Incident Management and Response Guide: Tools, Techniques, Planning, and Templates

By Tom Olzak, MBA, CISSP

Published by Erudio Security, LLC

Phone: 419-377-6844
Email: Tom.Olzak@v-cso.com

Web: v-cso.com

Section 1. Prepare

Section 1.01 Policy, Procedures, and Team
Section 1.02 Strategic Threat Intelligence
Section 1.03 Vulnerability Management
- (a) Unsecure Configuration and Coding
- (b) Training and Awareness
- (c) Access Control
- (d) Vulnerability Identification

Section 1.04 Section Summary

Section 2. Risk Management

Section 2.01 Risk Assessments
- (a) System Definition
- (b) Identify Existing Controls
- (c) Business Impact Analysis (BIA) and Calculating Risk
- (d) Risk Management Recommendations
- (e) Results Documentation and Presentation

Section 2.02 Section Summary

Section 3. Team Creation and Planning

Section 3.01 The Team
- (a) Computer Security Incident Response Team (CSIRT) Membership
- (b) CSIRT Responsibilities
- (c) CSIRT Response Tools and Resources

Section 3.02 The Plan
- (a) Step 1: Begin documentation and potential evidence preservation
- (b) Step 2: Determine if incident has occurred
- (c) Step 3: Prioritize the incident
- (d) Step 4: Report incident as specified in the incident response communications plan
- (e) Step 5: Obtain management decision about forensics preservation and collection
- (f) Step 6: Acquire, preserve, and document evidence as directed in Step 5
- (g) Step 7: Contain the incident
- (h) Steps 8 & 9: Eradicate the Incident and Recover
- (i) Step 10: Root Cause Analysis and Action Plan

Section 3.03 Section Summary

Section 4. Response

Section 4.01 Step 1: Begin documentation and potential evidence preservation
Section 4.02 Step 2: Determine if incident has occurred
Section 4.03 Step 3: Prioritize the incident and establish situational awareness
Section 4.04 Step 4: Report incident as specified in communications plan
Section 4.05 Step 5: Obtain management forensics evidence collection decision
Section 4.06 Step 6: Acquire, preserve, and protect evidence
Section 4.07 Step 7: Contain the incident
Section 4.08 Step 8: Eradicate the incident
Section 4.09 Step 9: Recover
Section 4.10 Step 10: Root cause analysis and reporting

Section 4.11 Section Summary

Section 5. Initial Response Forensics

Section 5.01 Forensics Overview
Section 5.02 Protecting Digital Evidence
Section 5.03 Securing a Potential Crime Scene
Section 5.04 Section Summary

Section 6. Works Cited

Figure 1: Risk Model
Figure 2: Attack Surface
Figure 3: Access Rights
Figure 4: Attack Tree
Figure 5: Controls Matrix
Figure 6: Qualitative Risk Calculator
Figure 7: Incident Handling Checklist
Figure 8: External Communication
Figure 9: VLAN Segmentation (Olzak, 2012(April))
Figure 10: Maximum Period of Tolerable Downtime
Figure 11: Dependent Processes
Figure 12: Root Cause Chain of Events
Figure 13: Five Whys
Figure 14: Incident Response Checklist
Figure 15: Digital Forensics
Figure 16: Initial Response Team Checklist

Section 1. Prepare

Incidents happen to us every day. We forget our password. One of our kids forgets their lunch. Our computer decides not to print. These are all small events that hinder our ability to move forward in our day. Security incidents are the same but usually have a greater impact.

A security incident is defined differently by various organizations. NIST defines an incident as "A violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices" (Cichonski, Millar, Grance, & Scarfone, 2012, p. 6). I find this too narrow.

In my experience, a security incident is an event, intentional or unintentional, that occurs outside what is expected in daily operations that can negatively affect business operation (processes), customers, investors, and employees. This expands the NIST definition by including anything that violates policy, regulations, laws, or ethics.

In other words, an incident is anything that can compromise the confidentiality, integrity, or availability (CIA) of data or the systems that support business processes. Confidentiality allows only authorized individuals or applications access to sensitive information. Integrity is the measure of the data's accuracy and authenticity. Availability ensures information is available to authorized entities when and where needed for business operation.

Incident response is a subset of an overall incident management program. The purpose of incident management is to prepare for various types of incidents and then respond when they occur. Incident management has four goals:

1. Development and management of an incident management policy and supporting procedures (details in Section 3)
2. Creation, training, and management of an incident response team (details in Section 4)
3. Preparation
    a. Strategic Threat intelligence
    b. Vulnerability management
    c. Risk management (details in Section 2)
4. Incident response to reduce or prevent business impact (details in Section 5)

Section 1.01 Policy, Procedures, and Team

The incident management policy forms the foundation for your organization's ability to prepare for and respond to the unwanted and unexpected.

An incident management policy template is available for download at http://bit.ly/2tTSKsg. The policy should create an incident management program and assign responsibilities for incident management and response.

In Section 3, I address creating the incident response team, plan, and procedures. For now, it is enough to understand the need for documented and up-to-date incident response procedures. You never want to face an incident without a clear approach to mitigating or preventing negative business impact.

Section 1.02 Strategic Threat Intelligence

Strategic threat intelligence (STI) provides your organization with information about probable threats and associated tools and techniques used by the threat agents. A threat agent is a specific incident of a threat. For example, a threat is potential for the theft of customer payment information by exploiting vulnerabilities. A threat agent would be a specific cybercriminal using certain tools and techniques to exploit weaknesses in your network to steal the information.

Without understanding how you might be attacked, it is impossible to perform comprehensive risk assessments. Information about potential threats and threat agents is available from

- Government and public sources
  - US-CERT Alerts (http://bit.ly/2pUj5oY)
  - The CyberWire (https://thecyberwire.com/)
  - Threat brief (http://threatbrief.com/)
  - Twitter feeds of top security professionals (http://bit.ly/2pWCG7u)
- Your vendors
  - IPS vendor
  - SIEM vendor
  - Threat analytics vendor
  - Microsoft
  - Apple

Section 1.03 Vulnerability Management

Managing vulnerabilities is ongoing. It allows us to identify and assess risk when associated with relevant threat agents. For example, we discover missing a patch during a vulnerability scan for Microsoft Windows that is currently exploited by one or more threat agents. Another example might be failing to block all nonessential SQL Server® traffic passing through a firewall or by unsecure configuration of VLAN access control lists. Vulnerabilities are usually caused by

- Unsecure configuration of operating systems, network devices, and applications
- Unsecure coding practices or developer mistakes

- Lack of user training and awareness
- Insufficient attention to authentication, authorization, and accountability in access controls

(a) Unsecure Configuration and Coding

Operating systems, such as Windows® and Windows Server®, have security baselines provided by Microsoft (http://bit.ly/2vcW6ML). Following these baselines is a good start. Network device vendors also provide guidance on how to securely configure their products. This guidance is also supported by security best practices, such as blocking everything on a firewall and opening only what is necessary for business operation. Cisco provides detailed information about hardening IOS devices at https://www.cisco.com/c/en/us/support/index.html.

Securely configuring applications and reviewing coding practices should not cause major concerns, if the System/Software Development Life Cycle (SDLC) minimally includes risk assessments and security requirements testing. For detailed information about integrating security into the SDLC, see *NIST SP 800-64 R2 Security Considerations in the System Development Life Cycle* (http://bit.ly/2kxni2y).

(b) Training and Awareness

Humans are the biggest vulnerability you face. Relying on user behavior to maintain confidentiality, integrity, and availability is a control of last resort: a control on which you should rely only when reasonable and appropriate technology controls leave gaps. Training and awareness activities, starting with a strong and communicated Acceptable Use Policy (download policy template from http://bit.ly/2pUtdOx), help to manage human vulnerabilities. For detailed information about developing and managing security training and awareness in your organization, see *NIST SP 800-50 Building an Information Technology Security Awareness and Training Program* (http://bit.ly/2qNzgII).

(c) Access Control

Controlling access to information resources is not easy. It requires reasonable and appropriate verification of any person or application attempting to access a resource (authentication). (The entity attempting to access a resource is called the subject, and the resource being accessed is called the object.) This is followed by authorization based on analysis of user roles to properly apply segregation of duties, need-to-know, and least privilege.

Segregation of duties prevents any single person from performing all tasks associated with a business process. Need-to-know ensures a person assigned a business role only can see the information necessary to perform related tasks. Least privilege limits what users in a role can do with data they access. Here is an example…

1. A user logs into the network and his/her identity is established (authentication)
2. The user is granted access to the payroll system because of his/her role (course authorization)
3. The user is granted access to specific tasks or data within the application, based on his/her role in the organization (fine authorization based on segregation of duties)
4. Once the user selects a specific task, he or she is only allowed to perform specific actions on the data (least privilege)
5. Database limits what the user sees to only what is necessary to perform an assigned task (need-to-know)

The final component of access control is accountability. Accountability ensures you understand what subject accessed an object, what was done to the object, and when the action happened. Collection of logs and log auditing is the foundation of accountability.

The strength of access control depends on the sensitivity of the resource protected: the resource's classification. We classify data based on its value to the organization and the negative impact on the organization if the CIA of that data is compromised. For example, we might classify data as

- **Public**: anyone can access and see the information with no negative impact on the business
- **Confidential**: moderate damage to the organization will occur if the data's confidentiality, integrity, or availability is compromised
- **Restricted**: severe damage to the organization will occur if the data's confidentiality, integrity, or availability is compromised

Any device through which data passes, is stored, or is processed is given the classification associated with the most sensitive classification of data involved. If, for example, a server contains restricted and public data (which is never a good idea), the server is classified as restricted. You should consider strong access control (multifactor authentication and encryption) for critical resources.

For a detailed discussion of access control, see *Identity Management and Access Control* (http://bit.ly/2q0unas). Download the University System of Georgia segregation of duties matrix template from http://bit.ly/2pXCeGF as sample tool for planning roles.

(d) Vulnerability Identification

You must know if you are open to attack. One of the best ways to do this is with regular vulnerability scanning. Nessus, for example, is an up-to-date tool widely used to scan networks for known vulnerabilities. A vulnerability management program also includes penetration testing and third-party security

program reviews. All of this begins with a vulnerability management policy and associated procedures (download policy template from http://bit.ly/2rjaikx).

A valuable tool for knowing what vulnerabilities you potentially have in house is the National Vulnerability Database (https://nvd.nist.gov/).

Section 1.04 Section Summary

The purpose of incident management is to prepare for various types of incidents and then to respond when they occur. Incident management has four goals:

1. Development and management of an incident management policy and supporting procedures
2. Creation, training, and management of an incident response team
3. Preparation
4. Incident response to reduce or prevent business impact

A security incident is an event, intentional or unintentional, that occurs outside what is expected in daily operations that can negatively affect business processes, customers, investors, and employees. It is anything that can compromise the confidentiality, integrity, or availability of data or systems that support business processes

Section 2. Risk Management

Managing risk is the first step in information assurance, and it is a critical piece of incident management. In both cases, risk assessments and subsequent risk acceptance, avoidance, transference, or mitigation are the foundation of preventing and responding to threats agents. If the incident response team does not run the organization's information risk management program, its members should at least be involved in every risk assessment. The formulaic risk model I use for our discussion of incident management related to human attacks is shown in Figure 1.

$$\text{Risk} = \frac{(\text{Means} + \text{Motive}) \times \text{Opportunity}}{\text{Controls}} * \frac{\text{Business Impact}}{\text{Incident Response}}$$

Figure 1: Risk Model

Section 1 explains threats and vulnerabilities. In Figure 1, the set of vulnerabilities available to enable an attack are categorized as opportunity. The probability that a threat agent can or will successfully take advantage of an opportunity to reach its objective is a key component of risk. Means are the skills necessary to successfully reach the intended target. A human threat agent is usually motivated by the financial, political, or other value of the attack target. Natural disasters need no motivation.

As the strength and tested effectiveness of controls increase, means and motivation must also increase. This serves to shrink the number of possible threat agents; probability of occurrence for human attacks tends to decrease. This decrease is caused by the increased effort (cost) to reach the target and the decrease in return on investment for the threat agent. Decrease in probability is also related to the difficulty something like a worm would have spreading across your organization and affecting availability, for example. If a threat agent's motivation is high, and she is highly skilled, a lower but still present probability of successful vulnerability exploits exists.

Once a threat agent gains entry to your network or one of your systems, potential for negative business impact arises. According to Gartner (2017), business impact includes "…the potential effects (financial, life/safety, regulatory, legal/contractual, reputation and so forth) of natural and man-made events on business operations" (para. 1)

How quickly we detect, contain, and manage an attack affects the extent of the impact. This is the purpose of incident response. If your organization has a documented incident response plan and a trained incident response team, you can prevent serious harm when the inevitable intrusion occurs.

As with all security activities, risk management begins with a management approved and supported policy. Shon Harris provides a great article about what goes into a risk management policy at http://bit.ly/2q9BgHB.

Section 2.01 Risk Assessments

Risk management helps prevent and prepare for incidents. The most valuable tool in this process is the risk assessment. A risk assessment looks closely at each system, the your network, and other organizations where your data is stored or processed. Perform risk assessments

- During the initiation and development/acquisition phases of the SDLC (http://bit.ly/2kxni2y)
- When deemed necessary by a Change Advisory Board (http://bit.ly/2f20um5)
- When new vulnerabilities are discovered in your systems or network, or when announced by a third-party
- When threat intelligence reveals a new threat, threat agent, or tools and techniques
- At least once per year for systems touching highly sensitive data or supporting critical business processes

An assessment consists of 10 steps divided into two phases:

Phase I: Assess

1. System definition
2. Threat identification
3. Vulnerability identification
4. Attack path controls assessment
5. Business impact analysis
6. Risk determination
7. Controls Recommendations

Phase II: Manage

8. Action plan and proposal creation and presentation
9. Implement approved controls or transfer risk
10. Measure to ensure steps taken work as expected and adjust where necessary

(a) System Definition

System definition begins with system decomposition. System decomposition breaks down a system into the various components of its attack surface (Olzak, 2011). A system is the collection of devices and media used to access, process, store, and move information for a related set of business processes. For example, the infrastructure supporting payroll processes is the payroll system. In some cases, you might want to assess only parts of the system. However, you should assess complete systems at least annually.

A system's attack surface is not a single piece. Instead, it is an aggregate of multiple attack surfaces. Figure 2 shows a very simple model. In this model, the network attack surface can be further broken down into each network device (switches, routers, firewalls, etc.) and cabling. The device attack surface includes the operating system and applications it hosts. I placed the device attack surface over the network attack surface because today's most popular and destructive attacks target users and their devices.

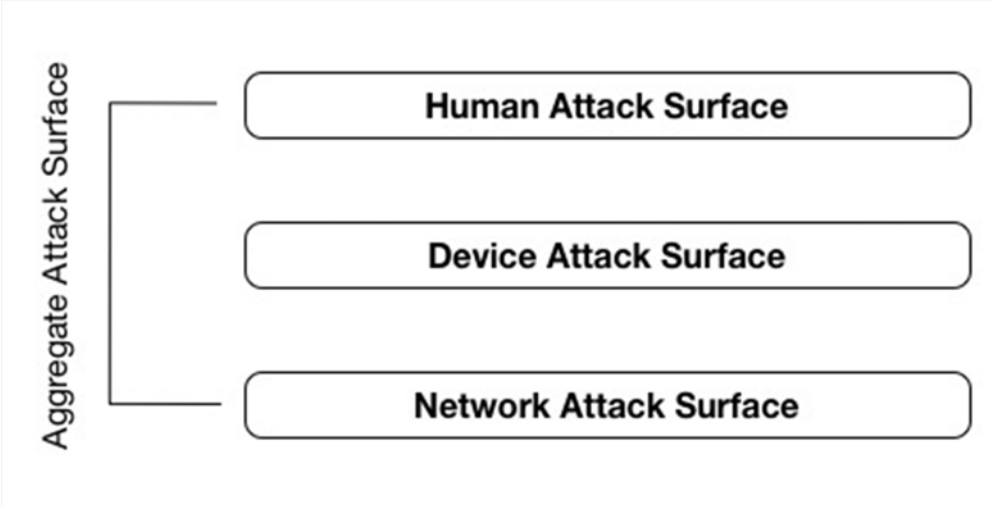

Figure 2: Attack Surface

When assessing attack surfaces, consider the following (Olzak, 2017)

- Entry points where the system receives information.
- Exit points where the system provides information to other systems:
  - Direct exit points exchange information with external systems.
  - Indirect exit points provide information to direct exit points.
- Data channels, protocol-enabled pathways over which information travels.
- Untrusted data items, persistent entities attackers use to control systems or extract data. Examples include cookies, files, malicious database records, and registry entries. Attackers cause exit points to read from untrusted data items or use entry points to write into untrusted data items (Manadhata, Karabulat, & Wing, n.d.). They are used by threat agents to own a device or system.

Protecting information transit points and channels; and defending against untrusted data items requires strong access rights between subjects and objects. See Figure 3.

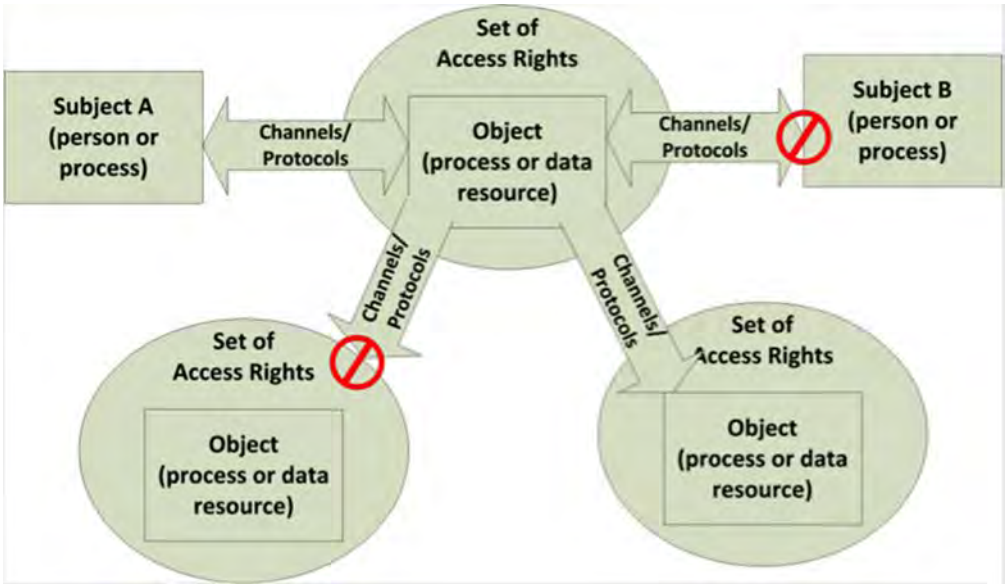

Figure 3: Access Rights

Any access between an object and a subject should be controlled with consistent rights management. "Access rights identify subjects, the objects they can access, and what they can do after access is granted" (Olzak, 2017). This does not just apply to users and the resources they access; it also applies to applications, services, protocols, and anything else that attempts to access an object for any reason.

Information about the system or network assessed can come from several sources:

- Existing documentation
- Interviews
- Questionnaires
- Network scans

(b) Identify Existing Controls

Identify existing controls and potential vulnerabilities by walking through probable attack paths using network and data flow diagrams to create attack trees. An attack tree helps visualize how a threat agent might gain access to an intended target. Figure 4 shows an attack tree with a database server as the target. This example does not show all possible attack paths. For a detailed description of how to use an attack tree, including addressing probability of successful attacks, see *Risk Management* (http://bit.ly/2rCPNM7).

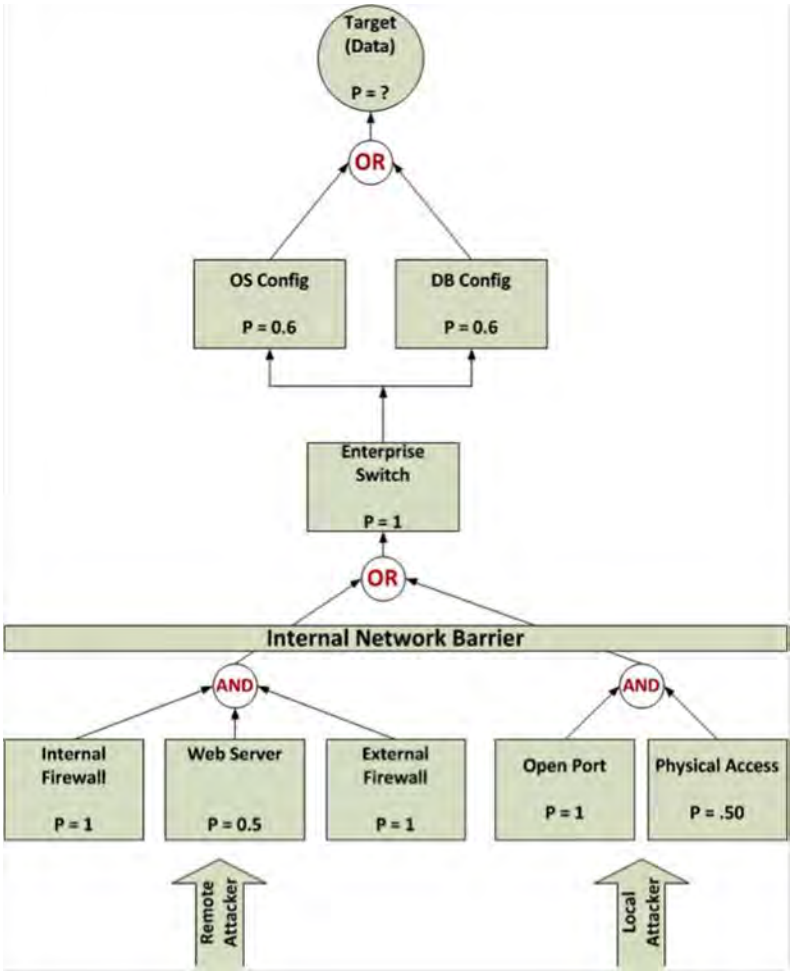

Figure 4: Attack Tree

In addition to attack trees, I recommend creating a controls matrix. A controls matrix lists all controls implemented, how they are configured, and what they protect. Figure 5 is a screen shot of a controls matrix template you can download from http://bit.ly/2pVWAV3. See *Use a security controls matrix to justify controls and reduce costs* (http://tek.io/2pWwnFA) for a detailed explanation on how to use the matrix.

Enterprise Security Controls Matrix September 30, 2016	Control															
Layer/Required Control	Endpoint Policy Manager	AV Client	IDP	Web Filter	Spam Filter	Email Encryption Solution	Firewalls	F5 Big IP	Group Policy Objects	Blackberry Enterprise Server	Policy	Layer 3 Switch	Aggregate			
Network Level Controls																
Data Center segmentation																
Intrusion detection/prevention																
Extrusion detection/prevention																
Rogue device detection/prevention																
Data leakage detection/prevention																
Isolated public ingress to Web servers and other outward facing services																

Figure 5: Controls Matrix

(c) Business Impact Analysis (BIA) and Calculating Risk

Use a BIA to determine the severity of the negative impact on a business if an incident occurs. Many variables affect business impact, including (Olzak, 2012)

- Maximum tolerable downtime (http://tek.io/2rDmgC5)
- Impact on employees
- Impact on investors
- Impact on customers
- Impact on current and future earnings potential
- Sanctions due to non-compliance with regulatory requirements

A BIA can be qualitative or quantitative. How you approach the BIA affects how you approach an overall risk assessment. A quantitative assessment uses actual dollar amounts to estimate business impact. A qualitative assessment uses some type of scale to estimate damage. Hybrid analysis is a combination of the quantitative and qualitative approaches. A qualitative risk calculator, downloadable from http://bit.ly/2pYNh6r, is shown in Figure 6. This calculator is just one approach to qualitative assessments, which are educated guesses based on experience and collaboration. For a detailed discussion of risk assessments, see *Risk Management* (http://bit.ly/2rCPNM7).

If you choose to download the calculator, the *System Sensitivity* cells are linked to a worksheet that calculates this value. The yellow column also contains a formula. Other worksheets provide guidelines for scoring the other columns. Change these to conform to your business operations, security framework, and management's appetite for risk. And remember, a threat agent usually must bypass two or more vulnerabilities to reach the target.

System Risk Calculator Date:

System: *[Enter system/resource name]* Team Members:

Threat Agent	Vulnerability	Likelihood		Business Impact			
		Means	Motive	System Sensitivity	Vuln. Severity	Controls & Response	Risk
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
				0			0
						System Risk Value	0

Figure 6: Qualitative Risk Calculator

Approaches to performing a BIA differ between organizations. However, we must always focus on the same things regardless of how our procedures look. According to Ross (2010), avoid the following 10 BIA mistakes:

1. **Considering the impact of interrupted applications, not business processes**. Remember, the impact is to business operations if a system is not available due to compromise or failure. Unavailability impacts business processes that feed and use the failed system. If you take your order entry system offline because of an attack, for example, no product ships. Customers are not happy, and revenue is lost.
2. **Considering applications in isolation**. Again, few applications operate in isolation. Most share information with other applications that enable multiple business processes. When performing a BIA for a system or a network device, look at all affected systems and related processes.
3. **Paying too little attention to financial impact**. Financial impact is a measure of how an incident affects your organization's bottom line on a profit and loss statement. This includes all costs, including

    a. Loss of short term revenue
    b. Regulatory fines
    c. Civil action by customers, shareholders, etc.
    d. Identity theft management
    e. Cost of recovery
    f. etc.

Costs associated with an incident must be calculated with the help of all affected areas of the business. This is necessary even if you use a qualitative or hybrid approach to your analysis.

4. **Paying too much attention to financial impact.** In addition to hard dollar costs, other costs affect the long-term health of a business following an incident, including loss of reputation and customer confidence; and loss of competitive advantage, especially when intellectual property is involved.
5. **Failing to distinguish enterprise applications.** Applications that serve the entire organization fall into this category. Examples include legal and document management systems.
6. **Failing to recognize data center applications.** Systems/solutions only used by IT are often ignored during risk assessments. Be sure you include these in your assessments.
7. **Confusing a risk assessment with a BIA.** A BIA is a subset of a risk assessment, but it can stand on its own. Even if you have no idea what might cause the unavailability of a system or business process, a BIA is something to consider: at least to establish value to the organization.
8. **Confusing risk acceptance with a business impact analysis.** Do not allow business managers to simply accept risk because they do not want to spend the time working with you to create a BIA. This is one more instance where support of C-level management for incident management is irreplaceable.
9. **Pre-determining BIA results.** Ross writes that a business manager can correctly estimate loss without a formal BIA about 80 percent of the time. This is the same as saying that one in five business processes or applications is inaccurately analyzed. Even when pursuing a qualitative analysis, it is important to take time to walk through estimated costs.
10. **Backing into a BIA result.** Sometimes, managers choose to understate the financial, reputational, and operational impact of an incident because the perceived impact is too high. This undermines the ability to effectively prepare for and manage incidents.

(d) Risk Management Recommendations

How you manage risk is largely determined by management's risk appetite: the level of risk managers are willing to assume to achieve business objectives. Part of creating an incident management program is meeting with the organization's business risk management team or senior management to understand acceptable levels of risk. This helps provide workable recommendations at this point in the risk assessment.

Once we know the risk, recommend one of the following:

- **Accept the risk.** If the cost of the risk is lower than any mitigation or transfer solutions available, we usually recommend risk acceptance.
- **Mitigate the risk.** If the cost of risk is higher than the cost of mitigation solutions, we usually recommend mitigation. Recommending mitigation requires a detailed analysis of our existing controls to determine if they can be reconfigured to reduce risk. It also requires analysis regarding how we might use fewer new controls by integrating them into the existing framework. In other words, never simply throw new controls at risk without a thorough analysis of what you have and what you need. Finally, any controls we recommend should be reasonable and appropriate for business operation.
- **Transfer the risk.** Transferring risk typically means purchasing incident loss insurance. Many insurance carriers now offer this. Purchasing insurance might be something done in addition to mitigation. For example, you might purchase insurance to cover costs associated with customer identity theft protection in addition to implementing additional technical controls. Together, transference and mitigation work to reduce risk to acceptable levels.
- **Avoid the risk.** Sometimes, risk is avoided by simply not doing something by removing existing procedures/technology or by not implementing a new solution. In my years as a director of security, management chose to avoid risk only once. Never count on avoidance. Our job as security professionals is to find ways to safely enable solutions that management deems necessary to reach the organization's objectives.

(e) Results Documentation and Presentation

Provide detailed documentation for how you conducted the assessment and your results. The details help the risk mitigation team be more effective. The NIST *Risk Management Guide for Information Technology Systems*, SP 800-30 (http://bit.ly/2rLdVfJ) provides an excellent template. However, details are something management usually does not care about. They only want to see the risks and what you believe needs to be done to manage the risks.

In addition to a detailed assessment document and a technical presentation, create a presentation for management. This presentation provides a high-level explanation of what you did and the risks discovered. At the opening of the presentation, let the attendees know you want them to decide on your recommendations. Be clear about how your recommendations are financially and operationally reasonable and appropriate.

The final document resulting from an assessment is the action plan. The action plan is the result of management's approval of your recommendations. It

includes what is to be done, who is responsible, and status. You can download from http://bit.ly/2rtKAtT the template I use.

Section 2.02 Section Summary

Incident management is inseparable from risk management. In addition to creating and practicing a response plan, the incident management team should be involved in every risk assessment. In my opinion, the team should manage the assessments as part of their day-to-day operations.

Risk is assessed by first understanding the system or network analyzed and then walking through all potential threat paths. This should occur when a new threat emerges or when new vulnerabilities are discovered. In any case, risk assessments for critical systems and sensitive data should happen at least annually.

Your risk management recommendations must be reasonable and appropriate for the organization's budget and operations. Management must see the short- and long-term financial and non-financial impact of simply accepting risk: or worse, doing nothing.

Section 3. Team Creation and Planning

In this section, I walk through details of creating and managing an incident response team and plan. The purpose of the plan is to

- Rapidly detect anomalous network, system, or device behavior (situational awareness)
- Minimize loss and destruction
- Mitigate exploited weaknesses
- Restore services
- Gather forensic evidence when reasonable and appropriate

Carnegie Mellon's incident response plan is a good start for any organization. It is available for download at http://bit.ly/2s7fCEn.

Section 3.01 The Team

Before planning starts, you need an incident response team. As I wrote in Section 2, this team is responsible for more than simply responding to incidents. It has a role in all risk management, incident prevention, and incident preparation activities. Consequently, the team makeup must include representatives from all technical teams, organization operations teams, and other relevant stakeholders.

(a) Computer Security Incident Response Team (CSIRT) Membership

The following list of team members is general and only a start. Each organization is unique, and the makeup of the team depends on whom must be involved to ensure effective incident management.

- Incident manager
- Security analyst
- Computer forensics investigator
- Server engineer
- Network engineer
- Server administrator
- Network administrator
- Business analyst for each department/line of business
- Software developer
- Data center operator
- Inside legal counsel
- Human resources
- Public relations

Depending on the organization, some of these members might be outside support vendors. All CSIRT members should participate in preparation and planning.

These team members serve as team leads in their respective areas. When an incident occurs, you will likely need more than one network engineer, for example. Also, consider training two individuals for each team role: a primary and a secondary. The primary might not always be available during an incident.

Identify a subset of the team as your initial responders. The initial response team, including an on-call responder, perform the first response steps as described later in this section and in Section 5.

One set of members, the business analysts, act as bridges between the CSIRT and the business departments and lines of business. In larger organizations, these positions already exist, providing day-to-day project and IT support functions to ensure technology effectively supports each department, lines of business, and overall tactical and strategic objectives. Business analysis are often missing in smaller organizations. In such cases, a representative from each department and each line of business is a necessary alternative. The business analyst or business representative is the point of communication between the CSIRT and the business. This is an irreplaceable and critical part of planning, preparation, and response.

Once an incident response policy creates the CSIRT, the team begins creating plans and procedures to meet its responsibilities.

(b) CSIRT Responsibilities

Many people believe the CSIRT sits around waiting for the next incident. Not true. The incident response team is responsible for

- **Risk management**. As shown in Section 2, the CSIRT is either directly responsible for managing information resource risk or provides support for those who are.
- **Incident prevention and preparation.** Conducting or participating in penetration tests and vulnerability management is a good start. The CSIRT should also be involved in the change management process. This ensures the risk management controls and procedures identified in the SDLC and risk assessments are maintained in a way that supports incident management.
- **New threat and vulnerability advisory distribution.** Threat intelligence and vulnerability research daily reveal new ways attackers try to attack your organization. The CSIRT is responsible for identifying new threats and vulnerabilities, performing analysis to determine associated risk to the organization, and distributing this information to appropriate IT and business teams**.**

- **Incident detection and response.** The CSIRT is responsible for monitoring for and assessing anomalous behavior of systems, devices, networks, and users. The CSIRT declares incidents when appropriate and executes the incident response plan.
- **Education and awareness.** Educating employees about the importance of safe use of information resources, policy compliance, and regulatory compliance should already be happening within your organization. However, many organizations do not address in training sessions what business employees, IT staff, and managers should do if they suspect an incident or if notified of one. This is a big miss. The CSIRT should manage security training and awareness or be directly involved in content and delivery, including how to report anomalous behavior.
- **Information sharing.** Whether an attack is successful or not, consider sharing all information gathered during initial and incident response analysis with both internal and external entities, including

    - Stakeholders
    - Regional and state law enforcement agencies
    - Federal agencies
    - Interest and industry groups

    In addition to incident information, share incident management findings about threats, risks, and other incident related. This allows a broad defense against threat agents.

(c) CSIRT Response Tools and Resources

Part of planning and preparing is putting together a set of tools and supporting resources that enable the CSIRT when an incident occurs, including a command center; jump kit; forensics lab (commonly outsourced); incident response forms with documented procedures and checklists; and external resource contacts.

(i) Command center

When an incident occurs requiring more than quick eradication and recovery, the CSIRT will gather in a central location for analysis, information sharing, and leadership. This command center is usually a previously designated conference room or training facility with minimally

- Whiteboards and markers
- Speaker phones
- Multiple tables for team and sub-team coordination and information sharing
- Hardwired connection to the internal network
- Isolated access path to the Internet for research, support, and reporting

The command center is the central point of response communication and operations. It is where the team and others will find the incident manager. It is also where all incident activity coordination and logging take place.

(ii) Jump kit

A jump kit is a forensics bag of tools a responder can quickly grab and head out the door. It should contain everything necessary for at least initial response evidence preservation, as described in Section 5, including

1. Journal for taking notes (who, what, when, where, how, and why) about every facet of the incident, including physical access
2. Contact list for all CSIRT members and external support
3. Up-to-date antimalware on USB drive or CD
4. Crime scene tape (http://amzn.to/2qgV1Nu)
5. Duct tape or other adhesive
6. Evidence bags (http://amzn.to/2rUBqTE)
7. Faraday bags for immediate collection of cell phones, tablets, and other wireless mobile devices (http://amzn.to/2qkFuuZ)
8. Evidence tags (http://amzn.to/2rAhwAK)
9. Chain of custody forms (http://bit.ly/2qkzr9K)
10. Digital camera with extra batteries
11. Sketch book with pencils and pencil sharpener
12. A laptop with an industry and judicially acceptable (stands up in court) forensics solution, such as EnCase (http://bit.ly/1SRrdxM)
13. Hard drive duplicator with write-block capabilities (http://amzn.to/2rAAJSX)
14. Miscellaneous cables, connectors, adaptors, etc.

The contents of your jump kit will vary from this list depending on whether your in-house team performs detailed forensics activities or whether you outsource them. At the very least, your kit should contain items 1 through 11 in the list above.

(iii) Forensics lab

Not every organization needs a forensics lab. I worked for a large organization, and we did not have one. Instead, we outsourced forensics analysis when needed. However, I provide a description of what a lab should include for those organizations deciding to retain this function in house. You can also use this list when assessing the credibility and effectiveness of a potential forensics vendor.

- Strong access control to the lab that minimally includes logging authorized personnel who enter and when
- A server for organizing and retaining investigation results (not connected to the Internet)
- A lab network isolated (preferably air gapped) from the organization's network with an Internet connection separate from the rest of the

organization and the lab administrative network (Internet connection should be only for administrative systems, never for systems used for evidence analysis or that are evidence themselves)

- Administrative systems for Internet access and lab management functions, connected to a network isolated from analysis systems
- Systems for analysis (virtual is a good idea) running various operating systems:
    - Windows desktop
    - Windows Server
    - Mac OS
    - Linux
- Drive duplicators with write blockers
- Readers for various types of media (e.g., SIMs and flash memory)
- Media wiping equipment
- Assortment of drive cables
- Miscellaneous cables and adapters
- Variety of drives of different types
- Accepted forensics software, such as EnCase and Forensics Tool Kit (http://bit.ly/2qnSYX6) running on non-admin lab systems
- Securable physical storage for separating and maintaining evidence chain of custody
- Video or audio equipment for recording findings, evidence, etc.
- Jump Kit (see *Jump kit* above)
- Certified computer forensics investigators

(iv) Procedures and checklists

Specific procedure content is unique to your organization, so I do not go into much detail. However, I provide an incident checklist (Figure 7) with recommendations for how to prepare for each line item. You can download the checklist from http://bit.ly/2qfUZtk. This checklist forms the basis for your response plan.

Step	Action	Completed	Responsible Person
	Detection and Analysis		
1	Begin Documentation and potential evidence preservation		
2	Determine if incident has occurred		
2.1	Analyze precursors and indicators		
2.2	Look for correlating information		
2.3	Perform research (e.g., search engines, knowledgebase		
3	Prioritize the incident (functional impact, information impact, recovery effort, etc.) and establish situational awareness		
4	Report incident as specified in communications plan		
5	Obtain management decision about forensics preservation and collection		
	Containment, Eradication, and Recovery		
6	Acquire, preserve, and document evidence as directed in Step 5		
7	Contain the incident		
8	Eradicate the incident		
8.1	Identify and mitigate critical exploited vulnerabilities		
8.2	Remove malware, inappropriate materials, untrusted data items, and other components across all affected devices, databases, etc.		
9	Recover		
9.1	Return affected systems to normal operation		
9.2	Confirm normal system operation and business process execution		
9.3	Monitor to ensure eradication was complete and critical vulnerabilities eliminated (continue situational awareness)		
	Post-Incident Activity		
10	Root cause analysis and create action report		

Figure 7: Incident Handling Checklist

Section 3.02 The Plan

Planning begins by working with all stakeholders to develop an overall approach to preparing for and responding to an incident. The discussion that follows is a general overview. Your plans should include various attack scenarios that affect how you approach planning and preparedness. Appendix A of the *NIST Computer Security Incident Handling Guide SP 800-61 r2* (http://bit.ly/1MYR74v) provides a good set of scenarios.

I approach planning by preparing to execute each of the 10 steps in the matrix in Figure 7. This ensures every step is thought through, documented, and practiced.

(a) Step 1: Begin documentation and potential evidence preservation

Provide the on-call responder with the means to immediately begin creating an incident log. This might be a document, spreadsheet, or other template already prepared and ready for use. Further, procedures and an associated contact list is necessary to begin preserving evidence in the data center, in the office, or at remote locations. Initial response evidence preservation requires training for business managers and IT personnel. You do not want the on-call responder to take time detailing preservation steps.

(b) Step 2: Determine if incident has occurred

Tools should be in place to enable immediate review of precursors and indicators. Precursors are log or other events that occur before an incident. They

provide insight into the potential for an attack. These might include social engineering attempts, phishing emails, unusual network or system activity, etc.

Indicators are evidence that an attack is in progress. Correlated log entries are a good way to identify indicator patterns of certain types of attacks, including unexpected movement of data, unexpected user access to resources, unauthorized log modifications, unusual/specific activity at the firewall or IPS, etc.

Crest provides a great document for how to configure and manage incident management logging at http://bit.ly/2qjOP7p. CrowdStrike provides a detailed look at indicators at http://bit.ly/2rUDYC8.

Once the responder gathers the precursors and indicators, he should research his findings using a knowledgebase or the Internet. Research sites should already be identified for quick access. This research provides insight into what is happening and next steps. Resources include

- Your antimalware, IPS, and SIEM vendors
- US-CERT (https://www.us-cert.gov/ncas)
- SANS Internet Storm Center (https://isc.sans.edu/dashboard.html)
- Fee-based cyberattack intelligence services

This step should be completed in minutes. The longer it takes to declare an incident, the larger the impact.

(c) Step 3: Prioritize the incident

Not all incidents are the same. Some might be remediated in minutes. Others might take days, and the potential impact across incidents differs. How to respond to each incident, or to multiple incidents at the same time, requires prioritization. Prioritization affects who is contacted and how response is initiated.

Your plan must include a quick guide for how to prioritize incidents. The on-call responder and initial response team must quickly assess the seriousness of the incident and, again, decide within minutes how to proceed.

Using a prioritization matrix is one approach. Use of a matrix begins with prioritizing the urgency and impact of the incident. Table 1 below is a template showing what this might look like (Wikipedia, 2017). This approach prioritizes an incident based on overall impact on the organization and how fast that impact might occur. You can download this template and the templates for Tables 2 and 3 from http://bit.ly/2qWEqkU.

Table 1: Prioritization Categories

Incident Urgency	
Category	**Description**
High (H)	- The damage caused by the incident increases rapidly - A minor incident can be prevented from becoming a major incident by acting immediately - Work that cannot be completed by staff will cause immediate harm to the organization
Medium (M)	- The damage caused by the incident increases considerably over time - Work that cannot be completed by staff will cause moderate harm to the organization
Low (L)	- The damage caused by the incident only marginally increases over time - Work that cannot be completed by staff will cause little or no harm to the organization
Incident Impact	
Category	**Description**
High (H)	- A large number of staff are affected or not able to do their jobs - A large number of customers are affected or acutely disadvantaged in some way - The financial impact of the incident is likely to exceed $9999999 - The damage to the reputation of the business is likely to be high - Someone has been injured
Medium (M)	- A moderate number of staff are affected or not able to do their jobs - A moderate number of customers are affected or acutely disadvantaged in some way - The financial impact of the incident is likely to reach $9999999 but not exceed $9999999 - The damage to the reputation of the business is likely to be moderate
Low (L)	- Few or no staff are affected or not able to do their jobs - Few or no customers are affected or acutely disadvantaged in some way - The financial impact of the incident is not likely to reach $9999999 - Little or no damage to the organization's reputation

Table 2 is the actual matrix used to determine the priority of the incident. Table 3 provides guidance on how quickly to respond and the expected recovery period. None of this information is likely to be a perfect fit for your organization. Adjusting the downloadable tables is the first step in integrating this into your response plan. The adjustment process requires close collaboration with business representatives and IT to ensure reasonable and appropriate response expectations.

Table 2: Incident Prioritization Matrix

		Impact		
		H	M	L
Urgency	H	1	2	3
	M	2	3	4
	L	3	4	5

Table 3: Incident Priorities

Priority Code	Description	Target Response Time	Target Resolution Time
1	Critical	Immediate	1 hour
2	High	10 minutes	4 hours
3	Medium	1 hour	8 hours
4	Low	4 hours	24 hours
5	Very Low	1 day	1 week

Finally, the matrix is a good general approach, but you will not always need it if certain types of incidents occur. Working closely with the business during response planning, you should quickly know when a response is critical because of the business services or processes affected. One or both of the following conditions will usually result in a high priority response (Wikipedia, 2017):

- Certain (groups of) business-critical services, applications or infrastructure components are unavailable and the estimated time for recovery is unknown or exceedingly long (*specify services, applications or infrastructure components, e.g., the customer facing order entry website is down*)

- Certain (groups of) Vital Business Functions (business-critical processes) are affected and the estimated time for restoring these processes to full operating status is unknown or exceedingly long (*specify business-critical processes, e.g. payroll during a payroll cycle*)

As part of this step, begin aggressive situational awareness activities. Situational awareness (SA) is the ability to understand the current state of a system and what has changed. SA is a continuous process supported by solutions like security information and event management (http://bit.ly/2qHKiyh); and identity governance and administration (http://bit.ly/2qa1cCR). Without SA, you can never quickly detect unwanted behavior and respond before your organization suffers serious damage, nor can you effectively manage an incident in progress.

(d) Step 4: Report incident as specified in the incident response communications plan

Once you determine an incident is in progress or has occurred, communicating what you know and what you are doing about it to the right people is important. Communication includes the rest of the CSIRT, previously identified managers, and external support organizations. Figure 8 (Cichonski, Millar, Grance, & Scarfone, 2012) depicts external entities normally included in a communication plan.

Figure 8: External Communication

How and when each of these entities is informed is up to your team's public relations (PR) representative and C-level management. As a responder, your responsibility should be to inform your PR team member and members of the management team listed in your communication plan. In addition, bringing in the necessary software and support vendors is incident manager's responsibility. When approved by PR or management, the CSIRT will communicate directly with external teams within guidelines documented in the communications plan.

Communication does not start with an incident. Rather, the CSIRT should have an ongoing working relationship with all outside entities as part of incident preparation. When contacted, external teams should already have familiarity with your organization and your team. They should have been involved in incident response exercises. No one should have to ask questions that are not specific to the incident and its characteristics. Again, time is critical.

Structured guidelines for creating an incident response communication plan are available for download from http://bit.ly/2qu6jwZ.

(e) Step 5: Obtain management decision about forensics preservation and collection

Most incident response guidance requires preservation of evidence. In my experience, this is a secondary consideration for management. What management wants is a quick return to normal operation while mitigating business impact. This does not mean you should not be prepared for evidence gathering, but there comes a point in the response when management should decide whether collecting evidence is more important than recovery.

As you read earlier, we immediately assume when an incident occurs that we must preserve all evidence. This mindset must continue until management decides otherwise. Include in incident planning what is needed to understand what happened and how without major recovery delays.

Continuous, comprehensive logging; event correlation; and retention and protection of the results usually provide what we need at the post incident root cause and action plan creation meetings. The information also provides first steps for law enforcement if a path to prosecution is taken.

In addition to logs, we can also seize relevant user devices without a significant delay in recovery. If we virtualize our servers, previous response planning can allow isolation and preservation of the incident-related servers while bringing up replacement virtual servers to restore business operations.

If we think through all probable scenarios during planning, evidence collection is often possible with what management might consider reasonable impact on the bottom line. Be sure to include forensics considerations in your preparation activities.

(f) Step 6: Acquire, preserve, and document evidence as directed in Step 5

If your CSIRT has its jump kit and internal/external forensics lab, it is ready to take on this step. Also, part of planning is ensuring your forensics investigators are certified and able to collect, analyze, and protect evidence so the results stand up in a court of law. See Section 5.

(g) Step 7: Contain the incident

Containment is the most important part of loss minimization and evidence preservation. For physical attacks, this translates into delaying an attacker long enough for law enforcement or other human intervention. Containment for logical attacks requires isolation of the affected systems and network segments. Isolation protects uninfected systems during malware attacks and helps prevent a cybercriminal from extracting data during a breach. It also helps prevent unwanted alteration of digital evidence.

Although containment involves processes unique to each incident, the overall approach to containment is strategic. It is "…a function that assists to limit and prevent further damage from happening along with ensuring that there is no destruction of forensic evidence that may [*sic*] be needed for legal actions against the attackers" (InfoSec Nirvana, 2015). Using scenario planning, assess the need for containment, how containment is achieved, and what you must do prior to an incident to prepare.

(i) Physical incidents

I do not spend much time on physical incidents in this guide. However, a brief look at physical incidents is important. Sometimes, a physical intrusion precedes a logical attack. Also, a device on your network might be used to launch or further an attack against your organization or one remote on the Internet.

The purpose of physical security is first to deter intruders with fences, guards, signs, etc. Second, we delay intruders by placing layers of barriers between them and the targets. Examples of barriers include gates, fences, walls, and locks. The length of required delay depends on the response time for arresting or otherwise intervening to stop the intrusion.

Barriers alone are not enough. SA also applies to physical attacks. Alarms, cameras, and other types of sensors help track and apprehend an intruder. Also, containing a crime scene and related evidence is necessary if management decides to prosecute. If a device or system is used for an attack or is the target of an attack, ensure barriers (crime scene tape and human oversight) prevent access by anyone not directly involved in the response process.

For a detailed look at physical security for protecting your information resources, see *Physical Security: Managing the intruder* (http://bit.ly/2q9I7AV).

(ii) Logical incidents

Containing logical incidents requires addressing isolation alternatives during the SDLC and all risk assessments. If we have not planned for isolation, we end up running through the data center unplugging cables (hopefully labeled cables), hoping for the best. If you have ever done this, you know isolation is iffy and recovery can take longer.

One of the most effective methods of isolation is use of VLANs. In addition to controlling day-to-day access, VLANs provide the segmentation and device isolation needed to prevent, deter, and contain an attack. See *VLAN Network Segmentation and Security* (http://bit.ly/2ggAuVA).

Figure 9 shows a network segmented with VLANs. All database servers are on a single VLAN, with users and application servers on another. All external traffic arrives and exits on other VLANs. VLANs are configured to prevent devices on the same VLAN from communicating with each other unless explicitly allowed, so some isolation is already built in. This is a simple example. In the real world, I would likely separate sensitive data, public data, control data, and different business processes onto different VLANs.

Figure 9 shows how easy it would be to isolate various segments of the network by reconfiguring one or two switches. Segmentation is also possible using routers in addition to switches. Segmentation strongly supports containment whether you operate in a traditional, virtual, or hybrid environment.

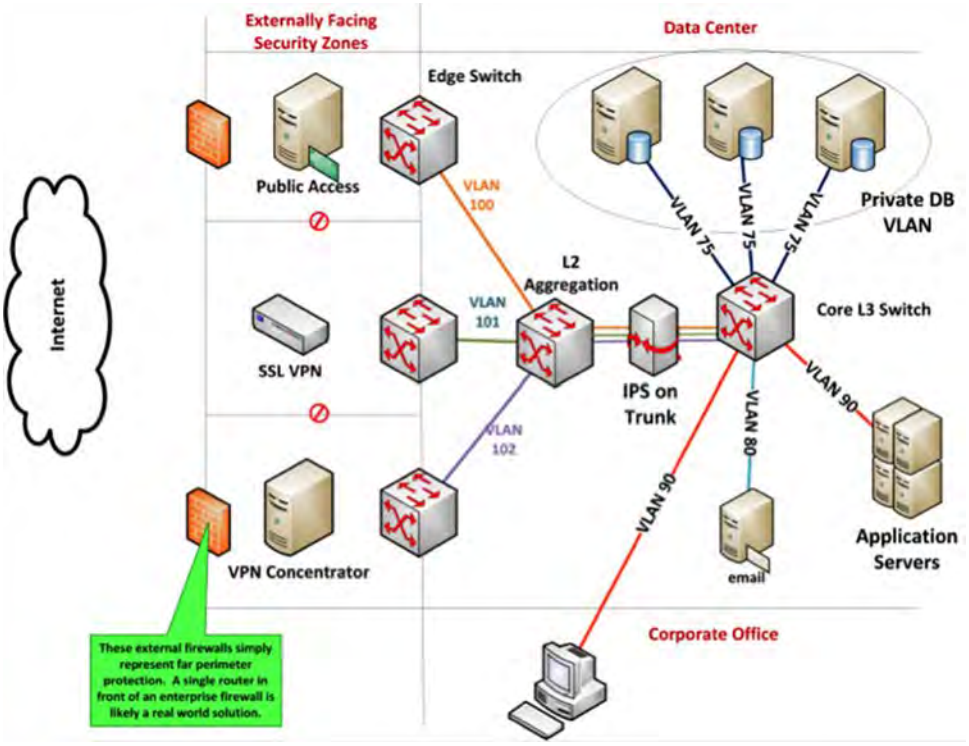

Figure 9: VLAN Segmentation (Olzak, 2012(April))

Quick containment using network devices requires preconfigured reconfigurations stored and easily accessed by the response team. This allows for rapid isolation once the threat agent's actions are analyzed. One way to ensure fast reconfiguration across all relevant devices is with a software defined network solution (http://bit.ly/2q7lwsq).

For end-user devices, the most effective isolation approach is unplugging them from the network. Place cell phones, tablets, and other mobile cellular access devices in Faraday bags. Do not power them off. Have a plan in place to block affected user devices from connection to wireless access points if too large to place in Faraday protection.

Your containment approach should enable the CSIRT to stop data extrusion or the spread of the attack quickly and as narrowly as is reasonable and appropriate for business operations and risk. How you do this is unique to the combination of your technology, your budget, legal ramifications, and management's appetite for risk.

(h) Steps 8 & 9: Eradicate the Incident and Recover

Steps taken to eradicate an incident depend on the type of incident and the tools and techniques used by the attacker. Scenario planning and comprehensive threat intelligence ensure you identify all malware, untrusted data items, inappropriate materials, unwanted registry entries, etc.

(i) Eliminate exploited vulnerabilities

The first step in eradication is making sure the incident does not happen again in the same way. Achieving this requires elimination of the vulnerabilities exploited by the threat agent. Identification of these vulnerabilities should be apparent through established threat intelligence research and results from your log management solution.

An expeditious process for fixing vulnerabilities is already part of a well-designed change management process. Patches and reconfiguration of controls or systems are quickly assessed, documented, and applied without going through complete change management sign off.

If you make a change that does not work as expected, reverse it before trying something else. Do not throw multiple changes at the incident without analysis of what works, what does not work, and removal of anything no longer needed. Otherwise, you will not know what actually saved your organization. Further, post-incident cleanup will take much longer than necessary.

(ii) Remove the unwanted

The most effective, and often quickest, approach to eradication on user devices is a complete wipe and reinstall. In many incidents, this is the only way to be certain all unwanted entities are removed from affected devices. Planning for this requires creation of user device images, including different configurations based on line of business, department, business role, etc. Image creation is part of incident planning and preparation. Initially, this can be very time consuming. Once done, however, including image management in the change management procedures makes keeping images up to date relatively easy.

Using server images is also effective, but using virtualized servers or containers for your critical servers is often a better option. Bringing up a virtual machine to replace a server isolated in an incident quickly achieves both eradication and recovery for that server and supported business processes. Before placing new servers in production, be sure to eliminate any identified vulnerabilities found in the compromised servers.

Finally, we need to ensure data integrity and the absence of untrusted data objects in our databases and on our file servers. The first preparation step is prevention of integrity compromise with strong authentication and authorization controls. Next, back up and back up often. Your back up timeline should represent the longest you want to be down following any kind of business continuity event (http://bit.ly/2rakdXj). Protect backups from any type of incident that might occur at any facility. This usually means retaining them off site or in the cloud.

Cloud database back up services, like those provided by Microsoft (http://bit.ly/2rcjXJ4), Oracle (http://bit.ly/1N3eQ3j), and other cloud service

providers enable both database and flat file backups that provide data integrity and reasonable recovery times. Solutions like Carbonite (http://bit.ly/2bj85fH) can protect even the smallest business with offsite, protected data. However, faster recovery times for your most time sensitive business processes might require maintaining synchronized database servers at a colocation (co-lo) or in the cloud.

A co-lo is a data center placed at least 25 miles from your primary data center that contains infrastructure supporting your critical business processes. For disaster recovery purposes, your co-lo and data center should be in different power grids, flood plains, weather corridors, etc. For attack purposes, the connection between the co-lo and data center must be tightly controlled. Consider allowing no remote user access unless the data center or one of the critical systems becomes unavailable.

Synchronize data between the co-lo and data center so that a simple change to DNS, VPN, or other remote access methods allow customers and remote sites access with little interruption in service delivery. Further, office staff at the data center location must have a way to easily access the co-lo servers.

Always assume your data center compromise can easily pass to your co-lo. SA for your co-lo is also necessary.

When wipe-and-replace or redundant systems are not available or possible, removal of unwanted items requires research into what happened, tools used by the attacker, actions taken by the attacker, files and executables installed, and any other changes made to registries, configuration files, etc. Once you complete this time-consuming process, the CIRT must create a procedure and associated tools to reverse all attacker actions. Team members and other recruited IT personnel must then follow the documented procedure to eradicate the threat. Unless the attack scope was very small, this approach might extend recovery time beyond one or more business process maximum tolerable downtimes (http://tek.io/2rDmgC5).

(iii) Recovery

Recovery is focused on returning business processes within MTDs defined in BIAs. Recovery time includes the time necessary to restore the infrastructure and the time needed to rebuild data sets. This is also known as the maximum period of tolerable downtime (Olzak, 2013), as shown in Figure 10. The RTO (recovery time objective) is how long it takes to restore supporting technology.

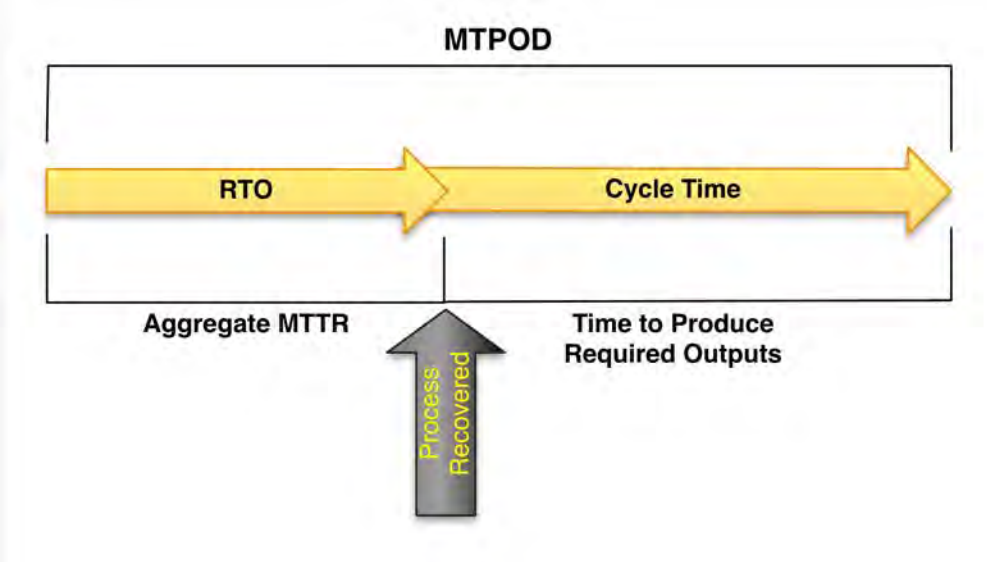

Figure 10: Maximum Period of Tolerable Downtime

We previously looked at eradication methods that also begin the recovery process. Solutions like a co-lo or a cloud-based redundancy solution can quickly return business processes to normal. Other approaches take more time. Regardless of how you approach eradication and recovery, be sure to work with management to understand your recovery time options. This includes considering the MTDs of processes affected by a downed system: both upstream and downstream. See Figure 11 (Olzak, 2013).

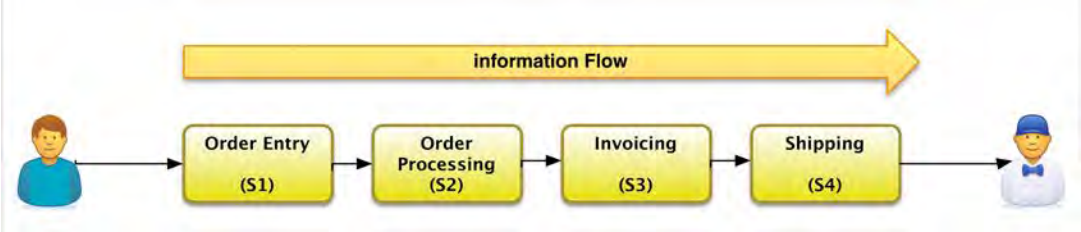

Figure 11: Dependent Processes

Failure of any one of these processes breaks a chain required to provide product to customers. The MTD for any process in this chain is the shortest MTD across all processes: the chain's MTD.

Once you recover systems, work with the business to confirm correct operation. Check not only whether the technology works as expected, but also ensure data integrity. Having predefined reports to validate data accuracy is one way to quickly do this.

(i) Step 10: Root Cause Analysis and Action Plan

The last step in incident response is ensuring the same incident does not happen again in the same way. Also, you want to assess how well your team responded and whether you can improve; there are always opportunities for improvement.

Although you should have already blocked one or more of the vulnerabilities exploited in the current incident, you need to understand why those vulnerabilities existed and the failure of controls to detect and stop the attack. Root cause analysis is the primary tool for this.

Root cause analysis finds the fundamental, the root, causes of any event. It prevents treating symptoms. Treating only symptoms will not effectively prevent future, similar incidents.

Cause analysis begins with bringing together everyone involved in the incident and with the systems affected. The resulting meeting must prohibit finger pointing and assigning blame. That is not the purpose of the meeting. Trying to place blame causes attendees to get defensive and lose objectivity.

Root cause is found by following the chains of cause and effect leading to the incident, as shown in Figure 12 (Olzak, 2008). In many instances, more than one root cause exists. Analysis begins with the proximate cause and works back to the root causes. A proximate cause is the event and surrounding conditions that enabled the incident. The process used for this step-back process varies between organizations. I used two different methods, but I found the five-whys approach worked best.

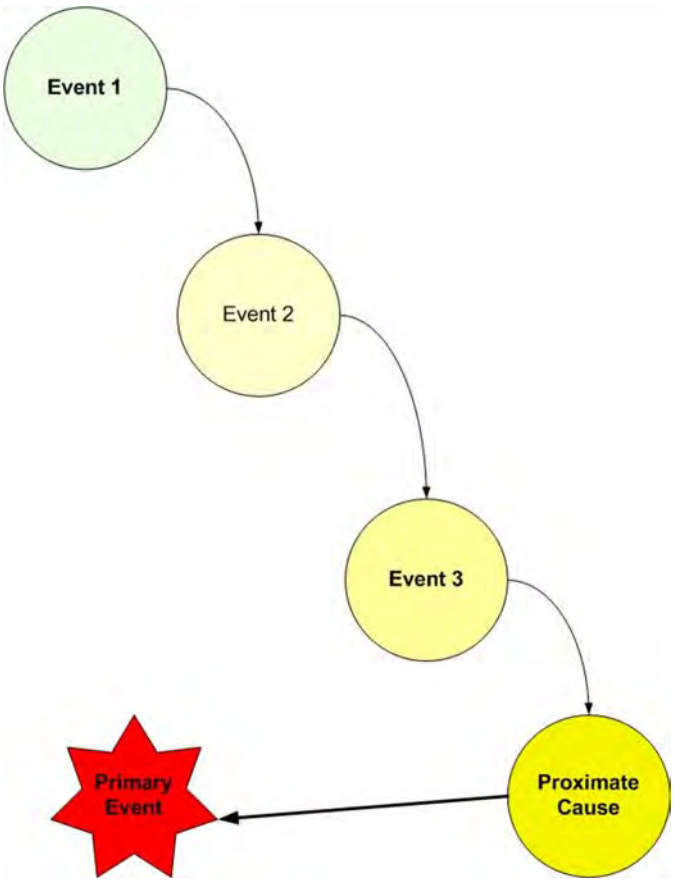

Figure 12: Root Cause Chain of Events

An example of a five-whys analysis is shown in Figure 13. In this example, ransomware crippled the organization because a user fell for a phishing attack. With five-whys, you begin by asking why your data was unavailable. The answer should include any actions taken, processes executed, and the conditions under which the actions and executions happened. When you arrive at the fifth why, the root cause is usually identified. If not, consider starting again. You either missed something or the answers are incorrect. The goal is to break the chain as far as possible from the proximate cause with new controls or procedures; or modifications to existing controls or procedures. However, a layered approach should multiple events along the chain.

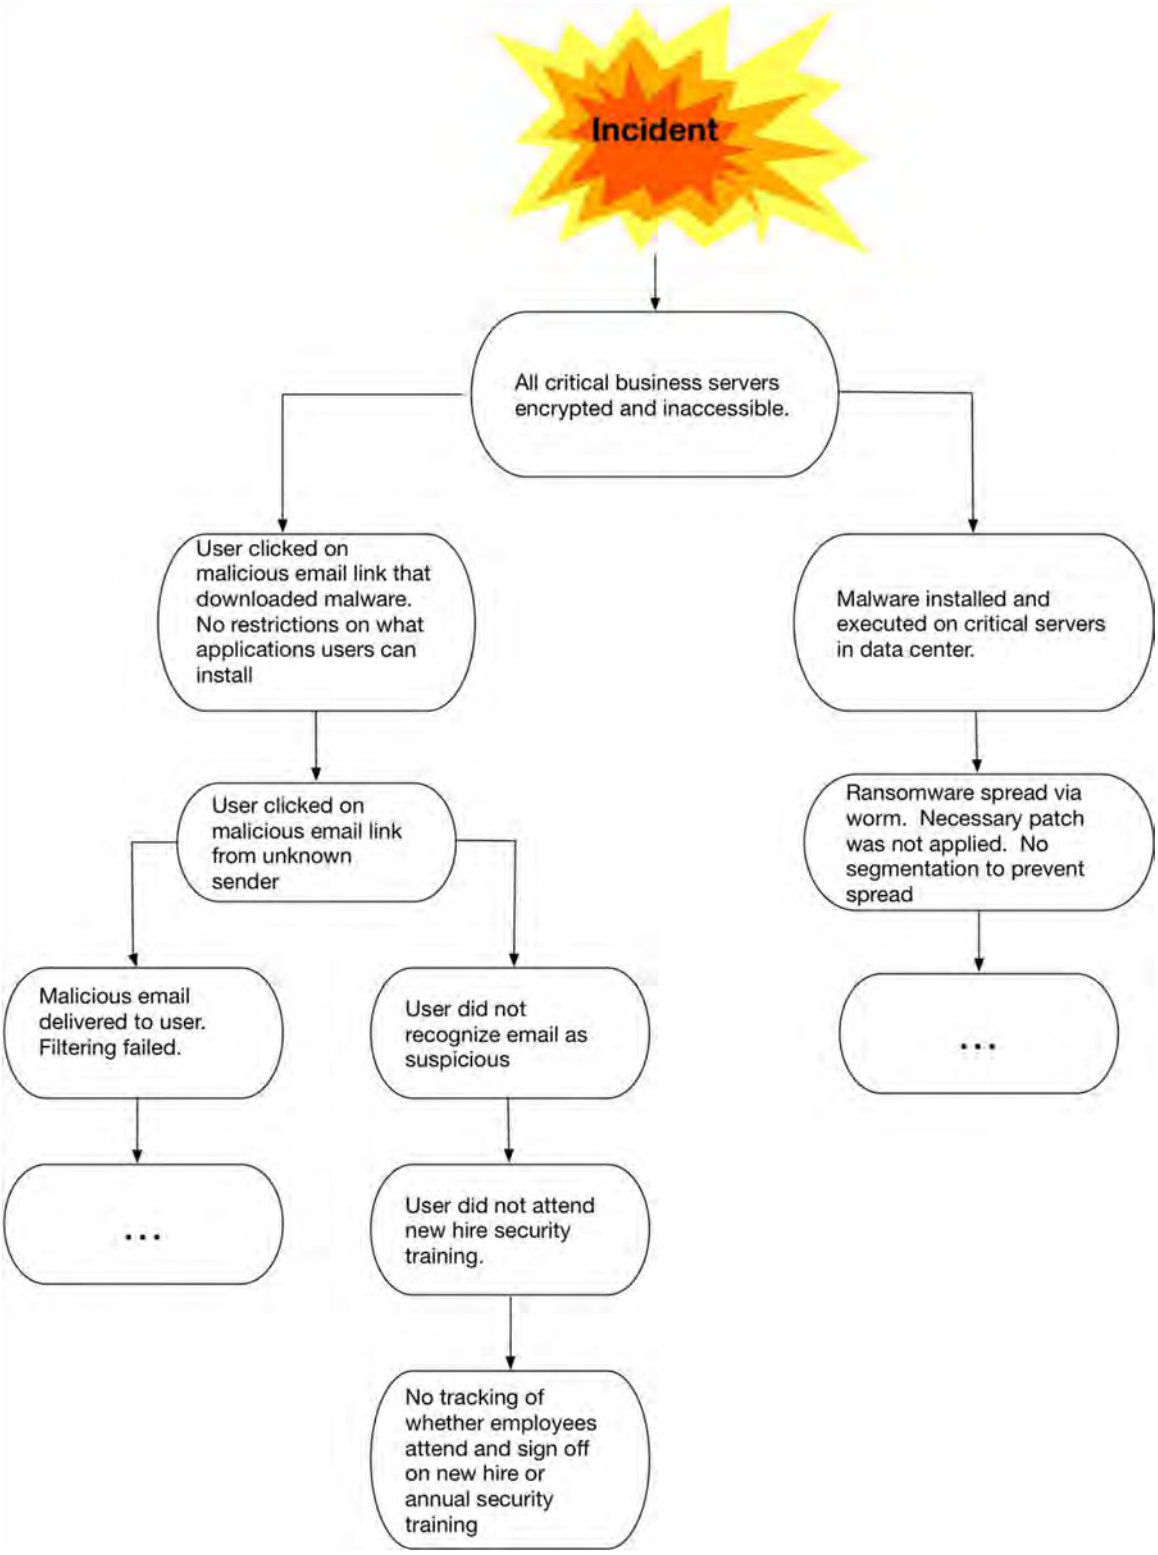

Figure 13: Five Whys

More than one root cause might exist. One of your 'why' answers might include two different causes. Your analysis must then branch off to address both. Consider each branch a separate set of five-whys. You might separate your team, so sub-teams address branches. When all branches are complete and all root causes identified, the entire team comes back together to complete the full analysis diagram.

You will not always know all the answers when first meeting. Consequently, it might take two or three meetings before you arrive at all root causes and create an action plan.

The action plan is part of the final report to management. It includes recommendations for eliminating root causes and improving response. A sample action plan is available for download from http://bit.ly/2rtKAtT. It should minimally include

- Action to take
- Priority of the action
- Plan to complete the action
- Action status
- Person or team assigned
- Date for expected completion

Finally, complete a full report on the incident. Convert your incident log into two stories: one for management and one for your technical teams. The report includes documents and presentations. The presentation to management includes a request for approval for the action plan and an assessment of risk if one or more actions are not approved. A template for a comprehensive but easy to reference incident report is available for download from http://bit.ly/2sm5k3h.

A report is also necessary when analysis of anomalous behavior is deemed not malicious. Referring to these reports during risk assessments or during root cause analysis might reveal previously unrecognized patterns. A shorter report is usually sufficient for this, and a template is available from http://bit.ly/2rk5Nui.

Section 3.03 Section Summary

Planning and creating the tools and procedures for managing an incident must happen before an incident occurs. This enables reasonable and appropriate prevention, detection, and response.

Training team members on the plan is not optional. Everyone on the CSIRT must understand his or her role and how to execute relevant procedures.

Tools for incident response are uniquely designed for each organization. Starting with templates helps ensure you cover all areas. In addition to the tools

provided in this section, the SANS Institute provides an alternative toolset at http://bit.ly/2qASIF1.

Section 4. Response

In this section, we walk through activities that might occur during a response. The walk-through assumes you planned and prepared as described in Section 3. I once again use the response checklist, shown in Figure 14, as our guide. This is a very high-level view of what a response might look like. Each response is unique to what is occurring, so scenario planning as described in Section 3 affects how a response happens and its effectiveness.

Step	Action	Completed	Responsible Person
	Detection and Analysis		
1	Begin Documentation and potential evidence preservation		
2	Determine if incident has occurred		
2.1	Analyze precursors and indicators		
2.2	Look for correlating information		
2.3	Perform research (e.g., search engines, knowledgebase		
3	Prioritize the incident (functional impact, information impact, recovery effort, etc.) and establish situational awareness		
4	Report incident as specified in communications plan		
5	Obtain management decision about forensics preservation and collection		
	Containment, Eradication, and Recovery		
6	Acquire, preserve, and document evidence as directed in Step 5		
7	Contain the incident		
8	Eradicate the incident		
8.1	Identify and mitigate critical exploited vulnerabilities		
8.2	Remove malware, inappropriate materials, untrusted data items, and other components across all affected devices, databases, etc.		
9	Recover		
9.1	Return affected systems to normal operation		
9.2	Confirm normal system operation and business process execution		
9.3	Monitor to ensure eradication was complete and critical vulnerabilities eliminated (continue situational awareness)		
	Post-Incident Activity		
10	Root cause analysis and create action report		

Figure 14: Incident Response Checklist

Section 4.01 Step 1: Begin documentation and potential evidence preservation

Upon notification of anomalous network or device behavior, initiate an incident log. Note

- Date and time of notification
- Person making the notification
- What the person reported
- Systems or networks initially affected

Notify relevant personnel to minimally

- Physically isolate affected user spaces if a crime is committed using a user device
- Avoid further logical or physical contact with affected systems or networks that would unnecessarily modify logs or wipe content: especially avoid powering down or resetting affected systems

Section 4.02 Step 2: Determine if incident has occurred

Using tools implemented during planning and preparation, look for incident precursors and indicators. Hopefully, you will immediately know what is happening or specifically what to examine thanks to automatic threat intelligence associated with your security tools. If not, the following list of things to check is from a poster provided by the SANS Institute at http://bit.ly/2rkg577. This resource includes tools for looking for these conditions in a Windows environment. A Linux version is also available.

- **Unusual log entries**.
  - Did a log activity unexpectedly stop?
  - Are there many failed login attempts or locked out accounts?
  - Are logs unexpectedly accessed or modified?
- **Unusual network usage.**
  - Have any new and unusual file shares appeared?
  - Are unusual sessions open on servers or user devices?
  - Is a large quantity of data moving in unexpected ways?
  - Are unexpected sessions open between internal systems or between internal and external systems?
- **Unusual files and registry keys.**
  - Has there been a major increase or decrease in disk free space?
  - Are there unusually large files?
  - Are there strange programs associated with system start up?
  - Is bulk file encryption occurring?
- **Unusual scheduled tasks.**
  - Are there unusual tasks running as admin, SYSTEM, or a blank user name?
- **Unusual accounts.**
  - Are there new, unexpected accounts in the administrator groups: local or domain?
- **Other.**
  - Are servers or user devices performing sluggishly?
  - Are there unusual system crashes?
  - Is there anything else happening that is unexpected given up-to-date network and system baselines you previously documented?

If you have the right tools in place (SIEM, IPS, firewalls, etc.), you will likely see much of what you need to know in a security management portal: at least you should. In any case, you will want to refer to the list above when determining what

you know and what you do not. Whiteboards with discovered information are a good tool for helping your entire team quickly gain insights in the incident.

Once you collect sufficient information, use your previously identified resources to research what might be the cause of your findings. If you find nothing malicious occurring, complete a short-form report and stop the incident process.

Section 4.03 Step 3: Prioritize the incident and establish situational awareness

Use a tool like the matrix in Section 3 to prioritize the incident based on urgency and impact. Assign a previously designated and trained initial CSIRT member to continuously monitor for conditions in Step 2 throughout the response process. This includes targets thought to be compromised and all critical devices and networks. Isolation does not always work as expected.

Section 4.04 Step 4: Report incident as specified in communications plan

Using the previously defined communication plan, notify the CSIRT members and appropriate management of a probable incident. Be sure to have available for distribution the initial response activities recorded in the log started in Step 1. The CSIRT establishes the incident command center and begins detailed analysis. Analysis never stops as additional information is gathered.

Section 4.05 Step 5: Obtain management forensics evidence collection decision

The initial response should have already taken steps to protect evidence. At this point, management must decide whether to continue forensics processes or focus on business process recovery. Information needed for this decision includes evidence collection impact on how long affected business processes might be down and what evidence is already available without delays (see Section 3). Also relevant is the probability that detailed evidence collection has value given the type of attack and the threat agent involved.

Section 4.06 Step 6: Acquire, preserve, and protect evidence

I address this process in Section 5.

Section 4.07 Step 7: Contain the incident

Although this step appears late in the process, it should be something that happens quickly once an incident is identified. For example, the on-call responder should have the skills and toolset to quickly isolate key network segments after performing Step 2. In Step 7, detailed analysis by the CSIRT and situational awareness information provide the need for additional containment activities.

Also in this step, virtual servers and newly imaged spare user devices can be activated to reduce business process downtime. Depending on how the incident is contained and the processes affected, recovery does not necessarily have to wait until after eradication. However, be sure these systems will not be compromised

again using the same vulnerabilities. This often requires patching or a quick reconfiguration of a network device or control.

Section 4.08 Step 8: Eradicate the incident

At this point, SA and additional analysis should provide enough information about the threat agent and related tools and techniques for eradication. The CSIRT documents a plan for eradication, including scripts and other tools for malware and other unwanted digital entities, and quickly trains response personnel on how to execute it.

The eradication plan includes predefined expeditious change documentation for exploited vulnerability management. Changes include quick modifications to existing network devices, operating systems, business applications, and security controls configurations.

Section 4.09 Step 9: Recover

With proper planning, recovery began in Step 7 and continued through Step 8. What has not occurred yet is verification of data integrity. Use tools and procedures selected and documented during preparation to verify or recover flat file and database data accuracy and authenticity. Remove all containment restrictions and work with business users to ensure affected business processes work as expected: producing valid results.

Section 4.10 Step 10: Root cause analysis and reporting

Gather all personnel involved in incident impact and response to perform an after-action root cause analysis. Complete a detailed response report and presentations for both management and technical teams. The report should include an action plan for management approval and detailed improvement of both security controls/procedures and response activities.

Section 4.11 Section Summary

This section provides a strong foundation for a documented response plan. Based on a checklist, it gives you general actions to take as you step through any type of incident. Again, the type of incident determines specific actions. Therefore, training with various attack scenarios is a necessary part of planning and preparation.

Section 5. Initial Response Forensics

A detailed discussion of digital forensics investigation is outside the scope of this guide. What is important in any response guide is how to initially preserve evidence for forensics investigations. That is what I cover in this section. For a deeper look at digital forensics investigations, see

- NIST SP 800-86 Guide to Integrating Forensic Techniques into Incident Response (http://bit.ly/2qKq9nP)
- NIST SP 800-101 Guidelines on Mobile Device Forensics (http://bit.ly/1odIMvB)
- Digital Forensics/Incident Response Forms, Policies, and Procedures (http://bit.ly/2sxbSML)
- Marshall University Forensic Science Center (http://bit.ly/2qKyqIu)
- NIST Crime Scene Investigation: A Guide for Law Enforcement (http://bit.ly/2rN9swT)

Section 5.01 Forensics Overview

Generally, forensics is the collection, examination, analysis, and reporting of evidence used in identifying and prosecuting perpetrators of a crime. Digital forensics is "the application of science to the identification, collection, examination, and analysis of data while preserving the information and maintaining a strict chain of custody" (Kent, Chevalier, Grance, & Dang, 2006, pp. ES-1).

The process of digital forensics, as shown in Figure 15, is the collection of digital media, the careful extraction of data from that media, correlation of the data to create meaningful information about the crime, and providing credible reports showing relevant evidence for or against one or more suspects. Throughout this process, initial responders and forensics investigators must ensure evidence integrity. Evidence integrity is ensured by

- Establishing a strict chain of custody as soon as potential evidence is collected
- Using only forensically acceptable methods of extracting data from media
- Never using original media for analysis
- Creating forensic copies of media for analysis with hash values calculated immediately after the copy is complete and before the start of analysis
- Allowing only authorized, tracked personnel access to the crime scene, forensics lab, and other areas where evidence is collected or analyzed
- Isolating all analysis systems from networks external to the lab, especially the Internet
- Using only forensics tools known to be acceptable to the forensics community and generally acceptable in legal proceedings

- Being able to demonstrate the professional, skilled status of the forensics investigators in legal proceedings

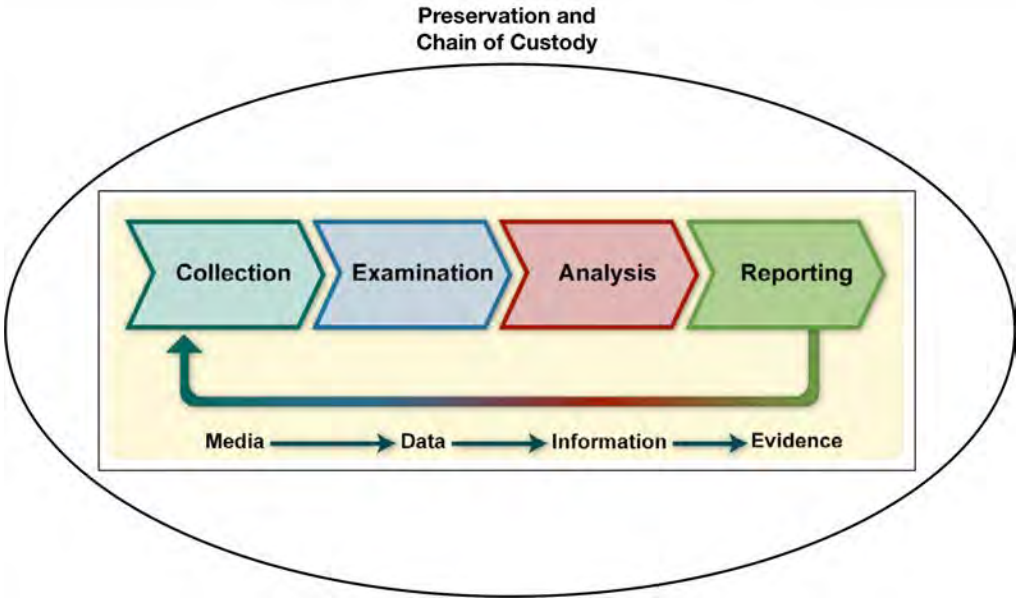

Figure 15: Digital Forensics

Evidence preservation and chains of custody begin with the initial responders. The rest of this section describes how they must work to preserve the integrity of evidence before arrival of forensics investigators.

Section 5.02 Protecting Digital Evidence

We have already discussed in previous sections the importance of securely maintaining logs and other information gathered during daily monitoring. These form the foundation for forensics work. However, we also need contents of swap files and memory, in some cases, to supplement our log information. Consequently, we must never allow anyone to reset or power off any potentially affected devices until management decides how far to proceed with evidence collection.

Ensuring proper handling of user devices during an incident requires training at least our business managers on what to do and not to do when an incident is suspected. In my experience, resetting or powering off a device is a common first step by management. Another management action is often sitting in front of a possibly compromised system, or one that was used in the commission of a crime, to "explore." All these actions must be stopped by policy and training.

IT personnel must also protect evidence in data centers. Once an incident is suspected, a hands-off policy must be enforced. The only exceptions are containment activities defined and directed by the response team. Reaching this outcome requires training and practice.

The hands-off conditions must continue until the digital forensics investigators take over or until management decides to forego detailed evidence collection.

Section 5.03 Securing a Potential Crime Scene

If an office, cubicle, conference room, or other physical space is suspected of use during an incident, you must secure it immediately. First steps include placing someone at the entrance to the area to block all access. Ideally, this would be a security guard. The initial securing of the scene is the responsibility of relevant managers and should take place before arrival of initial responders.

As quickly as possible, the CSIRT should dispatch initial responders to the site. The following steps taken upon arrival are modified recommendations from the NIST's *Crime Scene Investigation: A Guide for Law Enforcement* (http://bit.ly/2rN9swT). When performing these steps, the guiding principle is to avoid anything but minimal contamination and disturbance of evidence.

1. Begin log with notification of incident (date/time, address/location, type of incident, and parties involved) and then log every action taken and observation made at the scene
2. Be aware of any persons leaving the scene
3. Approach the scene cautiously, scan the entire area to thoroughly assess the scene, and note any possible secondary scenes
4. Ensure no one is still using any device or accessing any physical materials at the scene
5. Be aware of any persons in the vicinity that may be related to the crime
    a. Secure and separate suspects
    b. Secure and separate witnesses
    c. Determine if bystanders are witnesses and secure and separate as appropriate
    d. Exclude unauthorized and nonessential personnel from the scene, including managers demanding access
6. Make initial observations to assess the scene and ensure human safety before proceeding
7. Ensure human injuries are treated
8. Remain alert and attentive, and assume the crime is ongoing until otherwise determined
9. Treat the location as a crime scene until assessed and determined to be otherwise
    a. Use crime scene tape to identify and contain all related locations
    b. Log all persons entering and exiting the scene
        i. Time
        ii. Name
        iii. Contact information
        iv. Reason

10. Photograph the scene (and create sketches when photographs do not capture enough details of what you see)
    a. Walls
    b. Floor
    c. Desktops
    d. Computer and handheld device screens
11. Carefully place mobile devices (phones and tablets) into Faraday bags without powering them off and create a chain of custody form for each collected device
12. After photographing all connectors and original locations of the devices, unplug all network cables and ensure the CSIRT has blocked all wireless access for these devices
13. Wait for arrival of forensics investigators, and upon their arrival
    a. Provide detailed briefing of your actions and current state of the scene, witnesses, suspects, evidence, etc.
    b. Turn over all materials and evidence with proper chain of custody
    c. Assist as requested

Just as forensics investigators must be trained professionals, initial responders must have a thorough understanding of what steps to take in any situation. This, again, requires frequent scenario-based training. The checklist shown in Figure 16, and downloadable from http://bit.ly/2sn0ROz, is a good start for a reference and response management tool for initial responders.

Initial Response Team Checklist				
Incident:				
Activity	Date	Time	Assigned	Comments
Identify and obtain treatment for human injuries				
Meet relevant managers and gain their cooperation				
Identify entire scope of scene				
Ensure no one is using/accessing devices or materials within scene				
Mark boundaires with crime scene tape				
Secure and separate suspects				
Identify, secure, and separate witnesses				
Create crime scene access log				
Photograph scene				
Secure mobile devices				
Disconnect computer network connections				
Wired				
Wireless				

Figure 16: Initial Response Team Checklist

The checklist tasks are not necessarily listed in the order in which they are to be completed. A first response team lead should assign tasks to herself and other team members, and some tasks should be done simultaneously, if possible. These include identification and separation of witnesses, securing and separating suspects,

ensuring no unauthorized individuals are in or will enter the crime scene, and isolation of mobile and other devices from network access.

Section 5.04 Section Summary

First steps taken by business users and management are an important part of initial response. Part of security training, at least for managers, should be what to do and what not to do when they suspect an incident.

The CSIRT initial response team must work closely with management once they arrive on the scene. Managing employees, collecting evidence, and other activities need management cooperation. You are not law enforcement. Someone in authority must assist to avoid unnecessary confrontations and delays.

Before arrival of the forensics investigator, only perform those steps necessary to ensure human safety, preserve evidence, and gather witnesses/suspects. Log everything you do or see. Take photographs before touching anything. Create sketches in cases in which a photograph is not quite enough.

Section 6. Works Cited

Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012, August). *Computer Security Incident Handling Guide (NIST SP 800-61r2).* Retrieved May 18, 2017, from NIST (CSRC): http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

Gartner. (2017). *Business Impact Analysis*. Retrieved May 19, 2017, from Gartner IT Glossary: http://www.gartner.com/it-glossary/bia-business-impact-analysis

InfoSec Nirvana. (2015, March). *Part 4 - Incident Management*. Retrieved May 28, 2017, from InfoSec Nirvana: http://infosecnirvana.com/part-4-incident-containment/

Kent, K., Chevalier, S., Grance, T., & Dang, H. (2006, August). *Guide to Integrating Forensic Techniques into Incident Response.* Retrieved June 2, 2017, from NIST: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-86.pdf

Manadhata, P. K., Karabulat, Y., & Wing, J. M. (n.d.). *Report: Measuring the attack surfaces of enterprise software.* Retrieved May 19, 2017, from Carnegie Mellon: School of Computer Science: http://www.cs.cmu.edu/~wing/publications/ManadhataKarabulutWing08.pdf

Olzak, T. (2008, September). *Prevent recurring problems with root cause analysis*. Retrieved May 30, 2017, from TechRepublic: http://www.techrepublic.com/blog/it-security/prevent-recurring-problems-with-root-cause-analysis/

Olzak, T. (2011, June). *Manage the Enterprise Attack Surface*. Retrieved May 19, 2017, from CBS Interactive: http://www.techrepublic.com/downloads/manage-the-enterprise-attack-surface/2949257

Olzak, T. (2012, January). *Risk Management*. Retrieved May 20, 2017, from InfoSec Institute: http://resources.infosecinstitute.com/risk-management-chapter-2/

Olzak, T. (2012, April). *VLAN Network Segmentation and Security*. Retrieved May 23, 2017, from InfoSec Institute: http://resources.infosecinstitute.com/vlan-network-chapter-5/

Olzak, T. (2013, March). *The elements of business continuity planning*. Retrieved May 29, 2017, from TechRepublic: http://www.techrepublic.com/blog/data-center/the-elements-of-business-continuity-planning/

Olzak, T. (2016, May). *Ensure business continuity with change management*. Retrieved May 23, 2017, from CSO: http://www.csoonline.com/article/3067112/business-continuity/ensure-business-continuity-with-change-management.html

Olzak, T. (2017). *Attack Surface Reduction*. Retrieved May 19, 2017, from InfoSec Institute: http://resources.infosecinstitute.com/attack-surface-reduction/

Ross, S. (2010, October). *A business impact analysis checklist: 10 common BIA mistakes*. Retrieved May 21, 2017, from Search Disaster Recovery: http://searchdisasterrecovery.techtarget.com/feature/A-business-impact-analysis-checklist-10-common-BIA-mistakes

Wikipedia. (2017, January). *Checklist Incident Priority*. Retrieved May 26, 2017, from Wikipedia: https://wiki.en.it-processmaps.com/index.php/Checklist_Incident_Priority