redistribution of information received from the community.

As given in SP 800-61, having contact lists of key personnel is important when responding to an incident. If contact information must be supplied to a community, be sure to understand the degree of control that is provided over the visibility of this information to external users, community partners, and operators of the community (e.g., moderators, administrators). In bi-directional information sharing and coordination communities, the need for individual contact information may be necessary but a balance must be maintained between visibility, accessibility, and privacy. Participants in communities employing the hub-and-spoke model may not know other community members and only interact with the community's moderators or administrators. In addition to keeping contact information for selected peer organizations within an information sharing community, alternate communications mechanisms should be identified in case an incident compromises, disrupts, or degrades the community's primary communication channels.

4.1.5 Support for an Information Sharing Capability

The threat intelligence that an organization receives should be applied as part of an overall computer network defense strategy, not simply in response to a known incident. An organization should have personnel, infrastructure and processes in place to collect and analyze the information from both internal and external sources. This information should be used proactively throughout the incident response life cycle to design and deploy better protective measures, to more effectively perform signature and behavior-based detection, and to inform containment, eradication, and recovery operations. An organization will incur costs related to its participation in information sharing and coordination activities but may avoid larger costs from successful attacks. It is important for an organization to approach processes and technology in a way that is sustainable based on their resourcing levels and overall goals. Human and IT resources should be applied in a way that maximizes their benefit. Once a sustainable approach is developed, it is important to ensure that adequate funding exists to cover personnel; training; hardware, software, and other infrastructure needed to support ongoing data collection, storage, analysis, and dissemination; and any membership or service fees required for participation in these communities.

4.2 Participating in Sharing Relationships

An organization must establish operational practices that are compatible with those of the information sharing communities in which it is a member to make the most effective use of this additional information. Some practices are related to the types of information that are exchanged, the information's structure, the mechanisms for exchange, or semantics; others focus on the protection of information exchanged within the information-sharing community, or with the governance of the community.

Participation in an information sharing community encompasses a number of related activities:

- Engaging in on-going communication
- Implementing access control policies for shared information
- Storing and protecting threat intelligence, incident data, corrective measures, and evidence
- Consuming and responding to alerts and incident reports
- Consuming and analyzing indicators, TTPs, and corrective measures/course of actions
- Creating written records
- Performing local data collection
- Producing and publishing indicators, TTPs, and corrective measures/course of actions
- Producing and publishing incident reports

The following sections expand on each of these activities.

4.2.1 Engaging in On-going Communication

Information sharing communities use a variety of methods for communicating, depending on the nature of the information to be shared and the speed with which it must be disseminated; some methods, such as email lists or portals, make it possible to participate in a relatively passive, low-cost manner for some organizations. Other methods, such as conferences and workshops, require dedicated staff and travel. For organizations that actively produce information for other community members, communication costs are likely to be relatively higher. Communications may be event-driven, e.g., in response to the actions or behavior of an adversary, or they may be periodic, such as bi-weekly reviews, teleconferences, and annual conferences.

Message volume and frequency can vary widely across information sharing communities and largely depends upon the volatility of the attributes being observed, the importance that the community places on having the most current information, and the intended audience of the information. High volume sharing communities may publish summary information or digests (i.e., instead of sending individual messages, a collection of messages are sent that cover a specified period of time) to reduce the frequency of message traffic. Some recipients may be seeking only summary data (e.g., rollups) and have no need for detailed information. For an organization that has recently joined an information sharing community, just keeping up may be a significant effort, particularly until the organization has developed the skillsets needed to evaluate messages received (or found on a portal). In the early phases of participation, an organization may wish to focus on studying any best-practices guidance offered by the community, observing the messages sent by more experienced members, and querying databases made available by the community.

An organization's personnel should possess the technical skills needed to effectively communicate within their information sharing communities. The specialized skills required for incident handling and coordination are acquired over time through hands-on experience and training. Organizations should seek to minimize turnover within this team to foster enduring information sharing relationships, minimize knowledge loss, and preserve investments in training. Stability within the incident coordination team facilitates the formation of trusted professional relationships that span different CSIRTs and organizations — relationships that can be crucial during incident response.

In addition to developing technical skill sets and professional relationships, information sharing communities should employ communications protection measures when coordinating. Some communities issue authentication credentials for a web portal that can be used for coordination; in this case, the security of the portal itself (and the implementation of secure communication channels between clients and the portal) provides communications security. Other communities may issue or rely on a certificate hierarchy allowing participants to use public key cryptography to allow message senders to encode messages so that only designated receivers can decrypt. Other communities may use a web of trust model[22], in which certificates are distributed without a single hierarchy. Other communities may use dedicated physical networks, virtualized networks (e.g., peer, overlays), or a message bus as a secure media for conducting coordination activities. Protecting communications among participants is extremely important, particularly when the messages may contain information about techniques used by an adversary, PII, proprietary, or other sensitive information.

When one or more organizations are under attack or have been compromised, it is important for defenders to establish a means of secure communications, ideally physically and logically separate from the enterprise's infrastructure. An alternative cellular phone provider and externally managed collaboration portal are examples of such independent communication channels. If one believes that telecommunications services may be subject to eavesdropping, one may consider encrypting the voice channel as well. It is important to establish these communications amongst defenders before an incident takes place. Alternate data communications channels to share breaking threat indicators in the event of compromise may also be necessary to avoid eavesdropping by an adversary.

In addition to managing the communications mechanisms in a secure way, it is also necessary to ensure the efficient dissemination of information within the organization. Drawing on some of the key concepts presented in NIST SP 800-39, coordinated incident management processes should aim to operate seamlessly across all tiers of the organization at the (i) organization level; (ii) mission/business process level; and (iii) information system level. Inter-tier and intra-tier communication should be employed to create a feedback loop for continuous improvement and to help ensure that all stakeholders in the intrusion response are fully informed and effectively engaged in decision-making processes.

Decision-making in support of incident handling follows a similar model, where multiple incident response decision-making loops are executed concurrently with coordination and communication occurring in and between organizational tiers[23]. The established roles, responsibilities, and scope of authorities conferred to participants determine, to a large extent, how information sharing and coordination occurs within an organization. For example, operations personnel may be permitted to make decisions regarding configuration changes without seeking approval from the management or legal teams, provided the changes do not negatively affect customers or business partners, or prevent the organization from satisfying its business, legal, or regulatory obligations. The goal is to provide information that can be acted upon by stakeholders in the incident response process across all organizational tiers. The information provided can be used to inform policy changes at the organizational level, process changes at the mission/business level, or actions at the information system level, including patching, system configuration changes, introducing additional access control rules, removing devices from the network, or making network architecture changes.

[22] The web of trust concept was introduced by Pretty Good Privacy (PGP).

[23] Information regarding the Coordinated Incident Handling model is available in the IEEE publication titled *Operationalizing the Coordinated Incident Handling Model*

4.2.2 Implementing Access Control Policies for Shared Information

In order to address the risk of unauthorized disclosure of information, organizations should establish and enforce access control policies appropriate for the information being protected. The organization must ensure that access controls are in place, functioning as intended, and that processes are in place to establish oversight and accountability for the controls. Access control policies should take into consideration the information sharing rules and handling requirements established by the organization (see Section 4.1.3, Establishing Information Sharing Rules) and those expressed in information sharing agreements executed with partners. Multi-national organizations need to consider the national or regional policies related to privacy and information sharing when establishing and enforcing access control policies (e.g., sharing between business units operating in different countries). Additionally, access to information of a certain categorization or classification may be limited by business unit or department, role, or group membership.

When exchanging information with external entities, organizations must protect and distribute two basic types of information:

- Information produced within the organization (i.e., locally-produced)
- Information received by the organization from external sources

4.2.2.1 Locally-Produced Information

Locally-produced information may contain sensitive information, including critical business information, technical information that could reveal vulnerabilities in an organization's computing infrastructure, and information that is protected under regulation or law. Information that is determined to be sensitive must be protected through the implementation of security controls or mechanisms and through the enforcement of the organization's information sharing rules. Sensitive information can be protected through a variety of means, including:

- Authentication mechanisms that verify the identify of a user, process, or device through the use of usernames and passwords, cryptographic keys, tokens, biometric characteristics, or other authenticators.
- Encryption capabilities that protect sensitive data (including authenticators) by converting the plaintext information into ciphertext using a cryptographic algorithm.
- Authorization controls that grant access privileges to an authenticated user, program or process.
- Sanitization actions that remove, replace, redact, encrypt, or mask specific data elements.

When sharing incident and indicator information with peer organizations, sharing partners, or the public, an organization may wish to anonymize the data to some extent, depending on the context and agreed-to sharing arrangements. For phishing and other attacks, it is natural to look for instances of the targets' names, email or account names, in the body as well as the subject and attachments of the message. Organizations may also not wish to share the fact that they have been attacked, so reports may employ pseudonyms such as "USBUS1". If this is the case, then any artifacts of the attack, such as packet captures or files should be examined for revealing target IP addresses, domains, and URLs.

If sharing is a regular practice, then a review/release process should be established according to agreed-upon guidelines to mitigate inadvertent identity disclosures. When incident data contains PII, consult the

organization's privacy official to determine appropriate measures for redacting or anonymizing PII prior to sharing the information. Section 4.2.3 of NIST SP 800-122 provides guidance for anonymizing PII.[24]

4.2.2.2 Information Received from External Sources

In addition to the protections specified by governing law and regulation pertaining to privacy and other protected classes of information, an information sharing community may impose more restrictive terms of use on information shared within the community. These restrictions will vary by community, some being relatively simple and low-cost such as a verbal agreement to limit distribution of the information to the incident response team personnel within your organization; other agreements may be more formal and contain clauses enumerating specific obligations such as permitted/prohibited uses; ownership of intellectual property and community-submitted content; use of linkages or references to information; and obligations to outside organizations such as law enforcement or regulatory agencies.

Formal sharing communities generally employ a framework agreement that specifies the responsibilities of participants in both legal and technical terms. Such a community may rely on federally managed administrative systems for establishing trust, such as the federal system for the protection of classified information and the clearance processes that support it. For example, one way to share information pertaining to the protection of unclassified systems is to exchange possibly sensitive vulnerability and protected information from those systems using a separate, classified, network. Such formalized sharing relationships can achieve high levels of trust since the community-specific restrictions can dictate that community information be viewed and processed only by cleared staff and only on highly-protected systems. Some communities may also impose need-to-know rules, and require that a participant's incident coordination staff be individually authorized to access community information.

A somewhat less formal approach is to require that participants sign an NDA and that participant incident coordination staff hold clearances. In this context, information exchanged should be labeled with handling guidance, e.g., that the information should remain in the community, be released openly, or shared without source attribution.

A less formal approach is to require all community members to sign a memorandum of understanding (MOU), so that all participants can be considered to be trusted to the extent that they have agreed to the terms of the MOU, and then to use access control lists (or equivalent group-oriented mechanisms) to specify which community members should have access to specific messages shared with the community.

Some communities may also adopt an information sensitivity marking convention such as the US-CERT Traffic Light Protocol (TLP)[25] depicted in Figure 4-2.

[24] Another useful source for anonymization criteria can be found in the Health Insurance Portability and Accountability Act (HIPAA) regulations at §164.514(b). These criteria are only required under certain circumstances but are a useful set of criteria for other applications.

[25] Traffic Light Protocol, http://www.us-cert.gov/tlp

Figure 4-2: US-CERT Traffic Light Protocol

The TLP specifies a set of restrictions and a color code for indicating which restrictions apply to a particular record. In the TLP, red specifies the most restrictive rule, with information sharable only in a particular exchange or meeting, not even with a participant's own organization. The amber, green, and white color codes specify successively relaxed restrictions.

4.2.3 Storing and Protecting Evidence

As part of the information management, consideration should be given to how evidence is to be stored and protected. Basic questions to consider include:

- Is an appropriate backup policy in place and exercised?[26]
- Who is permitted access to the information?
- What qualifications will be required for system administrators that have access to the data? Background investigation? Citizenship?
- How long should the data be retained?[27]

Evidence should be collected and preserved using best practices for data preservation following chain of custody requirements, and other laws pertaining to the submission of evidence. A more detailed treatment of forensic techniques related to chain of custody and preserving information integrity are available in NIST SP 800-86 and section 3.3.2 of NIST SP 800-61, Revision 2.

Common security controls[28] should be employed where appropriate:

[26] NIST SP 800-34 Rev. 1, *Contingency Planning Guide for Federal Information Systems* provides guidance regarding IDPS principles and technologies.

[27] For federal agencies, National Archives and Records Administration (NARA) General Records Schedule (24) Item 7, "Computer Security Incident Handling, Reporting and Follow-up Records" requires that these records be destroyed/deleted 3 years after all necessary follow-up actions have been completed. Research conducted by Mandiant indicates adversaries have maintained access to victim networks for close to five years. The complexity of evaluating incident data and potential difficulties at connecting a series of related incidents that initially appeared unrelated, coupled with the potentially lengthy timeframes on which adversaries may operate signal the need to re-evaluate the 3-year retention period for incident handling data.

- Data in transit should be protected by encryption.
- Physical media such as CD's and DVD's should also be encrypted if that is the mechanism for exchanging data.
- Strong, two-factor authentication should be employed for portal or server access to data.
- Web portals and file servers should employ strong cryptographic protocols to provide communications security.
- Access to data should be logged and audited regularly.
- Intrusion detection should be deployed.[29]

Malware samples require special storage, access, and handling procedures. Malware samples are often preserved to support offline analysis and as evidence for an ongoing investigation or legal proceeding. Organizations often store not only the malware sample, but also accompanying metadata, artifacts, and analysis results. A malware sample that is not safely quarantined or sandboxed during unpacking and storage could propagate to enterprise networks and systems. Additionally, care must be taken to ensure that antivirus and anti-malware products do not inadvertently detect and remove an organization's malware collection. Common practice is to store malware samples in an isolated, protected file system or database as password-protected compressed files to avoid being inadvertently wiped by antivirus products during transit.

In the case of commercial threat intelligence services, the provider usually retains the rights to the intelligence collected at each customer point-of-presence and can use that information to improve intelligence and defenses. A threat intelligence sharing community may find that some members may wish to make use of the community's data for research or even product development. Each community should consider these data use cases when drafting their membership charter.

Organizations should determine the appropriate retention policies for information about attacks[30]. Multiple types of information with varying policies may be involved. There are motivations to retain detailed information for an indefinite period of time, since this provides historical value as well as helps new members or sharing partners understand the persistence and evolution of different adversaries. Other considerations, such as financial, legal, contractual, or regulatory, may require one to limit data retention to a fixed period of months or years. The retention policy for shared repositories should be determined by its members, in consultation with the appropriate records management personnel and legal counsel for each organization, and made explicit in any information sharing agreements. Once the retention schedule is satisfied, organizations must either archive or destroy the incident data in accordance with the applicable policies.[31]

For consortiums or organizations in specific industries or fields, there may be additional guidelines for storing and handling information. For example, organizations that are subject to HIPAA have

[28] NIST SP 800-53, Revision 4, *Security and Privacy Controls for Federal Information Systems and Organizations* provides a catalog of security and privacy controls and a process for selecting controls to protect organizational operations and assets from a diverse set of threats.
[29] NIST SP 800-94, *Guide to Intrusion Detection and Prevention Systems (IDPS)* provides guidance regarding IDPS principles and technologies.
[30] Federal agencies are subject to the National Archives and Records Administration (NARA) General Records Schedule as well as agency-specific retention policies.
[31] Draft NIST SP 800-88, *Guidelines for Media Sanitization*, provides guidance to assist organizations in making risk-based decisions regarding the sanitization and disposition of media and information.

requirements for safeguarding protected health information (PHI). If there are any discrepancies between the organization's obligation to protect certain information types and how that information is handled during the incident data sharing process, the key stakeholders and information owners as well as the organization's counsel must work collaboratively to identify the appropriate course of action.

An incident-coordinating or threat-sharing collaborative entity may well become a target of attack in and of itself. Therefore, measures should be taken to ensure that the infrastructure is adequately protected and monitored, that hosts and applications are maintained with current security patches and configurations, and that applications are free of common coding flaws[32].

4.2.3.1 Information Stored by a Community Portal

Some communities provide a portal that maintains stored information for sharing. For these communities, it is necessary for participant organizations to access the portal to find, analyze, download, and upload shared information. Access to a shared portal may be triggered by significant events, such as alerts, may be periodic, or both. It is important for organizations to carefully manage and protect all credentials used to access the portal, to clearly understand the notification mechanisms used by a community, and to regularly visit the portal to contribute content, download new information, and to participate in coordination activities within the community. Organizations should understand that interaction with a shared portal requires a level of ongoing effort.

Each community portal may implement a specific set of data access and retention policies. In order for organizations to have confidence that shared information is available and appropriately preserved, organizations should understand the access control policy of a shared portal and its data retention policies. In order for participants to trust a portal's ongoing availability and performance, a community should have a written SLA for the portal which specifies expected availability, the security posture of the portal, expected outages, acceptable usage policies, and any remedies for failure to perform.

4.2.3.2 Information Stored by an Organization

If an organization stores shared information on its own computers and networks, the organization should institute practices that minimize the likelihood of data loss, protect the data from unauthorized access, and provide mechanisms for search and analysis. During an incident, it may be important to access shared information quickly; consequently, the information should be available and readily accessible to authorized incident handling personnel. An organization should ensure that shared information is ensconced on systems that are well protected and available during an incident.

It is important to understand that shared information may be voluminous and that a storage system is required that can scale and that also provides for the confidentiality, integrity, and availability of the information. An organization should formulate a data retention policy for shared data that balances cost with the need to retain historical information. One possibility is to deploy a database system, within an organization's network, that uses replication to preserve shared information in the event of hardware failures, and to compress or reduce older records on a schedule.

[32] The NIST Software Assurance Metrics and Tool Evaluation (SAMATE) project seeks to develop standard evaluation measures and methods for software assurance. http://samate.nist.gov/index.php/SAMATE_Publications.html

4.2.4 Consuming and Responding to Alerts and Incident Reports

An information sharing community may send out alerts or incident reports to its members. An alert generally provides technical information that receivers can use to understand their degree of exposure to a particular vulnerability, the potential impacts of a problem (e.g., application crashes, data exfiltration, hijacking), and recommended steps to effectively mitigate the problem. An incident report documents a problem in greater detail and categorizes an incident by type. It is important to understand that both alerts and incident reports may contain sensitive information and may, if publicly disclosed, reveal to adversaries some of the defensive capabilities of members if the information sharing community. Incident reports in particular may contain sensitive information that should be shared only with community members with which a high level of trust has been established. In either case, a participant in an information sharing community must appropriately protect the information in an alert or report and must independently decide how to respond.

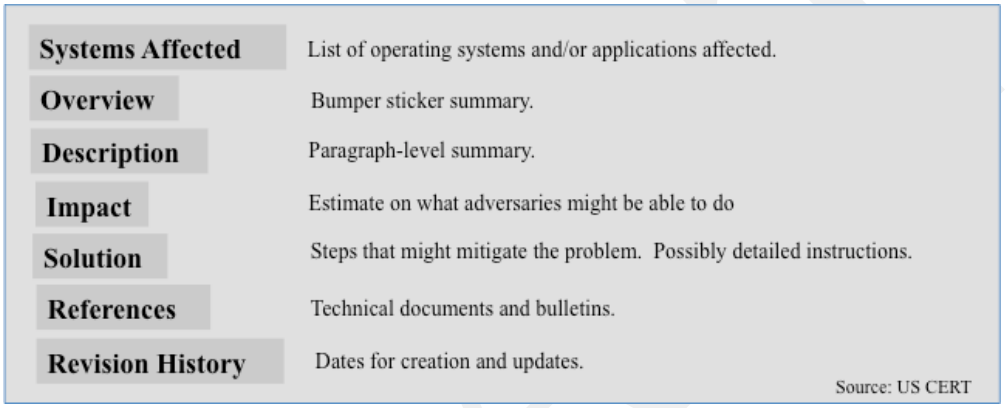

Figure 4-3: US CERT Alert

Figure 4-3 depicts an alert as documented by US CERT. This kind of alert identifies the types of systems that could be affected by a problem, provides a short overview of the nature of the problem, provides an estimate of the negative effects of the problem (e.g., system crash, data exfiltration, application hijacking),[33] possible steps to ameliorate the problem, and pointers to other sources of relevant information.

When an organization participating in an information sharing community receives an alert, the organization should evaluate how to respond based on the answers to six key questions.

1. *Does the alert apply to my organization's information technology assets?* An organization should compare the affected products identified in an alert with the information technology products deployed within their organization, preferably in an automated manner. If the organization does not use the products described in the alert, it may not be directly affected but it could still be impacted in unforeseen ways. If the alert applies to an organization's information technology assets, the remaining questions in this list should be considered.

2. *Are the suggested mitigations, if provided, both safe and effective?* An organization can approach this in two basic ways: (i) directly assess, analyze, and test the efficacy of the proposed mitigations, or (ii) if the source is deemed trustworthy and the suggested course of action seems viable, accept the mitigations as proposed. Organizations should consult multiple sources to arrive at an overall

[33] A more extensive list of potential effects is given in the MITRE Common Weaknesses and Vulnerabilities Types.

judgment about the accuracy of an alert and the technical competency and the degree of diligence demonstrated by the submitter and base any mitigation decisions on information that is well-understood and comes from a trusted source. Organizations should seek out personal connections with competent technical personnel within the community. These connections can be developed through participation in community events formed around common technical or research interests, at information sharing conferences, or through collaborative incident response. A source's past and ongoing participation in an information sharing community can also be used to gauge their reputation: Have the source's recommendations in the past proven to be both safe and effective? Do the alerts issued by the source display a high degree of quality and technical knowledge? Some communities have rigorous membership processes that require prospective members to be sponsored by a current member and demonstrate a high degree of technical competency. In such cases, membership in the community itself attests to the trustworthiness of the source.

3. *Does my organization have access to the skills to implement the mitigation guidance?* Performing the mitigation steps may require specific, and sometimes scarce, technical skills. Mature organizations may already possess these skills, but less capable organizations may not have personnel with the requisite skills. Improving technical skills through training or bringing on contracted staff with the appropriate skills and experience is a time-consuming process, it is therefore important for an organization to establish (perhaps contractual) relationships with an appropriate consulting entity or service provider who can respond quickly if needed. In the longer term, it is important for an organization to understand the skill sets that are needed to respond to the alerts and incident reports flowing through a community, and to develop or hire staff with the skills to meet these needs.

4. *What would be the costs of mitigation?* Mitigation strategies vary in their costs and impacts on an organization's ability to execute its mission. Some mitigation techniques, like filtering traffic from a specific set of IP addresses, are relatively low-cost and low-risk, but others, such as retiring vulnerable software versions, may be disruptive to implement. An additional consideration is the level of confidence that an organization has regarding the mitigation's effectiveness and side effects. A configuration change to a firewall, for example, may have unanticipated side effects to the mission. An organization should scrutinize mitigation techniques carefully, organize them using a change tracking process, perform pre-deployment testing when time permits, and preserve the ability to reverse mitigation techniques that turn out to be too costly or ineffective.

5. *Given my organization's mission and the possible infeasibility of mitigation strategies, should I perform mitigations at all?* When mitigation strategies cannot be realistically adopted because of cost or because the needed skills are not available, it may be necessary tolerate the additional risk posed by the problem described in an alert. An organization should consult the NIST Risk Management Framework (SP 800-37) for guidance on how to operate with known risks through maintaining a security plan and performing periodic security assessments to determine effectiveness of security controls. A supplementary strategy is to strategically reduce services where mitigation is difficult but where the mission can be achieved with reduced service levels.

6. *Is this alert associated with a campaign or wave of attacks?* An organization should evaluate the alert in the context of observed events, both current and historical. Through the analysis of information from local data sources and external sharing partners an organization may be able to correlate indicators; reveal meaningful patterns or sequences of indicators; or identify indicators that are common across multiple incidents. Organizations with advanced incident response capabilities may also be able to expose similarities in the adversary's TTPs; the specific types of organizations, systems, devices, or information targeted; or observe behaviors that are commonly exhibited by the adversary. When an analyst observes multiple incidents with the consistent appearance of specific indicators, TTPs, and behaviors within the attack lifecycle it is likely that the incidents are related and

possibly part of a larger campaign by an adversary. By shifting the focus from tactical detection and remediation (i.e., single event-oriented) to the detection of campaigns, network defenders can devise courses of action that prevent, or at a minimum, make it harder for the adversary to achieve their goals. When an organization is able to enrich the information it receives from its sharing partners (e.g., by identifying additional related indicators or behaviors) or through its analysis has reason to believe that a campaign or wave of attacks is underway, it should share this information with its partners if possible and appropriate. By sharing this information, the knowledge maturation cycle can continue, improving the overall fidelity of detection methods and related mitigation strategies.

Figure 4-4 depicts an incident report as described by US CERT. An incident report presents a more complete view of a problem. As shown in the figure, an incident report will generally characterize an incident by type, give a range of dates when it was active, provide source information, describe functional impacts, describe vulnerable system types, and summarize the impacts and resolution strategies. Much of this information may be very sensitive, and information-sharing communities tend to distribute incident reports primarily in trusted venues.

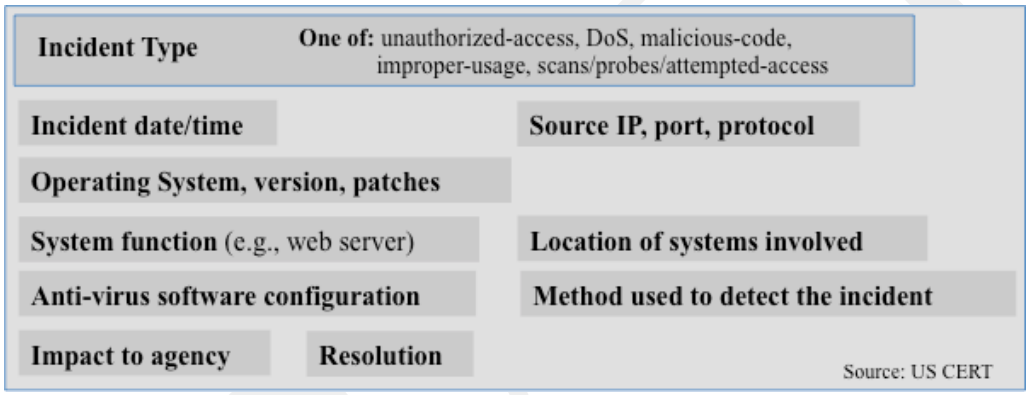

Figure 4-4: US CERT Incident Report

4.2.5 Consuming and Analyzing Indicators

A key aspect of consuming and analyzing indicators is that an organization must be able to monitor the same underlying observable events that are monitored and referenced in indicators by other participants in an information sharing community. If an information sharing community distributes an indicator about a particular set of observables, this will not help a receiving organization unless that organization can configure its systems to also monitor that set (or a significant subset) of observables. An organization should therefore, at a minimum, gain access to technical skills (either organization personnel or contractors) that are sufficient to configure event collection mechanisms as needed to monitor observables of interest to the community, and to perform a threat analysis of the observables to understand how they may relate to the organization's mission.

When receiving indicator from external data sources a series of activities are generally performed to ensure that the information can be efficiently put into use by the receiving organization. These activities may include categorization, initial prioritization, decompression, decryption, validation, and content extraction. Categorization requires a review of the content metadata to determine the security designation and handling requirements for the content received. Sensitive information may require encrypted storage, more stringent access control, or limitations on distribution. Content like malware samples or artifacts may require special handling precautions to prevent their inadvertent introduction on production networks. Initial prioritization ensures that newly received information is processed in the most

advantageous manner and may be based on the perceived value of the data source, the overall confidence level of the data, an operational requirement that specifies that data sources be processed in a particular order, the degree of preprocessing required to transform the data into actionable information, or other factors.

Analysis of indicators includes a broad range of activities that are focused on the rapid identification of malicious actors and actions within an organization's systems and networks. By integrating and correlating data from internal sensors (e.g., antivirus, IDS/IPS, DLP) and network monitoring systems with data received from external sources an organization can expose and characterize relationships between indicators that allow cyber defenders to more effectively identify an adversary's activities and behaviors and rapidly apply effective mitigations. Analysis activities can also include identifying patterns of attack or misuse, contextual analysis that considers the conditions under which a pattern is observed, and incident timeline reconstruction. Indicator analysis processes should inform the selection of courses of action, defensive measures, and mitigation strategies.

4.2.6 Creating Written Records

An organization should produce and maintain written records throughout the incident response lifecycle. The written record produced by an organization should be able to answer the following key questions:

- What happened? When did the incident occur?
- How was it detected?
- Who took part in the incident response? When were they notified?
- What actions were taken in response to the incident? What was the rationale behind these actions?
- What was the overall impact of the incident?

By answering these questions, an organization will be better able to reconstruct the timeline and narrative of the response activity. This documentation is much easier to produce at the time of the incident, while the details of the incident are fresh in the minds of the participants; important details are often lost when events are documented ex post facto. It is important to capture information regarding indicators; the TTPs used by the adversary; the types of systems targeted/affected; and possible adversaries. When documenting decisions, describe the deliberations that led to the final decision. Document the amount of downtime suffered, the recovery/restoration process, and describe the mitigation strategies employed or other courses of action. Be sure to collect, preserve, and safeguard as much information as possible – this information may be necessary to support future legal action, for termination/disciplinary actions for insider threats, or to shape incident response policies and procedures. Any information that could be used to better protect the organization (and its sharing partners) in the future should be captured.

An organization should produce an after-action report that captures lessons learned for each phase of the response cycle (e.g., a particular indicator that, if observed, would have allowed the organization to act sooner and perhaps disrupt or stop the attack earlier in the cyber attack life cycle). Use the lessons learned to identify opportunities for improvement – focus on identifying and addressing weaknesses that were exposed in the response plan. The after-action report is an opportunity to formally document what went well during the incident response, and what did not. Based on the lessons learned, implement any changes to policy, management, and/or operational practices that are necessary. These changes could include identifying supplemental information; personnel training; or other protective or detective measures that would have allowed the incident to be prevented, responded to more rapidly, detected earlier, or

recovered from faster. In the aftermath of an incident, the overriding objective is to prevent a similar incident from occurring in the future.

4.2.7 Performing Local Data Collection

Organizations that have the resources to monitor their systems and networks should identify and configure local data collection capabilities. Commonly available data sources include the log files and alerts generated by network devices, security appliances, operating systems, antivirus products, applications, and intrusion detection/protection systems. Local data collection entails more than just enabling logging on these various sources; logging parameters must be configured to capture those events and alerts that provide the most value to the incident responder.

When configuring log collection parameters, consideration should be given to the volume of data that a particular setting is likely to produce. Log configuration should be actively tuned to bring relevant events into sharper focus, remove "noise" (i.e., data with little or no practical value) from the channel, and ensure that the data collection strategy is not so aggressive that it creates a self-imposed denial of service. This tuning may include establishment of alerting thresholds; determining what actions/accesses will or will not be logged; and defining baselines for network activity, system configurations, and filesystem or registry objects. A significant consideration is also to ensure that logging errors are appropriately handled by defining how the logging system should respond when specific errors are encountered (e.g., can't complete a "write" operation because the disk is full or network connectivity has been lost).

Local logging and monitoring practices can be refined and improved upon based on input received from sharing partners, after-action reports, red team exercises, and by reviewing the alerts/events generated by an organization's own security scans. The frequency and/or scope of information collection may on occasion be temporarily increased (e.g., additional objects are monitored, more frequent measurement of network/CPU/disk utilization, both successful and failed object/service accesses are logged) in response to an active incident or to assist with fault detection, isolation, and correction during troubleshooting of networks and systems.

Threat sharing organizations collect threat intelligence from a variety of sources, including open source, internal malware repositories, and key external partners, easily collecting thousands of indicators in a short time. Inevitably, there is a need to store and organize this information into in some kind of structured knowledgebase. Free-form methods such as wikis can be quite flexible and suitable for developing working notes, while ticketing systems are good for tracking response activity. Some form of structured database is useful for organizing and tracking intelligence, and above all, querying and analyzing the collected threat information. An organization's collections or knowledgebase should pay particular attention to any TTPs regarding known adversaries that have been targeted by them.

Organizations typically collect the following items in a knowledgebase:

- Source of the indicator
- Rules (e.g., NDAs) governing the use of or sharing of this indicator
- When the indicator was collected by the organization
- How long the indicator is valid
- Groups or adversaries associated with the indicator

- Aliases of different adversaries or attack groups
- TTPs commonly used by the adversaries or attack groups
- Employees or types of employees targeted in the attacks
- Systems targeted in the attacks

It is often desirable to consolidate the log files from multiple sources to a centralized logging server or analytics platform such as a SIEM platform. Such aggregation, correlation, and analytics capabilities can be implemented using locally deployed hardware and software, or deployed in a cloud through various types of commercial service offerings. Section 3.4, presents specific considerations when evaluating and selecting commercial service offerings. The use of an analytics platform, depending on its feature set, can make it easier to correlate disparate data sources, perform offline analysis, support trending, and visualization. The ability to graphically depict data sets offers a unique perspective that may expose patterns of relationships among the data elements that might otherwise go unnoticed.

As part of the data collection process, an organization must also establish and implement a data handling and retention strategy. The data handling guidelines will specify the access control requirements for the log files; stipulate the rules governing data capture and acceptable use (e.g., avoid capturing sensitive data or PII); and protect log data at rest (e.g., both online and offline storage), in memory (i.e., by protecting the logging and analytics services), and in transit using end-to-end encryption where messages are encrypted by the sender and decrypted by the recipient with no third party involvement (e.g., PGP) or server-to-server encryption such as SMTP over Transport Layer Security (TLS) that uses Public Key Infrastructure (PKI) for encrypting messages between mail servers. The retention strategy will define the period of time the data will be retained, its storage method (e.g., online vs. offline), and how it will be safely and securely disposed of when it is no longer needed.

4.2.8 Producing and Publishing Indicators

An organization's information technology systems produce numerous observables; these observables include indicators such as: malicious email messages; IP address, domain, and URL watch lists; and file hash codes. Security software often generates observables in the form of log files. For example, NIST SP 800-92 "Guide to Computer Security Log Management" describes logs for intrusion detection and prevention systems, remote access software, web proxies, vulnerability management software, authentication servers, routers, and firewalls, as well as the use of non-security-specific log collection mechanisms, such as syslog, among others.

Indicators can be produced organically, thorough local data collection and analysis activities, or through maturation or enrichment of indicators received from sharing community partners. There are three basic types of indicators: atomic, computed, and behavioral.[34] Atomic indicators are simple data elements that cannot be further decomposed (e.g., IP address). Computed indicators are derived from other incident data (e.g., hash value). Behavioral are composite indicators, consisting of atomic and computed indicators joined through combinatorial logic and perhaps enhanced through the inclusion of contextual information. Organizations with basic network monitoring capabilities should be able to produce atomic indicators and perhaps simple computed indicators from existing data sources. The generation of sophisticated computed

[34] Amin, R., Cloppert, M., Hutchins, E. *Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains*, Lockheed Martin

indicators or behavioral indicators often require more advanced tools and analytical processes, and greater technical expertise.

When producing and publishing indicators, it is important to include metadata that provides context for the indicator, describing how it is to be used and interpreted, how it was observed, how it relates to other indicators. Metadata may also include handling instructions, sensitivity designations, and provenance information (e.g., what tool was used to acquire the data, how the data was processed, who collected the data). Publishers of indicators should also consider assigning a confidence level to the information that it intends to share. The confidence level represents the degree of certainty that the publisher asserts for a specific data element, relationship, or data set. Users of the information may take this confidence level into consideration when using this information as basis for decisions. As indicators are created, aggregated, or enriched their sensitivity and classification should be reevaluated; in some cases it may be necessary to sanitize the data or place restrictions on its use or dissemination.

While there is a need to provide information to sharing partners in a timely manner, it is equally important to ensure that any content that is published is known to be of good quality prior to publication; inaccurate or imprecise indicators may result in high false positives/negatives rates, disrupting response activities and adversely affecting an organization's reputation within a sharing community. Incident data that is shared should be managed through a version control system, whereby new or updated content receives a unique release number that allows it to be efficiently identified and retrieved. Incident data often has a "shelf-life" that consists of the period of time from the initial creation of the data and when it is no longer considered useful or relevant. Organizations that publish incident data should implement data aging procedures and algorithms that ensure that the published data is topical, timely, and accurate.

At times some information may be shared with a community that turns out, on closer investigation, to be incorrect, perhaps due to a cut-and-paste error, or typo, or some information that is sensitive may be inadvertently shared. Therefore, some mechanism for retracting submissions should be included in the community knowledgebase. These can be simply a communication to the administrator for manual removal or perhaps a programmed feature. Automated submission mechanisms require hardening to ensure that the feature does not become an attack vector for the adversary that allows them to mask their presence by modifying or deleting information. Organizations that share indicators should provide a feedback mechanism that allows sharing partners to submit error reports; suggested improvements; or additional information about the indicators. This feedback plays in an important role in the enrichment, maturation, and quality of the indicators that are shared within a community.

Some information shared with a community may be marked as "currently under investigation" and requires that members not share beyond the collective, and do no active investigation (such as collecting malware samples from a suspect website, or even performing a DNS lookup on a suspect host-name) that might tip-off a potential adversary or otherwise compromise the investigative activities. At some point, such information may be downgraded, once an investigation is concluded, so it is useful to have some mechanism to change the marking or add a revised marking such as "downgraded to GREEN as of 12/20/2015."

The use of standard data formats for the exchange of incident data enables greater interoperability and speed when communicating with sharing partners. Information is commonly exchanged in unstructured formats (e.g., text documents, email) that require manual processing and interpretation. The use of structured data supports the exchange of data with minimal or no human intervention (i.e., automated or "machine-to-machine"). When evaluating standard data format, look to formats that are lightweight and easy to implement; formats that are very feature-rich can also be exceedingly complex and difficult to use in practice. Choose formats that are widely adopted, readily extensible (i.e., new data elements or features

can be incorporated with minimal engineering and design effort), scalable, and provide the requisite security features to protect the data.

4.2.9 Producing and Publishing Incident Reports

Once an incident has been resolved, a final report should be produced that provides a summary of the incident, the ensuing investigation, the findings, and recommended improvements.[35] Incident reports help ensure that key decision makers are apprised of the incident and have the information necessary to make important operational decisions (i.e., those impacting the fundamental interests of the organization). Organizations should sanitize incident reports shared with an external partner by removing sensitive information or incident details that are not relevant to an external entity.

4.3 Maintaining the Sharing Relationship

Once sharing relationships are established, continued participation in the sharing community is essential for fostering stronger ties to other members and for the continuous improvement of incident response practices. Participating in community conference calls and face-to-face meetings increases an organizations ability to establish and cultivate trust with other members – a trust that may be a catalyst for a more free and open exchange of information, broader participation, and increased collaboration over time. Community-sponsored training events provide opportunities for less mature organizations to gain practical insights from seasoned incident response practitioners.

Organizations are encouraged to conduct after-action (i.e., hotwash) discussions and evaluations after an incident. In particular, it is helpful for an organization to review the value of external information sharing and collaboration efforts, identify opportunities for improvement (e.g., address data quality or latency issues), and draw attention to tools, techniques, or internal or external information or threat intelligence sources that can be used to counter similar threats in the future. The amount of post-incident analysis needed may vary based on the size, complexity, and impact of the incident. Shortly after an incident, the participants in an incident should meet to discuss the following types of questions:

- Did the organization gain any important threat intelligence and indicators (from external organizations) that assisted with the subsequent detection of the IT security incident?
- Did threat intelligence and coordination information from the external organization provide any countermeasures that the organization used to minimize the damage of the incident?
- Did threat intelligence received from the external organization result in the detection of false positives?
- Were the countermeasures employed effective?
- Were the countermeasures cost effective?
- If the organization shared internal incident information with external information sharing communities, was that information useful to the community?
- Did the organization sanitize the information that it provided to the external communities? Was the level of redaction performed appropriate? Was enough information released to be useful? Were organizational, legal, contractual, and ethical obligations regarding sharing met?

[35] Appendix B-Incident-Related Data Elements, of NISP SP 800-61 provides suggestions of what information to collect for incidents.

- If an incident caused damage internal to the organization, was that information shared with the external communities? If not, why not? How did the organization decide what damage information to share or not share?

- If the organization lacked threat intelligence and countermeasure information during the incident, are there external information sharing and collaboration communities that could have provided the information?

- Did the organization share technical information collected internally? If so, how much effort was expended to sanitizing the information?

The hotwash findings can be used by the organization to improve security measures, update policies and procedures, identify training needs, and to improve the organizational incident handling processes. An organization may also choose to selectively share relevant hotwash findings with their sharing communities to help improve the overall effectiveness of the community's incident response practices.

The ongoing maintenance of a sharing relationship requires that an organization's information sharing rules be reevaluated on a regular basis. Some of the events that can trigger the need to reexamine information sharing rules or practices include:

- Changes to regulatory or legal requirements
- Updates to organizational policy
- Introduction of new information sources
- Risk tolerance changes
- Information ownership
- Operating/threat environment
- Organizational mergers and acquisitions

4.4 Recommendations

The key recommendations presented in this section are summarized below:

- Define the overall goals, objectives, and scope of the information sharing initiative
- Obtain formal approval from the management, privacy officials, and legal teams and the support of key organizational stakeholders before sharing information
- Perform an information inventory that identifies the primary types of information that an organization currently possesses, the information owner, the sensitivity of the information, the protection requirements for the information, and the location of the information
- Enumerate risks of sharing incident and threat-intelligence data and identify appropriate mitigation strategies for each phase of the information life cycle
- Develop a process for reviewing and protecting data types and content that is likely to contain sensitive information

- Document the circumstances and rules under which information sharing is permitted by evaluating the risks of disclosure, the urgency of sharing, the trustworthiness of the information sharing community, and the methods used by the community to safeguard shared information
- Develop a list of data types and content, such as adversary indicators, that can be shared quickly with relatively minor review
- Identify peers and other organizations with whom coordination and information sharing relationships would be beneficial
- Ensure that the resources required for ongoing participation in a sharing community are available (e.g., personnel, training, hardware, software, and other infrastructure needed to support ongoing data collection, storage, analysis, and dissemination)
- Establish points of contact and engage in on-going participation with the sharing community through established communication channels
- Procedures for markup and data handling should be documented and approved by management.
- Mark, store and track information regarding the sensitivity of data to be shared.
- Protect sensitive information through the implementation of security controls, access control measures, and through the enforcement of an organization's information sharing rules
- Provide role-specific training to personnel so they understand how to handle incident and threat intelligence data appropriately
- Store and protect evidence that may be needed in the future; to help diagnose a future attack, or perhaps to support legal proceedings
- Implement the organizational processes, procedures, and infrastructure necessary to consume, protect, and respond to alerts and incident reports received from external sources
- Prepare for incident and threat-intelligence activities as much as possible in advance of needing to share in response to an actual incident.
- Implement the organizational processes, procedures, and infrastructure necessary to consume and analyze indicators received from external sources
- Document and use standard data formats and protocols to facilitate the efficient capture and exchange of information
- Produce and maintain written records throughout the incident response lifecycle, allowing the organization to later reconstruct the timeline and narrative of the response activity
- Produce and publish indicators based on local data collection and analysis activities, or through maturation or enrichment of indicators received from sharing community partners
- Produce and publish incident reports to provide initial notification of an incident, interim progress reporting during an incident, and a final report after the incident has been resolved
- Track sources of data and with whom that data has been shared

5. General Recommendations

The general recommendations presented in this document are summarized below:

- Establish and actively participate in information sharing relationships as part of a proactive, ongoing cyber incident response capability
- Exchange threat information, tools, and techniques with sharing partners – in doing so, an organization benefits from the collective resources and knowledge of its sharing peers and is able to better defend its networks and share costs
- Increase the organization's cybersecurity posture and maturity by enhancing or augmenting local data collection, analysis, and management functions. By implementing such capabilities, an organization gains a more complete understanding of its systems and networks, and is able to use a broader and richer set of information available through external sharing partners
- Use a cyber attack life cycle, such as the Lockheed Martin kill chain to define a framework for active defense that makes effective use of information available through both internal and external sources
- Share information about both attempted and successful intrusions. Often, information related to attempted intrusions is less sensitive and requires minimal sanitization and analysis; therefore it can be shared more quickly
- Carefully evaluate potential sharing communities/partners and select an information sharing model and community that is best suited for an organization or industry sector
- An organization should perform a self-assessment to determine if they have the capabilities to effectively engage in an information sharing community
- Ensure that a basic, foundational computer network defensive capability is in place before engaging in information sharing and coordination activities
- As a new entrant in an information sharing community, use information from external sources to enhance existing internal incident response capabilities
- Mature organizations should expand internal data collection operations, perform analysis, and begin to develop and publish indicators and actionable threat intelligence
- An organization may need to consider outsourcing incident response functions in cases where the personnel and skills necessary to perform a task are not readily available within the organization, or in cases where developing or maintaining a specific security capability in-house is not financial advantageous
- Before implementing an information sharing program, define its overall goals, objectives, and scope; obtain formal approval from the management, privacy, and legal teams; and acquire the support of key organizational stakeholders
- Perform an information inventory that identifies the types of information that the organization currently possesses, the information owner, the sensitivity of the information, the protection requirements for the information, and the location of the information

- Document the circumstances and rules under which information sharing is permitted by evaluating the risks of disclosure, the urgency of sharing, the trustworthiness of the information sharing community, and the methods used by the community to safeguard shared information
- Identify peers and other organizations with whom coordination and information sharing relationships would be beneficial
- Ensure that the resources required for ongoing participation in a sharing community are available (e.g., personnel, training, hardware, software, and other infrastructure needed to support ongoing data collection, storage, analysis, and dissemination)
- Establish points of contact and engage in on-going participation with the sharing community through established communication channels
- Protect sensitive information through the implementation of security controls, access control measures, and through the enforcement of the organization's information sharing rules
- Store and protect evidence that may be needed in the future; to help diagnose a future attack, or perhaps to support legal proceedings or disciplinary actions
- Implement the organizational processes, procedures, and infrastructure necessary to consume, protect, analyze, and respond to indicators, alerts, and incident reports received from external sources
- Produce and maintain written records throughout the incident response lifecycle, allowing the organization to later reconstruct the timeline and narrative of the response activity
- Produce and publish indicators based on local data collection and analysis activities, or through maturation or enrichment of indicators received from sharing community partners
- Produce and publish incident reports to provide initial notification of an incident, interim progress reporting during an incident, and a final report after the incident has been resolved
- Enumerate risks of sharing incident and threat-intelligence data and identify appropriate mitigation strategies for each phase of the information life cycle
- To the extent possible, prepare for incident and threat-intelligence sharing activities in advance of an actual incident
- Develop a list of data types and content that can be shared quickly with minimal review
- Develop a process for reviewing and protecting data types and content that is likely to contain sensitive information.
- Employ standard data formats and transport protocols to facilitate the efficient and effective exchange of information.
- Mark, store and track information regarding the sensitivity of data to be shared.
- Provide role-specific training to personnel so they understand how to handle incident and threat intelligence data appropriately.

Appendix A—Incident Coordination Scenarios

The scenarios presented in this appendix introduce real-world applications of threat intelligence sharing and coordinated incident response. These scenarios are meant to provide insights into how sharing and coordination can increase the efficiency and effectiveness of an organization's incident response capabilities. These scenarios seek to demonstrate that by leveraging the knowledge, experience, and capabilities of their partners, an organization is able to enhance its cybersecurity posture. These scenarios represent only a small number of possible applications of sharing and collaboration. The dynamic nature of the threat landscape means that as the tactics, techniques, and procedures of the adversary change organizations must adapt their protection, detection, and response strategies.

Scenario 1: Nation State Malware Attacks Against a Specific Industry Sector

A nation-state regularly targets companies in a certain industry sector for several months. The attacks come in the form of targeted emails that carry weaponized attachments containing a software exploit that, upon opening, launches malware on the victim's system. Once compromised, these systems contact servers controlled by the adversary to receive further instructions and to exfiltrate data.

The individual companies form a formal threat-sharing collective, where they establish a central forum to post information about different attacks. The posts describe details relevant to detecting and defending against the threat, such as the sender addresses of phishing emails, samples of malware collected from the attacks, analysis of exploit code used by the attackers, and IPs and URLs involved with the attacks.

As soon as one company's security team identifies a new attack, they quickly share the information with their peers. One company has advanced malware analysis capabilities and is able to extract additional information about the adversary and the infrastructure used for command and control from a malware sample collected by another company, and shared via the forum. By sharing the malware sample, the community is able to benefit from the malware analysis capabilities of one of its peers and to quickly and efficiently detect attacks that individually they likely would not have been able to find until well after the adversaries had penetrated their enterprises. In this scenario, an attack faced by one company becomes another's defense.

Scenario 2: Campaign Analysis

Cybersecurity analysts from companies in a business sector have been sharing indicators and malware samples in an online forum over the past few years. Each company performs independent analysis of the attacks and observes consistent patterns over time, with groups of events often having a number of commonalities, such as the type of malware used, the domains of command and control channels, and other technical indicators. These observations lead the analysts to suspect that the attacks are not fully random.

The forum members hold a technical exchange meeting to share data, insights, and analyses of the different attacks. What emerges from the combined data sets and joint analyses is the identification of several distinct sets of activities that are likely attributable to common adversaries or attacker groups, each with their own TTPs, target sets, and time table.

This scenario demonstrates how a broader set of data helps reveal collective action and campaigns by an adversary and the TTPs used by specific adversaries or campaigns.

Scenario 3: Distributed Denial of Service Attack Against Industry Sector

A hacktivist group targets a select set of companies for a large-scale distributed denial of service (DDoS) attack. The group employs a distributed botnet, loosely coordinated and controlled by members of the group. By analyzing the traffic generated by the botnet, one company is able to determine that the attackers are using a variant of a popular DDoS tool.

The targeted companies are members of an ISAC. Using the ISAC's discussion portal, the companies establish a working group to coordinate their efforts to end the attacks. The working group contacts the ISAC's law enforcement liaison, who coordinates with federal and international authorities to aid in the investigation and gain court orders to shut down attacker systems.

The working group contacts various ISPs, and provides information to aid in identifying abnormal traffic to their network addresses. The ISPs identify the source networks for the bulk of the traffic and are able to place rate limits on these sources, mitigating the attack. Using network traffic collected by the ISPs, international law enforcement agencies are able to identify the command and control servers, seize these assets, and identify some members of the hacktivist group.

After a technical exchange meeting among the targeted companies, several companies decide to enlist the aid of content distribution providers, to distribute their web-presence and make their business systems more resilient to future DDoS attacks.

Scenario 4: Financial Conference Phishing Attack

A cyber crime group made use of a popular business practices conference's attendee list to select targets for a wave of phishing emails. The group was able to identify multiple members of the business offices and, in some circumstances, compromise those machines and authorize electronic payments to overseas businesses.

One company identifies the attack against their business office employees and during their investigation realizes that the recipients of the attack email had all attended the same conference six months earlier. The company's CIRT contacts the conference organizers, as well as representatives from other organizations that attended the conference. A conference call is arranged to share information about the attack. Separately, two other businesses stop the attack, but are unable to identify the source. Three other businesses check their mail and network traffic logs and are able to identify potentially compromised hosts using the shared indicators.

The companies agree to share information about future attacks via an informal email list.

Scenario 5: Business Partner Compromise

"Company A" and "Company B" are business partners that have established network links between the organizations to facilitate the exchange of business information. A cyber crime organization compromises a server at Company B and uses their access as a stepping-stone to launch attacks against internal servers at Company A. A system administrator Company A notices the unusual activity and notifies their security team, who identifies the source of the activity as coming from a Company B system.

Company A's security team, who had previously engaged in a joint incident response exercise with Company B, contacts Company B's incident response team and describes the activity they are seeing. Company B's team is able to isolate the compromised server and perform an investigation to identify the source of the breach and other possible compromises.

The initial attackers had identified a weakness in a web-facing application and used it to take control of the server. Company B developers quickly fix the code to close the security hole and also enable additional logging and intrusion detection signatures to be ready for future attacks.

Company B's security team also determines that some customer personal information was potentially exposed to the attackers, so those customers are contacted and informed of the event, and instructed to change their passwords.

Because the security teams of the two companies had participated in a joint exercise, they had established contacts, built trust relationships, understood each other's networks and operations, and were able to quickly resolve the issue and prevent further damage from occurring.

Scenario 6: US-CERT Provides Indicators, Receives Feedback

The US-CERT receives information from a variety of sources that a number of servers located in the U.S. are being used to carry out cyber attacks against other U.S. firms. A specific foreign actor controls the compromised servers. The US-CERT identifies the firms under attack and notes that they are predominantly in the aviation industry. The US-CERT contacts the security teams of these firms and shares initial information about the attacks, including URLs, malware, and the kinds of vulnerabilities being exploited.

A number of the U.S. firms are able to identify and remediate attacks. These firms, during the course of their investigation, are also able to identify new indicators associated with the attackers that the US-CERT was unaware of. The US-CERT is able to share these new indicators with the rest of the firms, anonymizing the sources, leading to a more comprehensive response to the threat.

Scenario 7: A Retailer Fails to Share

A large retailer is subject to a cyber attack by a criminal organization. Millions of credit card numbers and account information of users are stolen during a breach that goes undiscovered for several weeks. The retailer does not participate in sharing threat information, so the organization relies on its own security and detection capabilities. Their internal capabilities prove inadequate in the face of a sophisticated, targeted threat that uses custom malware.

The breach is discovered by credit card companies investigating a rash of credit card fraud. The commonality in the credit card fraud was purchases made from this one retailer. The credit card companies notify law enforcement as well as the retailer, who begin an investigation.

The damages are enormous. The company notifies their customers of the theft of personal information, but does not release details of how the attack was carried out. Consequently several other retailers are successfully attacked by the same methods in the weeks following the initial breach. The financial losses realized by the retailers, customer, and credit card issuers could have been avoided, at least in part, had these companies engaged in active sharing of threat information

Appendix B—Glossary

Selected terms used in the publication are defined below.

Alert: Timely information about current security issues, vulnerabilities, and exploits. [SOURCE: US-CERT]

Computer Security Incident: See "Incident".

Computer Security Incident Response Team (CSIRT): A capability set up for the purpose of assisting in responding to computer security-related incidents; also called a Computer Incident Response Team (CIRT) or a CIRC (Computer Incident Response Center, Computer Incident Response Capability).

Cyber Threat Information: Information (e.g., indications, tactics, techniques, procedures, behaviors, motives, adversaries, targets, vulnerabilities, courses of action, or warnings) regarding an adversary, their intentions, or actions against information technology or operational technology systems.

Event: Any observable occurrence in a network or system.

False Negative: An instance in which a security tool intended to detect a particular threat fails to do so.

False Positive: An instance in which a security tool incorrectly classifies benign content as malicious.

Incident: A violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices.

Incident Handling: The mitigation of violations of security policies and recommended practices.

Incident Report: A written summary of an incident that describes the steps in the investigation of the event, the findings, and the resolution.

Incident Response: See "Incident Handling".

Indicator: An artifact or observable that suggests that an adversary is preparing to attack, that an attack is currently underway, or that a compromise may have already occurred.

Information Life Cycle: The stages through which information passes, typically characterized as creation or collections, processing, dissemination, use, storage, and disposition. [SOURCE: OMB Circular A-130]

Information Sharing and Analysis Organization (ISAO): An ISAO is any formal or information entity of collaboration created or employed by public or private sector organizations, for the purpose of—(A) gathering and analyzing critical infrastructure information in order to better understand security problems and interdependencies related to critical infrastructure and protected systems, so as to ensure the availability, integrity, and reliability thereof; (B) communicating or disclosing critical infrastructure information to help prevent, detect, mitigate or recover from the effects of an interference, compromise, or incapacitation problem related to critical infrastructure of protected systems; and (C) voluntarily disseminating critical infrastructure information to its members, State, local, and Federal Governments, or any other entities that may be of assistance in carrying out the purposed specified in sub-paragraphs (A) and (B).

Malware: A program that is covertly inserted into another program with the intent to destroy data, run destructive or intrusive programs, or otherwise compromise the confidentiality, integrity, or availability of the victim's data, applications, or operating system. [SOURCE: NIST SP 800-83, Revision 1]

Precursor: A sign that an attacker may be preparing to cause an incident.

Profiling: Measuring the characteristics of expected activity so that changes to it can be more easily identified.

Signature: A recognizable, distinguishing pattern associated with an attack, such as a binary string in a virus or a particular set of keystrokes used to gain unauthorized access to a system.

Social Engineering: An attempt to trick someone into revealing information (e.g., a password) that can be used to attack systems or networks.

Threat: Any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destruction, disclosure, or modification of information, and/or denial of service. [SOURCE: NIST SP 800-30, Revision 1]

Threat Source: The intent and method targeted at the intentional exploitation of a vulnerability or a situation and method that may accidentally exploit a vulnerability. [SOURCE: NIST SP 800-30, Revision 1 and CNSSI No. 4009]

Vulnerability: A weakness in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source. [SOURCE: NIST SP 800-30, Revision 1]

Appendix C—Acronyms

Selected acronyms used in the publication are defined below.

ACSC	Advanced Cyber Security Center
AI	Asset Identification
AMC	Army Materiel Command
APWG	Anti-Phishing Working Group
ARF	Asset Reporting Format
ASLR	Address Space Layout Randomization
CAPEC	Common Attack Pattern Enumeration and Classification
CCE	Common Configuration Enumeration
CCIPS	Computer Crime and Intellectual Property Section
CEE	Common Event Expression
CERT®/CC	CERT® Coordination Center
CFM	Cyber Fed Model
CIO	Chief Information Officer
CIRC	Computer Incident Response Capability (Center)
CIRT	Computer Incident Response Team
CISO	Chief Information Security Officer
CPE	Common Platform Enumeration
CSD	Computer Security Division
CSIRC	Computer Security Incident Response Capability
CSIRT	Computer Security Incident Response Team
CSOC	Cyber Security Operations Center
CVE	Common Vulnerabilities and Exposures
CVSS	Common Vulnerability Scoring System
CWE	Common Weakness Enumeration
CybOX	Cyber Observable Expression
DDoS	Distributed Denial of Service
DEP	Data Execution Prevention
DFIR	Digital Forensics for Incident Response
DHS	Department of Homeland Security
DIB	Defense Industrial Base
DLP	Data Loss Prevention
DNS	Domain Name System
DOD	Department of Defense
DOE	Department of Energy
DoS	Denial of Service
ENISA	European Network and Information Security Agency
ES-ISAC	Electrical Sector Information Sharing and Analysis Center
FIRST	Forum of Incident Response and Security Teams
FISMA	Federal Information Security Management Act
GAO	General Accountability Office
GFIRST	Government Forum of Incident Response and Security Teams
GLBA	Gramm-Leach-Bliley Act
HIPAA	Health Information Portability and Accountability Act
HTCIA	High Technology Crime Investigation Association
HTTP	HyperText Transfer Protocol
IC	Intelligence Community
ICE	Immigration and Customs Enforcement

ICS-CERT	Industrial Control Systems Cyber Emergency Response Team
IDMEF	Intrusion Detection Message Exchange Format
IDPS	Intrusion Detection and Prevention System
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IODEF	Incident Object Description Exchange Format
IR	Interagency Report
IRC	Internet Relay Chat
ISAC	Information Sharing and Analysis Center
ISAO	Information Sharing and Analysis Organization
ISC	Internet Storm Center
ISP	Internet Service Provider
IT	Information Technology
ITL	Information Technology Laboratory
MAEC	Malware Attribute Enumeration and Characterization
MOU	Memorandum of Understanding
MSSP	Managed Security Services Provider
NASA	National Aeronautics and Space Administration
NCCIC	National Cybersecurity and Communications Integration Center
NDA	Non-Disclosure Agreement
NERC	North American Electric Reliability Corporation
NIST	National Institute of Standards and Technology
NTP	Network Time Protocol
OCIL	Open Checklist Interactive Language
OMB	Office of Management and Budget
OpenIOC	Open Indicators of Compromise
OVAL	Open Vulnerability and Assessment Language
PCI DSS	Payment Card Industry Data Security Standard
PHI	Protected Health Information
PII	Personally Identifiable Information
PKI	Public Key Infrastructure
POC	Point of Contact
RCERT	Regional Computer Emergency Response Team
RFC	Request for Comment
RID	Real-time Inter-network Defense
SCAP	Security Content Automation Protocol
SOX	Sarbanes-Oxley Act
SIEM	Security Information and Event Management
SLA	Service Level Agreement
SMTP	Simple Mail Transfer Protocol
SOP	Standard Operating Procedure
SP	Special Publication
STIX	Structured Threat Information Expression
TAXII	Trusted Automated Exchange of Indicator Information
TLP	Traffic Light Protocol
TLS	Transport Layer Security
TSA	Transportation Security Administration
TTP	Tactics, Techniques, and Procedures
URL	Uniform Resource Locator
US-CERT	United States Computer Emergency Readiness Team
USACE	United States Army Corps of Engineers

USCYBERCOM	United States Cyber Command
VERIS	Vocabulary for Event Recording and Incident Sharing
XCCDF	Extensible Configuration Checklist Description Format

Appendix D—Resources

The lists below provide examples of resources that may be helpful in establishing and maintaining an incident response capability.

Incident Response Organizations

Organization	URL
Anti-Phishing Working Group (APWG)	http://www.antiphishing.org/
Computer Crime and Intellectual Property Section (CCIPS), U.S. Department of Justice	http://www.cybercrime.gov/
CERT® Coordination Center, Carnegie Mellon University (CERT®/CC)	http://www.cert.org/
Cyber Fed Model (CFM)	http://web.anl.gov/it/cfm/index.html
European Network and Information Security Agency (ENISA)	http://www.enisa.europa.eu/activities/cert
Forum of Incident Response and Security Teams (FIRST)	http://www.first.org/
Government Forum of Incident Response and Security Teams (GFIRST)	http://www.us-cert.gov/federal/gfirst.html
High Technology Crime Investigation Association (HTCIA)	http://www.htcia.org/
InfraGard	http://www.infragard.net/
Internet Storm Center (ISC)	http://isc.sans.edu/
National Council of ISACs	http://www.isaccouncil.org/
United States Computer Emergency Readiness Team (US-CERT)	http://www.us-cert.gov/
Defense Industrial Base (DIB) Cyber Security / Information Assurance (IA) Program	http://dibnet.dod.mil
Advanced Cyber Security Center (ACSC)	http://www.acscenter.org

NIST Publications

Resource Name	URL
NIST SP 800-30, *Guide for Conducting Risk Assessments*	http://csrc.nist.gov/publications/PubsSPs.html#800-30
NIST SP 800-34 Revision 1, *Contingency Planning Guide for Federal Information Systems*	http://csrc.nist.gov/publications/PubsSPs.html#800-34
NIST SP 800-37 Revision 1, *Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach*	http://csrc.nist.gov/publications/PubsSPs.html#800-37
NIST SP 800-39 Revision 1, *Managing Information Security Risk; Organization, Mission, and Information System View*	http://csrc.nist.gov/publications/PubsSPs.html#800-39
NIST SP 800-53 Revision 3, *Recommended Security Controls for Federal Information Systems and Organizations*	http://csrc.nist.gov/publications/PubsSPs.html#800-53
NIST SP 800-61 Revision 2, *Computer Security Incident Handling Guide*	http://csrc.nist.gov/publications/PubsSPs.html#800-61
NIST SP 800-83, *Guide to Malware Incident Prevention and Handling*	http://csrc.nist.gov/publications/PubsSPs.html#800-83
NIST SP 800-84, *Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities*	http://csrc.nist.gov/publications/PubsSPs.html#800-84

Resource Name	URL
NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*	http://csrc.nist.gov/publications/PubsSPs.html#800-86
NIST SP 800-88 DRAFT, *Guidelines for Media Sanitization*	http://csrc.nist.gov/publications/PubsSPs.html#800-88
NIST SP 800-92, *Guide to Computer Security Log Management*	http://csrc.nist.gov/publications/PubsSPs.html#800-92
NIST SP 800-94, *Guide to Intrusion Detection and Prevention Systems (IDPS)*	http://csrc.nist.gov/publications/PubsSPs.html#800-94
NIST SP 800-115, *Technical Guide to Information Security Testing and Assessment*	http://csrc.nist.gov/publications/PubsSPs.html#800-115
NIST SP 800-122, *Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)*	http://csrc.nist.gov/publications/PubsSPs.html#800-122
NIST SP 800-128, *Guide for Security-Focused Configuration Management of Information Systems*	http://csrc.nist.gov/publications/PubsSPs.html#800-128
NIST SP 800-137, *Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations*	http://csrc.nist.gov/publications/PubsSPs.html#800-137

2466

2467 **Other Publications**

Resource Name	URL
6 U.S.C., Sec. 131, Definitions	http://www.gpo.gov/fdsys/granule/USCODE-2010-title6/USCODE-2010-title6-chap1-subchapII-partB-sec131/content-detail.html

2468

2469 **Data Exchange Specifications Applicable to Incident Handling**

Title	Description	Additional Information
AI	Asset Identification	http://csrc.nist.gov/publications/ PubsNISTIRs.html#NIST-IR-7693
ARF	Asset Reporting Format	http://csrc.nist.gov/publications/ PubsNISTIRs.html#NIST-IR-7694
CAPEC	Common Attack Pattern Enumeration and Classification	http://capec.mitre.org/
CCE	Common Configuration Enumeration	http://cce.mitre.org/
CEE	Common Event Expression	http://cee.mitre.org/
CPE	Common Platform Enumeration	http://cpe.mitre.org/
CVE	Common Vulnerabilities and Exposures	http://cve.mitre.org/
CVSS	Common Vulnerability Scoring System	http://www.first.org/cvss/cvss-guide
CWE	Common Weakness Enumeration	http://cwe.mitre.org/
CybOX	Cyber Observable eXpression	http://cybox.mitre.org/
MAEC	Malware Attribute Enumeration and Characterization	http://maec.mitre.org/
MARF	Message Abuse Reporting Format	http://datatracker.ietf.org/wg/marf/documents/
MMDEF	Malware Metadata Exchange Format	http://standards.ieee.org/develop/indconn/icsg/mmdef.html
OCIL	Open Checklist Interactive Language	http://csrc.nist.gov/publications/ PubsNISTIRs.html#NIST-

Title	Description	Additional Information
		IR-7692
OpenIOC	Open Indicators of Compromise	http://www.openioc.org/
OVAL	Open Vulnerability Assessment Language	http://oval.mitre.org/
RFC 4765	Intrusion Detection Message Exchange Format (IDMEF)	http://www.ietf.org/rfc/rfc4765.txt
RFC 5070	Incident Object Description Exchange Format (IODEF)	http://www.ietf.org/rfc/rfc5070.txt
RFC 5901	Extensions to the IODEF for Reporting Phishing	http://www.ietf.org/rfc/rfc5901.txt
RFC 5941	Sharing Transaction Fraud Data	http://www.ietf.org/rfc/rfc5941.txt
RFC 6545	Real-time Inter-network Defense (RID)	http://www.ietf.org/rfc/rfc6545.txt
RFC 6546	Transport of Real-time Inter-network Defense (RID) Messages over HTTP/TLS	http://www.ietf.org/rfc/rfc6546.txt
SCAP	Security Content Automation Protocol	http://csrc.nist.gov/publications/PubsSPs.html #SP-800-126-Rev.%202
STIX	Structured Threat Information Expression	http://stix.mitre.org/
TAXII	Trusted Automated Exchange of Indicator Information	http://taxii.mitre.org/
VERIS	Vocabulary for Event Recording and Incident Sharing	http://www.veriscommunity.net/
x-arf	Network Abuse Reporting	http://www.x-arf.org/
XCCDF	Extensible Configuration Checklist Description Format	http://csrc.nist.gov/publications/ PubsNISTIRs.html#NIST-IR-7275-r4

2470

Appendix E—Change Log