

O'REILLY®

Second
Edition

Defensive Security Handbook

Best Practices for Securing Infrastructure

Early
Release

RAW &
UNEDITED



Lee Brotherston
& Amanda Berlin

Defensive Security Handbook

Best Practices for Securing Infrastructure

SECOND EDITION

With Early Release ebooks, you get books in their earliest form—the author’s raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

Lee Brotherston and Amanda Berlin

Defensive Security Handbook

by Lee Brotherston and Amanda Berlin

Copyright © 2022 Amanda Berlin and Lee Brotherston. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://oreilly.com/safari>). For more information, contact our corporate/institutional sales department: 800-998-9938 or corporate@oreilly.com.

Acquisitions Editor: Jennifer Pollock

Development Editor: Shira Evans

Production Editor: Clare Jensen

Interior Designer: David Futato

Cover Designer: Karen Montgomery

Illustrator: Kate Dullea

April 2022: Second Edition

Revision History for the Early Release

- 2022-04-21: First Release
- 2022-07-05: Second Release

See <http://oreilly.com/catalog/errata.csp?isbn=9781098127183> for release details.

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *Defensive Security Handbook*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

978-1-098-12718-3

[LSI]

Chapter 1. Creating a Security Program

A NOTE FOR EARLY RELEASE READERS

With Early Release ebooks, you get books in their earliest form—the authors’ raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

This will be the 1st chapter of the final book. Please note that the GitHub repo will be made active later on.

If you have comments about how we might improve the content and/or examples in this book, or if you notice missing material within this chapter, please reach out to the editor at sevans@oreilly.com.

Creating or improving upon a security program can be a daunting task. With so many facets to consider, the more initial thought and planning that is put into the creation of this program, the easier it will be to manage in the long run. In this chapter, we will cover the skeleton of a security program and initial administrative steps.

Do not fall into the habit of performing tasks, going through routines, or completing configuration with the mindset of, “This is how we’ve always done it.” That type of thinking will only hinder progress and decrease security posture as time goes on.

Humans are allergic to change. They love to say, “We’ve always done it this way.” I try to fight that. That’s why I have a clock on my wall that runs counter-clockwise.”

Grace Hopper, “The Wit and Wisdom of Grace Hopper” (1987)

We recommend that when creating the program, you follow this chapter in order. While we attempted to group the remaining chapters accordingly, they can be followed as best fits a company.

Lay the Groundwork

It is not necessary to reinvent the wheel in order to lay out the initial groundwork for an information security program. There are a few standards that can be of great use that we will cover in Chapter 8. The National Institute of Standards & Technology (NIST) has a risk-based cybersecurity framework that covers many aspects of a program. The NIST Framework Core consists of five concurrent and continuous functions—Identify, Protect, Detect, Respond, and Recover. When considered together, these functions provide a high-level, strategic view of the lifecycle of an organization’s management of **cybersecurity risk**. Not only will a framework be a possible asset, so will compliance standards. Although poorly implemented compliance standards can hinder the overall security of an organization, they can also prove to be a great starting point for a new program. We will cover compliance standards in more depth in Chapter 8. While resources like these can be a phenomenal value add, you must always keep in mind that every organization is different, and some aspects covered may not be relevant (there are continuous recurring reminders of this throughout the book).

Establish Teams

As with many other departments, there are virtues in having the correct staff on the correct teams in regards to security. Open cross-team communication should be a primary goal, as without it the security posture is severely weakened. While smaller organizations may combine several of the following teams, or have a lack of them all together, this remains a good goal to populate a security department.

Executive team

A chief information officer (CIO) or chief information security officer (CISO) will provide the leverage and authority needed for businesswide decisions and changes. An executive team will also be able to provide a long-term vision, communicate corporate risks, establish objectives, provide funding, and suggest milestones.

Risk team

Many organizations already have a risk assessment team, and this may be a subset of that team. In the majority of organizations, security is not going to be the number-one priority. This team will calculate risks surrounding many other areas of the business, from sales to marketing and financials. Security may not be something they are extremely familiar with. In this case they can either be taught security basics case by case, or a security risk analyst could be added to the team. A risk framework such as **NIST's Risk Management Framework (RMF)** or the Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) Framework can assist with this.

Security team

The security team will perform tasks to assess and strengthen the environment. The majority of this book is focused toward this and the executive team. They are responsible for daily security operations, including managing assets, assessing threats and vulnerabilities, monitoring the environment for attacks and threats, managing risks, and providing training. In a large enough environment, this team can be broken up into a variety of subteams such as network security, security operations, security engineering, application security, and offensive security.

Auditing team

It is always a good idea to have a system of checks and balances. This is not only to look for gaps in the security processes and controls, but also

to ensure the correct tasks and milestones are being covered. As with the Risk Team, this may be a subset of a larger group.

However it is entirely possible that a Small to Medium Business (SMB) may combine one or (unfortunately due to things like budget constraints/etc) all of these roles into one. In those cases we definitely commiserate with you, as it happens all too often. As the company grows, and hopefully the security program also grows, these separate roles can then be planned and adequately filled.

Baseline Security Posture

The unknowns in any environment are going to be scary, but that shouldn't stop you from diving in. How will you know what level of success the program has had without knowing where it started? At the beginning of any new security program or any deep dive into an existing one, a baseline and discovery phase should be one of the first and foremost tasks at hand for all teams. Throughout this book we will cover asset management several times in different ways. The baseline of the security of the organization is just another step in that management. Items that should be gathered include:

- Policies, procedures, and incident response playbooks
- Endpoints—desktops and servers, including implementation date and software version
- Licensing and software renewals, as well as SSL certificate expiration dates
- Internet footprint—domains, mail servers, dmz devices, cloud architecture
- Networking devices—routers, switches, APs, IDS/IPS, and Network Traffic
- Logging and monitoring

- Ingress/egress points—ISP contacts, account numbers, and IP addresses
- External vendors, with or without remote access, and primary contacts
- Applications - any primary software applications either maintained by your company, or used in any aspect as primary business functions

Assess Threats and Risks

Assessing threats and risks will be incredibly different for each and every organization. Each internal and external footprint is unique when combined with the individual infrastructure involved. Assessing these includes both a high-level overview, as well as in-depth knowledge of assets. Without the knowledge of the threats and risks your organization faces, it is more difficult to custom fit technologies and recommendations to provide a suitable defense. As mentioned previously, a risk team or role is an essential part of creating an information security team. Much more detailed information can be found when researching Governance, Risk, and Compliance (GRC).

As we can't cover the entirety of GRC in this book we'll go over a general risk framework. There are a handful of risk management frameworks, but can be summarized in 4 steps: identify, assess, mitigate, and monitor.

Identify Scope, Assets, & Threats

Organizations should be concerned with a large amount of threats and risks that will cross industry verticals. Focusing on industry trends and specific threats will allow the security program to be customized and prioritized to become more efficient. Many organizations have put very little thought into what threats and risks they face on a day-to-day basis, and will continue to do so until they fall victim to them. Invaluable resources in this case are available through Information Sharing and Analysis Centers (ISACs), which are brought together by the **National Council of ISACs** to share sector-specific Information Security. "ISACs collect, analyze and

disseminate actionable threat information to their members and provide members with tools to mitigate risks and enhance resiliency.”¹

Not only should industry-specific threats be identified, but also overall trending threats such as malware, ransomware, phishing, and remote exploits. Three very important places to make note of are the OWASP top 10, Center for Internet Security (CIS) 18, and the standards outlined by the Cloud Security Alliance. The majority of the items on these lists will be covered in more depth in this book, but keeping up-to-date with them year to year should be a key part of any strategic plan.

Assess Risk & Impact

After the potential risks have been identified, assess these risks to determine if they apply to the particular environment. Tasks such as internal and external vulnerability scans, firewall rule audits, authentication/user permission assessments, and asset management and discovery will lend a larger picture to the type of overall risk exposure.

During the assessment step, you’ll want to analyze each identified risk and determine the likelihood of it negatively impacting the organization, how severe that impact might be, and what it would look like when executed. For example:

Threat: An attacker exploits a new vulnerability on an _____.

Vulnerability: unpatched

Asset: mail server

Consequence: to use RCE access to pivot to internal systems

Mitigate

Mitigation of risks is the meat and bones of why we’re all here; it’s also the purpose of the majority of this book. Options include avoiding, remediating, transferring, or accepting the risk. Some examples:

Risk avoidance

Dave decides that storing Social Security numbers for customers is an unneeded process and discontinues the practice.

Risk remediation

Alex starts turning off open ports, implementing stricter firewall rules, and patching endpoints.

Transferring of risk

Ian outsources credit card processing to a third-party as opposed to storing data on site.

Accepting risk

Kate knows that a certain endpoint has no access to other endpoints and runs a third-party application. This application has a low-risk vulnerability that is required for it to function. While nothing at that point can be changed or remediated with that vulnerability, the risk is low enough to accept.

NOTE

You should only accept risk as a last resort. If a risk ever makes it to this point, request full documentation from third-party vendors and the executive team, as well as documentation of processes that have been attempted prior to making this decision. Add at least an annual review of any accepted risks to ensure they are revisited accordingly.

Monitor

Keep track of the risk over time with scheduled quarterly or yearly meetings. Throughout the year, many changes will have taken place that affect the amount and type of risk that you should consider. As a part of any change monitoring or change control, determine if the change is affecting risk in any way. One way of tracking ongoing risk status is by using a risk register to document different scenarios, controls, and treatment plans, which also can be combined in any vulnerability management program.

Prioritize

Once threats and risks have been identified and assessed, they must also be prioritized from highest to lowest risk percentage for remediation, with a concentration on ongoing protection. This doesn't always have to be an expensive venture, however. A large amount of defensive mitigations can be performed at little or no cost to an organization. This enables many opportunities to start a security program without having a budget to do so. Performing the due diligence required to get the program off the ground for free should speak volumes to an executive team.

NOTE

Do not always take application, vendor, or third-party advice for prioritization. Every environment is different and should be treated as such. Prioritize tasks based on the bigger picture when all of the information has been collected.

This book wasn't written to be a sequential list of security tasks to complete. Prioritization can differ greatly from environment to environment. Just remember, if the environment is already on fire and under attack, don't start by creating policies or reversing malware. As a fire marshall, you shouldn't be worried about looking for the arsonist and point of origin when you haven't even put out the fire yet.

To determine the priority of some risks, you can use a risk matrix, where the overall risk level is calculated by taking "Likelihood times Impact"

5x5 risk matrix

5: VERY SEVERE	Medium 5	Medium high 10	High 15	Very high 20	Very high 25
4: SEVERE	Low 4	Medium 8	Medium high 12	High 16	Very high 20
3: MODERATE	Low 3	Medium 6	Medium 9	Medium high 12	High 15
2: MINOR	Low 2	Low 4	Medium 6	Medium 8	Medium high 10
1: NEGLIGIBLE	Low 1	Low 2	Low 3	Low 4	Medium 5
	1: RARE	2: UNLIKELY	3: POSSIBLE	4: LIKELY	5: HIGHLY LIKELY

SOURCE: MICHAEL COBB

©2020 TECHTARGET. ALL RIGHTS RESERVED. 

Figure 1-1. <https://www.techtargget.com/searchsecurity/tip/How-to-perform-a-cybersecurity-risk-assessment-step-by-step>

Create Milestones

Milestones will take you from where you are to where you want to be. They will be a general progression on the road to a secure environment. This is

heading a little into project manager (PM) duties, but in many cases companies do not have dedicated PMs. Milestones can be broken up loosely into four lengths or tiers:

Tier 1: Quick wins

The earliest milestones to meet should be quick wins that can be accomplished in hours or days—high vulnerabilities such as one-off unused endpoints that can be eliminated, legacy devices that can be moved to a more secure network, and third-party patches all could fall under this category. We will mention many free solutions as the sales process can take a significant time to complete.

Tier 2: This year

Higher vulnerabilities that may need to go through a change management process, create a change in process, or be communicated to a significant amount of people might not end up in Tier 1. Major routing changes, user education implementation, and decommissioning shared accounts, services, and devices are all improvements that also require little-to-no-budget to accomplish, but can take a little more time due to the need for planning and communication.

Tier 3: Next year

Vulnerabilities and changes that require a significant amount of planning or that rely on other fixes to be applied first fall into this tier. Transitioning entire business functions to a cloud service, domain upgrades, server and major infrastructure device replacements, monitoring, and authentication changes are all good examples.

Tier 4: Long-term

Many times a milestone may take several years to accomplish, due to the length of a project, lack of budget, contract renewals, or difficulty of change. This could include items such as a network restructure, primary software replacement, or new datacenter builds.

It is helpful to tie milestones to critical controls and risks that have already been identified. Although starting with the higher risks and vulnerabilities is a good idea, they may not be easy fixes. In many cases, not only will these items take a significant amount of time and design, but they may require budget that is not available. All aspects need to be taken into account when creating each tier.

Use Cases, Tabletops, and Drills

Use cases are important for showcasing situations that may put critical infrastructure, sensitive data, or other assets at risk. Brainstorm with data owners and leaders to plan ahead for malicious attacks. It is best to come up with around three different use cases to focus on in the beginning and plan on building security mitigations and monitoring around them. Items such as ransomware, DDoS (Distributed Denial of Service), disgruntled employee, insider threat, and data exfiltration are all good examples of possible use cases. After several use cases have been chosen they can be broken down, analyzed, and correlated to each step of any one of the security frameworks that are covered in this book, or additionally others that may end up being created after we're done writing. One example of a common framework used to map use cases is [Lockheed Martin's Intrusion Kill Chain](#).

The Intrusion Kill Chain, sometimes called the Cyber Kill Chain, is “a model for actionable intelligence when defenders align enterprise defensive capabilities to the specific processes an adversary undertakes to target that enterprise.” It is composed of seven steps as described in the [Lockheed Martin whitepaper](#):

1. Reconnaissance: research, identification, and selection of targets, often represented as crawling internet websites such as conference proceedings and mailing lists for email addresses, social relationships, or information on specific technologies.
2. Weaponization: coupling a remote access trojan with an exploit into a deliverable payload, typically by means of an automated tool

(weaponizer). Increasingly, client application data files such as Adobe Portable Document Format (PDF) or Microsoft Office documents serve as the weaponized deliverable.

3. **Delivery:** transmission of the weapon to the targeted environment. The three most prevalent delivery vectors for weaponized payload are email attachments, websites, and USB removable media.
4. **Exploitation:** After the weapon is delivered to victim host, exploitation triggers intruders' code. Most often, exploitation targets an application or operating system vulnerability, but it could also more simply exploit the users themselves or leverage an operating system feature that auto-executes code.
5. **Installation:** installation of a remote access trojan or backdoor on the victim system allows the adversary to maintain persistence inside the environment.
6. **Command and Control (C2):** Typically, compromised hosts must beacon outbound to an internet controller server to establish a C2 channel. APT malware especially requires manual interaction rather than conduct activity automatically. Once the C2 channel establishes, intruders have "hands on the keyboard" access inside the target environment.
7. **Actions on Objectives:** only now, after progressing through the first six phases, can intruders take actions to achieve their original objectives. Typically, this objective is data exfiltration, which involves collecting, encrypting and extracting information from the victim environment; violations of data integrity or availability are potential objectives as well. Alternatively, the intruders may only desire access to the initial victim box for use as a hop point to compromise additional systems and move laterally inside the network.

This whitepaper has a good amount of information that can be used for creating use cases as well.

Table 1-1 is an example of a step-by-step kill chain use case we've created for a ransomware attack.

*T
a
b
l
e*

*l
-*

.

R

a

n

s

o

m

w

a

r

e

u

s

e

c

a

s

e

Kill chain step Malicious action Defensive mitigation Potential monitoring

Reconnaissance

Attacker obtains email addresses, technologies used, and creates an organizational profile based on that information.

Create policies around sharing internal information on sites such as LinkedIn or using corporate email addresses for nonbusiness use.

Have corporate emails been seen in breaches elsewhere? How many emails are found with OSINT?

After a major breach has been seen on the news run a password reset. Even though they shouldn't, employees will reuse passwords for other services and sites.

Weaponization

Attacker creates a malicious exploit to send to the victim, or uses a current exploit.

Knowledge and awareness of threats currently being used by attackers will allow for better constructed and tuned mitigation steps.

n/a

Delivery

A user receives a phishing email.

Assess which attachment types are needed in the organization. File types such as *.js* can be extremely harmful and are rarely exchanged from external sources.

Instill the idea of “trust but verify” to your users.

Implement Ad-Blocking

Implement mailing blacklists and greylists such as Spamhaus and dnsbl to block known malicious mail servers.

Filetypes of a certain size known to be malicious and associated with ransomware. (Flag *.scr* files over 22 MB and *.js* over 15 MB.)

Exploitation

Endpoint downloads a JavaScript file or Word document with malicious macro.

Disable macros and malicious filetypes via group policy.

Ensure any endpoint protection is up-to-date and installed.

Monitor proxy logs for unexpected file retrievals (e.g., JavaScript is the first file fetched from that host, host is on a threat intel list, etc.)

Use proxies or IDS (if cleartext) to monitor for known deobfuscation strings.

Installation

The payload is executed on the end user's device. (Lucky, Cerber, and CryptoWall use the built-in Windows Crypto API to handle the encryption.)

Keep backups (that are not permanently attached) so that encrypted files can be restored easily.

High increase in Windows Crypto API over short amount of time.

Depending on OS, you can use "filesystem firewalls" such as **Little Flocker** to permit access to files on

Excessive numbers in a domain or low % of meaningful strings in domain.

per-process basis. That means that you can permit read access to MS Word, but not IE, for example.

There are experimental techniques that can be used to block crypto-based ransomware (e.g., **Decryptonite**)

Command & Control (C&C)

The ransomware contacts a C&C server on the internet to transmit the decryption key.

Implement sinkhole DNS and autoblock outbound connections to known malicious IP addresses using Dynamic Block Lists (DBL).

Connection to known C&C servers.

Actions & Objectives

The malware starts encrypting the files on the hard disk, mapped network drives, and USB devices. Once completed, a splash screen, desktop image, website, or text file appear with instructions for the ransom.

Implement Honey Directories—the ransomware goes into C:\\$\$ it sees another \$\$ directory, when it goes into C:\\$\$\\$\$ it sees another \$\$ directory, and so on.

Advanced file auditing can be enabled for alerting on an extreme increase in filesystem changes.

Many different defensive mitigations and specific detections can be added at each step of the kill chain for an overall decrease in risk at each layer.

Following the creation and implementation of security controls around use cases, and employing tabletop exercises and drills can serve as a proof of concept as well as assist in an ever growing collection of playbooks. A *tabletop exercise* is a meeting of key stakeholders and staff who walk step by step through the mitigation of some type of disaster, malfunction, attack, or other emergency in a low stress situation. A drill is when staff carries out as many of the processes, procedures, and mitigations that would be performed during one of the emergencies as possible.

While drills are limited in scope, they can be very useful to test specific controls for gaps and possible improvements. A disaster recovery plan can be carried out to some length, backups can be tested with the restoration of files, and services can be failed over to secondary cluster members.

Tabletop exercises are composed of several key groups or members.

- During a tabletop exercise there should be a moderator or facilitator who will deliver the scenario to be played out. This moderator can answer “what if” questions about the imaginary emergency, as well as lead discussion, pull in additional resources, and control the pace of the exercise.
- Inform the participants that it is perfectly acceptable to not have answers to questions during this exercise. The entire purpose of tabletops is to find the weaknesses in current processes to mitigate them prior to an actual incident.
- A member of the exercise should also evaluate the overall performance of the exercise, as well as create an after-action report. This evaluator should take meticulous notes and follow along any runbook or playbook to ensure accuracy. While the evaluator will be the main notetaker, other groups and individuals may have specific knowledge and understanding of situations. In this case, having each member provide the evaluator with their own notes at the conclusion of the tabletop is a good step.

- Participants make up the majority of this exercise. Included should be groups such as finance, HR, legal, security (both physical and information), management, marketing, and any other key department that may be required. Participants should be willing to engage in the conversation, challenge themselves and others politely, and work within the parameters of the exercise.

What to include in the tabletop:

- A handout to participants with the scenario and room for notes.
- Current runbook of how security situations are handled.
- Any policy and procedure manuals.
- List of tools and external services.

Post-exercise actions and questions:

- What went well?
- What could have gone better?
- Are any services or processes missing that would have improved resolution time or accuracy?
- Are any steps unneeded or irrelevant?
- Identify and document issues for corrective action.
- Change the plan appropriately for next time.

TABLETOP TEMPLATES

- [IncidentResponse.com](https://www.incidentresponse.com)
- [Microsoft's Incident Response Playbooks](#)
- The Federal Emergency Management Agency (FEMA) has a collection of scenarios, presentations, and tabletops that can be used as [templates](#).

Expanding Your Team and Skillsets

Finding a dedicated, passionate, and intelligent team can be one of the most difficult aspects of any professional's life.

What can you and your team do to expand knowledge and skillsets?

- Encourage staff to either set up a home lab or provide a lab for them. Labs can be used for testing out real-world scenarios, as well as practicing skills and learning new ones. Labs can be created at a relatively low cost by buying secondhand equipment. The best way to learn for the majority of people is hands-on, and with a lab there is no risk introduced into a production environment.
- Compete in or create Capture the Flag competitions (CTFs). CTFs are challenging, and they can provide cross training and team building, as well as increase communication skills. Most information security conferences have CTFs. If you are looking to expand a team, CTFs are a wonderful place to find new talent. Not only will participants be showing their level of knowledge, but also communication skills, how well they work with others in a team, and their willingness to help and teach others.
- Find or create a project. Automate something in the enterprise, find a need and fill it. It doesn't matter what the skillset, there will be a project out there that needs help. Documentation is needed on 99% or more of the open source projects out there.
- Attend, organize, volunteer, speak, sponsor, or train at an industry conference or local meetup. There are hundreds of them across the US and they almost always need volunteers. Just attending a conference has its benefits, but truly immersing yourself will push you further to learn and experience more. Many careers have been started by having a simple conversation about a passion over lunch or a beer. Networking is a game changer in our industry, but it's not the silver bullet for everyone. You can network all you want, but unless you are a desirable candidate it won't matter. Having a willingness and desire to

learn, listen, collaborate, and the ability to think for yourself are all ideal traits in such a fast-paced industry.

- Participate in mentoring. Whether as a mentor or mentee, structured or unstructured, mentoring can be a valuable learning process both on and off the job.

Conclusion

Creating an information security program is no easy task. Many programs are broken or nonexistent, adding to the overall lack of security in the enterprise environment today. Use this book as a guide to work through the different areas and to suit them to a custom-tailored plan. Organizational skills, a good knowledgeable, hard-working team, strong leadership, and an understanding of the specific environment will all be crucial to an effective program.

¹ <https://www.nationalisacs.org/about-isacs>

Chapter 2. Asset Management and Documentation

A NOTE FOR EARLY RELEASE READERS

With Early Release ebooks, you get books in their earliest form—the authors' raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

This will be the 2nd chapter of the final book. Please note that the GitHub repo will be made active later on.

If you have comments about how we might improve the content and/or examples in this book, or if you notice missing material within this chapter, please reach out to the editor at sevans@oreilly.com.

As a whole, asset management is not an information security function. However, there are definitely components to it that assist in strengthening the overall security posture. It is one of the most difficult verticals to cover. Without proper asset management, an environment cannot be protected to its full potential. It is impossible to protect assets that are unknown, and they can be a significant disadvantage in any troubleshooting or incident response. In larger and older networks, it is next to impossible to completely be aware of each and every device that is connected or every piece of software the users may have installed. However, with the correct organization and security controls in place, it becomes much easier.

The majority of this chapter will cover how best to find assets, tie all of the information together, and document it for ease of use and troubleshooting. Above all else, the two most important things to remember about asset

management are to ensure there is one source of truth, and that it is a process, not a project.

Additionally, each asset or group of assets must be assigned an owner and/or a custodian. An asset owner serves as a point of contact for the assigned asset, whereas a custodian has responsibility for the stored information. The assets are then categorized into different levels of importance based on the value of the information contained in them and the cost to the company if an asset is compromised.

Information Classification

The need for information classification has risen as the amount of data on digital storage has grown. Attackers use confidential data for their profit by selling it on the black market, to expose or cripple a specific organization, to commit fraud, or to aid in identity theft. While many industry compliance standards such as HIPAA and PCI DSS attempt to dictate the type of information that should be specifically guarded and segregated, that may not be the only data that is classified as confidential in an organization. There may also be contracts and other legal measures that must be consulted for classification and protection of certain data. Steps to correctly classify data can be described as follows:

1. Identify data sources to be protected. Completion of this step should produce a high-level description of data sources, where they reside, existing protection measures, data owners and custodians, and the type of resource. Obtaining this information can be difficult, but can be an added part of the documentation process as data owners and custodians are assigned and documented.
2. Identify information classes. Information class labels should convey the protection goals being addressed. Classification labels like Critical and Sensitive have different meanings to different people, so it is important that high-level class descriptions and associated protection measures are meaningful and well-defined to the individuals who will

be classifying the information, as well as those who will be protecting it.

3. Map protections to set information classification levels. Security controls such as differing levels and methods of authentication, air-gapped networks, firewalls/ACLs, and encryption are some of the protections involved in this mapping.
4. Classify and protect information. Now the hands on work comes in! All information that has been identified in step 1 should now be classified as dictated in step 2, and protected as in step 3.
5. Repeat as a necessary part of a yearly audit. Data footprints are ever expanding. As new software is installed or upgraded with add-ons, data grows or changes in scope. A yearly audit of the current data footprint in the enterprise will be required to ensure data continues to be protected as documented.

Asset Management Implementation Steps

The asset management process can be separated out into four distinct steps: defining the lifecycle, information gathering, change tracking, and monitoring and reporting. Assets can be added to the environment at an alarming rate via scripted virtual machine roll outs or business acquisitions, refreshed to new hardware or software versions, or removed altogether. There are several enterprise-level tools that assist in identifying data on systems. A solution should be implemented that will track an asset from as early as possible until its eventual decommissioning.¹

Defining the Lifecycle

There are many lifecycle stages in between delivery and decommissioning: an asset may be moved; the person it's assigned to may no longer be employed; it may require repair or replacement; or it may go inactive while its assigned user is on leave of absence. Define lifecycle events and document them. Each department or person involved in each step should

understand when and how assets are tracked at every point of their lifecycles. This assists with ensuring that any unplanned deviation from documented processes is caught. Following is a map of a very basic asset management lifecycle:

Procure

This is the procurement step of the lifecycle where assets are initially added to be tracked. At this point, the initial device information, such as serial number, PO, asset owner, criticality, and model name and number, can be added to the tracking system.

Deploy

When an asset is deployed by a sys admin, net admin, helpdesk member, or other employee, the location of the device can now be updated and any automated software data population can be tested. Remember: prior to deploying assets, they should be scanned for viruses and vulnerabilities or built with a custom secure image (if applicable) before being attached to the network. Too many times, do assets arrive from vendor's that are already pre-infected, or shipped with outdated software that have security flaws.

Manage

The management lifecycle step can contain many subsections depending on the level of documentation and tracking that is decided upon. Items can be moved to storage, upgraded, replaced, or returned, or *may* change users, locations, or departments.

Decommission

Decommissioning assets is one of the most important steps of the lifecycle due to the inherent security risks regarding the disposal of potentially confidential data. When deciding on disposal options, different classifications of data can be tied to varying levels. There are many different ways to destroy data, and these have varying levels of

security and cost. For example, for a physical spinning HDD options would include:

- Staging for disposal
 - A single pass wipe: drives can be reused and provide a residual value return.
 - Multiple wipes: increases the level of security; this still protects the residual value but adds cost.
 - Degaussing: removes the ability to resell the drive and has no visual indicator that it has worked, but is a cheaper alternative to shredding and is often more practical.
 - Full disk encryption: drives can be reused; this increases the level of security.
- Physical disposal
 - Crushing/drilling/pinning: these are other low-cost options, all of which show physical evidence that they have been completed and deter ordinary criminals. These methods destroy the value of the unit and the data is still present on the platters despite them not being able to spin up.
 - Shredding: the most secure form of data destruction and typically used for the highest levels of secure data, but it is expensive and destroys the ability for resale.
 - Removed as asset from inventory

Information Gathering

Information gathering contains the most hurdles and has a huge amount of complexity from environment to environment. As stated previously, obtaining a good software package coupled with well thought-out processes will lead to gathering as much information as possible on network-

connected assets. If you do not have the budget for a full blown asset management and discovery platform, there are a handful of good free, cheap, or open source projects available such as **Rumble** or Netdisco.

Figure 2-1 shows how Rumble can be used for asset inventory, and it is currently free for smaller organizations.

Asset inventory

[New Scan](#) [Import](#) [Export](#) [Reports](#) [Modify](#)

Search assets...

Build a query

Show 20 per page

Comm Tags Scan Del Merge Cols Reset

Address	Up	Name	OS	Type	MAC	Vendor	Age	HW
<input type="checkbox"/> 192.168.0.8	●	ACCESS-CONTROL ⁺¹	Raspbian Linux 9.0	Server	B8:27:EB:04:BD:13	Raspberry Pi ...	2012-02-11	Raspberry Pi
<input type="checkbox"/> 192.168.0.19	●	APCFD796D	APC AP9631 6.5.0	UPS	00:CO:87:FD:79:6D	AMERICAN POWE...	1998-04-22	APC Smart UPS Web Manage
<input type="checkbox"/> 192.168.100.69	●	AS-P_11THFLOOR	Linux	BACnet				Schneider Electric SmartX C
<input type="checkbox"/> 192.168.100.70	●	AS-P_18THFLOOR	Linux	BACnet				Schneider Electric SmartX C
<input type="checkbox"/> 192.168.100.71	●	AS-P_21STFLOOR	Linux	BACnet				Schneider Electric SmartX C
<input type="checkbox"/> 192.168.30.47	●	BREAKROOM	Ubiquiti Linux	WAP	F0:9F:C2:69:D4:6B ⁺⁹	Ubiquiti Netw...	2014-12-17	Ubiquiti UAP-HD 4.0.80.108
<input type="checkbox"/> 192.168.30.202	●	BRN3C2AF4ABE1C6 ⁺¹	Brother NC-8900h ZD	Printer	10:5B:AD:4A:69:45 ⁺²	Mega Well Lim...	2018-07-07	Brother DCP-9042CDN
<input type="checkbox"/> 192.168.30.117	●	DATACENTER	Ubiquiti Linux	WAP	80:2A:A8:93:D2:B1 ⁺⁴	Ubiquiti Netw...	2014-12-17	Ubiquiti UAP-AC-LR 4.0.80.1
<input type="checkbox"/> 192.168.0.41 ⁺³	●	DATASETS	Synology Linux DSM	NAS	00:11:32:8B:6D:5C	Synology Inco...	2004-04-25	Synology NAS
<input type="checkbox"/> 192.168.100.117	●	DCC1_2_3000		BACnet				Tridium NiagaraAX Station
<input type="checkbox"/> 192.168.0.251	●	DESKTOP-DFC6JIQ	Windows 10 (1903)	Desktop	00:0C:29:58:E3:A0	VMware, Inc.	2003-01-21	
<input type="checkbox"/> 192.168.30.116	●	DOWNSTAIRSCONFERENCE	Ubiquiti Linux	WAP	F0:9F:C2:D0:BA:64 ⁺⁴	Ubiquiti Netw...	2014-12-17	Ubiquiti UAP-AC-Mesh-Pro 4
<input type="checkbox"/> 192.168.0.238	●	IPMI	Super Micro ATEN Linux	BMC	0C:C4:7A:70:65:0D	Super Micro C...	2013-10-24	Super Micro IPMI
<input type="checkbox"/> 192.168.0.244	●	IPMI	Super Micro ATEN Linux	BMC	0C:C4:7A:70:65:7C	Super Micro C...	2013-10-24	Super Micro IPMI
<input type="checkbox"/> 192.168.100.185	●	KCC-BOILER-PLANT ⁺¹		BACnet				Reliable Controls Corporatio
<input type="checkbox"/> 192.168.0.220	●	LABDC01	Windows Server 2016 (RTM)	Server	00:0C:29:D9:E3:45	VMware, Inc.	2003-01-21	
<input type="checkbox"/> 192.168.0.30	●	M4300-DC	Linux	Switch	10:DA:43:D7:53:2B ⁺²	NETGEAR	2015-12-22	Netgear M4300-8X8F
<input type="checkbox"/> 192.168.0.31	●	M4300-OFFICE		Switch	10:DA:43:D7:59:14 ⁺²	NETGEAR	2015-12-22	Netgear M4300-8X8F
<input type="checkbox"/> 192.168.0.32	●	M4300-WORKBENCH		Switch	10:DA:43:D7:53:60 ⁺²	NETGEAR	2015-12-22	Netgear M4300-8X8F
<input type="checkbox"/> 192.168.0.5 ⁺²	●	MACMINI-EE7C7B ⁺²	Apple macOS 10.15	Desktop	F0:18:98:EB:7C:7B	Apple, Inc.	2017-12-23	Apple Mac Mini (Late 2018)

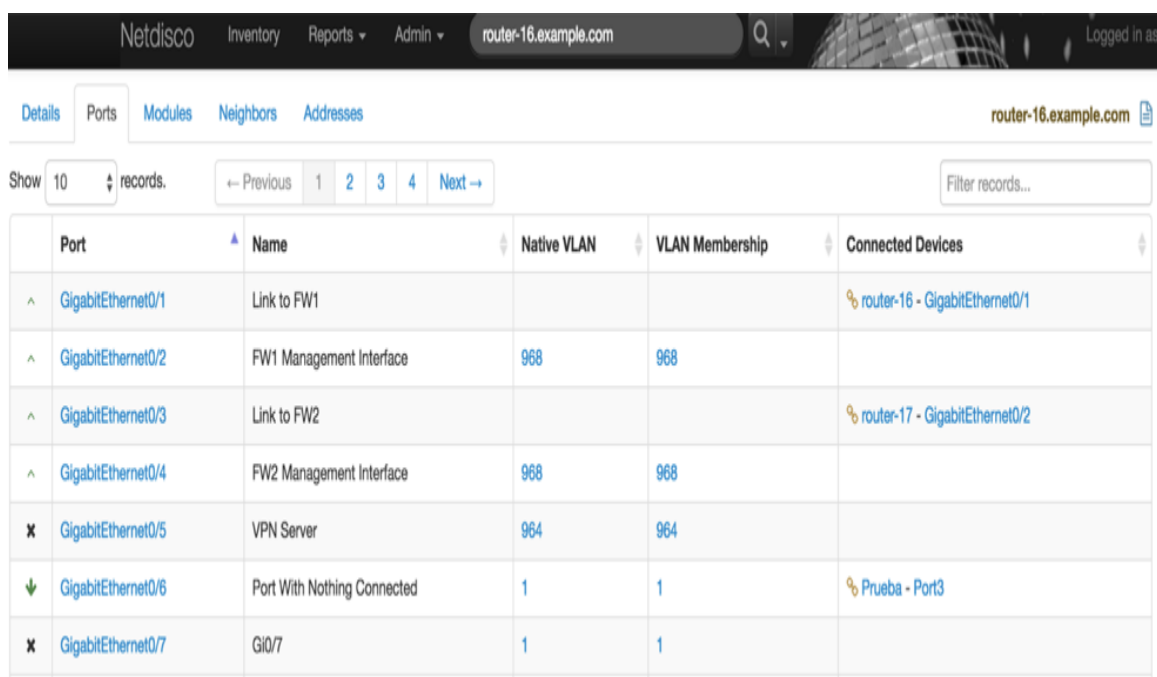
1 to 20 of 52 results

Previous **1** 2 3 Next

Figure 2-1. Rumble web interface

Figure 2-2 shows how Netdisco displays an example router and what ports are configured.

Netdisco is an SNMP-based L2/L3 network management tool designed for moderate to large networks. Routers and switches are polled to log IP and MAC addresses and map them to switch ports. Automatic L2 network topology discovery, display, and inventory.²



Port	Name	Native VLAN	VLAN Membership	Connected Devices
GigabitEthernet0/1	Link to FW1			router-16 - GigabitEthernet0/1
GigabitEthernet0/2	FW1 Management Interface	968	968	
GigabitEthernet0/3	Link to FW2			router-17 - GigabitEthernet0/2
GigabitEthernet0/4	FW2 Management Interface	968	968	
GigabitEthernet0/5	VPN Server	964	964	
GigabitEthernet0/6	Port With Nothing Connected	1	1	Prueba - Port3
GigabitEthernet0/7	Gi0/7	1	1	

Figure 2-2. Netdisco web interface

There are also manual steps and several methods that can aid initial collection of information.

Address Resolution Protocol (ARP) cache

An ARP cache is a collection of entries that are created when an IP address is resolved to a MAC address. Pulling the ARP cache from routers and switches will provide a list of IP and MAC addresses connected to the network.

Dynamic Host Configuration Protocol (DHCP)

Whatever device or software platform that has the role of handing out DHCP addresses will contain all IP address reservations and possibly hostnames.

Nmap

"**Nmap** ("Network Mapper") is a free and open source (**license**) utility for network discovery and security auditing. Many systems and network administrators also find it useful for tasks such as network inventory, managing service upgrade schedules, and monitoring host or service uptime."

Running `nmap` against the entire range of networks can provide an amazing amount of information as it is a very comprehensive scanning tool. A simple scan to get started is this:

```
nmap -v -Pn -O 192.168.0.0/16 -oG output.txt
```

This command provides a verbose output (`-v`, or `-vv` for even more verbosity), and assumes all hosts are online, skipping discovery (`-Pn`) with operating system detection (`-O`) on 192.168.0.0 through 192.168.255.255, and outputs the results in a grepable (searchable) format (`-oG`) to *output.txt*.

PowerShell

PowerShell is a very versatile and powerful tool in Windows environments. Starting with Windows 2008R2, the Active Directory Module was introduced. Prior to this, `dsquery`, `adsisearch`, and `ldap` commands were used to obtain information from AD.

```
Get-ADUser -filter *
```

can be used to return an entire list of users within an AD domain.

There are many filters that can be added to return subsets of this list as well. To retrieve all domain computer accounts, run

```
Get-ADComputer -Filter 'ObjectClass -eq  
"Computer"' | Select -Expand DNSHostName.
```

Simple Network Management Protocol (SNMP)

SNMP can provide a great amount of information on networked devices. By default, most SNMP version 2 read and write strings (the passwords used to query devices) are set to “public” and “private.” SNMP settings should be changed to custom strings, and if possible switched to SNMP version 3, which supports username and password authentication. Many software packages, such as NetDisco, which we mentioned earlier in this chapter, use SNMP to gather data.

Vulnerability management software

Data from a vulnerability scanner can be added into the asset management system. This assists in tracking risks and adding to the overall amount of useful information about assets.

Windows Management Interface (WMI)

WMI can be used to pull almost all important information from a Microsoft Windows device. Specific information can be gathered on components such as the CPU, memory, and disk, as well as system information such as OS, processes, and services. The command `WMIC` can be performed on a system locally, or the WMI API can be used with most programming languages—for example, remotely connecting with PowerShell uses the `Get-WmiObject` cmdlet.

Change Tracking

Keeping track of changes in hardware, software, and performance is a necessary step to having an up-to-date inventory. Certain modifications can significantly change the security risk for a device. Below are two different

examples of how changes can introduce new security vulnerabilities to an environment without the proper tracking.

Steven runs a section of the company that has just spun up a new service line. He has been working with marketing and they give him the okay to buy a domain and spin up a Wordpress site. Before configuring the site for the public, he installs Wordpress on his own machine and begins to install add-ons. First off, Wordpress itself has had a history of vulnerabilities, but add-ons can be written by anyone and can introduce vulnerabilities from privilege escalation to XSS (Cross-Site Scripting). An endpoint change-tracking client could trigger an alert on the addition of unapproved or unlicensed software.

Mark decides he's had enough with this company and quits. Before leaving, he removes an expensive piece of software from his laptop to keep the license key for his own personal use. When his equipment is repurposed, his replacement will more than likely need to use that same software. The asset management software should have the list of software that was present, as well as any corresponding licensing keys, the vendor contact information, and software renewal dates.

Monitoring and Reporting

Monitoring and reporting on assets provides notifications of upcoming software licensing renewals and hardware warranty expirations. Trends can be discovered with the right amount of information, taking guesswork out of creating yearly budgets and equipment procurement plans. This information can also be used to assist in any equipment refresh processes.

A helpful security measure to implement is the monitoring and alerting of any unapproved devices. In a perfect world, an alert would fire when a device MAC shows up that isn't located in the asset management tracking program. As mentioned at the beginning of this chapter, this level of control for assets can be extremely difficult in larger, more complex, and older environments.

Alerts may also be created for lack of certain software or system settings if, for example, an endpoint has no antivirus or management software installed, isn't encrypted, or has unauthorized software. More than likely this will be done with some sort of endpoint monitoring software, but can also be accomplished in a more manual route with some software. Microsoft Endpoint Configuration Manager (formerly SCCM) has the ability to report on installed software as well.

Asset Management Guidelines

In addition to the steps involved in implementing asset management, there is also a solid set of guidelines to keep in mind during the implementation process.

Automation

To accomplish these steps effectively, attempt to automate as many of them as possible. If any person along the chain of custody of an asset finds they are repeating a manual process, the question, "Can this be automated?" should be asked. The process should pull authoritative information from trustworthy locations on as many assets as possible. DNS can pull in hostnames and IP addresses; DHCP can tie MAC addresses to those IP addresses; and a vulnerability scanner may find entire networks that were previously unknown. Adding barcodes early on in the lifecycle can greatly assist with automation as well. Everything that can be automated leads to a more efficient process.

One Source of Truth

As there are many different ways to gather information about devices, such as DNS, DHCP, wireless connections, MAC address tables, software licenses, nmap scans, etc., it is important to select a software that will easily integrate with the technologies already present. Having conflicting information in several different locations like spreadsheets and SharePoint

is not conducive to a complete picture regarding current assets. When choosing a software or method, it should be well communicated that it alone is the one source of truth regarding assets, and any deviation should be dealt with accordingly (author's note: As much as I didn't want this to sound menacing, it totally does).

Organize a Company-Wide Team

Assets will enter the company from a variety of different areas. The purchasing department is the obvious first choice; third-party vendors may bring their own equipment; or there may be a BYOD (bring your own device) policy, which is a whole other can of worms. Types of departments that would benefit from being on an asset-management team include purchasing, receiving, helpdesk, communications, maintenance, and system administrators.

As with most other processes and procedures, it is close to impossible to plan for every possibility. Plan for the unplanned. If a member of the helpdesk team or another group that has access to the asset management software happens upon a device that has not been documented, there should be a process to deal with this. Not only should the asset then be added to the software, but the cause should also be investigated. Are there assets entering the organization in a different department or by different means that have yet to be added to the asset management process?

Executive Champions

The organizational team should also contain one or more members of the executive staff as a champion to assist in process and procedure changes that will cross through several departments. Larger organizations normally have difficulty with communicating changes and additions to procedures to the correct individuals, while smaller companies seem to resist change. Creating a well thought-out and communicated directive from someone other than security or IT staff will greatly increase the success. This

executive member will also be able to see the effects of proper asset management in the form of cost savings and avoidances.

Software Licensing

When it comes to software license management, knowing what you are entitled to have deployed is often more important than what you actually have deployed. More often than not, organizations fail software audits for over-deployment because they can't prove exactly what they have the right to have deployed. Ending up behind on software licensing can be a very expensive mistake. Not only will an updated list of currently installed software eliminate the risk of paying for software that isn't being used, but it also ensures that there are no licensing overage fees or fines.

Define Assets

Define criteria for what constitutes a critical asset—many times they may be the device where critical data lies, as well. It may be a specific type of hardware or appliance such as a head end firewall or fiber switches or certain custom software packages. Discovery and inventory will produce a large asset list. Some assets will require more oversight or management than others.

Documentation

Proper and complete documentation is an integral part of asset management. Creating and maintaining it should be a continual process from day one. Documentation is used to set clear directions and goals, as well as offering a continual reference as needed.

Spend sufficient time creating documentation and provide detailed descriptions of all security projects, including charts and statistics. These documents can be a major advantage when showing management where the security budget went. Another benefit of documentation is the knowledge that will be gained by all parties while creating it. Potential security holes or

weaknesses in the program may also become apparent during this process. Every mistake is a learning opportunity; document problems or mistakes so they are not repeated.

What should be documented? There are several levels of documentation depending on the size and scope of an environment. The following sections present a starting point as to what documentation will be beneficial.

Networking Equipment

Many automated scanning tools can provide a detailed overview of networking equipment, examples of which are outlined in the following lists:

- Hostname
- Licensing information
- Location
- Management IP
- Software, hardware, and firmware versions
- Warranty information

Network

- Ingress/Egress point public IP addresses for all sites
- ISP Account Information and Contacts
- Performance baselines of network traffic over a day/week/month period
- Default Gateways

Servers

- Applications and Roles
- Department or group that manages/supports
- Hostname
- iLO address
- IP address(es)
- Is remote access allowed?
- Is there PII or other sensitive data?
- OS version
- Open ports
- Performance baselines of CPU, Memory, & Disk
- Warranty Information

Desktops

- Hostname
- Department
- Patch level

Users

Not only should the individual accounts be documented, but also what users have access to each. Many times this can be tracked in an Identity and Access Management (IAM) solution or in a more manual process:

- Database Administrator Accounts
- Domain/Enterprise/Schema Admins (and other admin level accounts)
- Root and Administrator Accounts

- Service Accounts

Applications

- Administrative Users
- Licensing
- Servers and Appliances involved
- Type of Authentication
- Workflow of data (explain further)

Cloud Assets

Many cloud providers have their own automated asset management and tracking systems built into their solution by default. However if you have assets spread over multiple providers or less advanced providers where you have to track your own items you'll want to remember the following:

- Resources
 - Compute Engine virtual machines (VMs)
 - Cloud Storage buckets
 - App Engine instances
- Policies
 - IAM policy
 - Access Context Manager policy

Other

- Certificates and Expiration dates
- Domains and Expiration dates

Just as important as the documentation itself is the consistency and organization of the information. Naming in a consistent fashion assists in locating and understanding the roles of assets. For example:

- ORG1-DC1-R2B-RTR3 = Organization 1, Datacenter 1, Row 2, Rack B, Router 3
- SVC_ORG2-SQL10-SNOW = The service account for Organization 2, SQL Server 10, the Snow Application
- ORG3-FL1-JC-AP3 = Organization 3, Floor 1, JC Hall, Wireless Access Point 3

Conclusion

Classify, organize, automate, define, gather, track, monitor, report, document, rinse, lather, repeat. The messy world of asset management can be a daunting task without a solid plan and the understanding that it's not a one-time project. Having as much information in one place as possible about endpoint and infrastructure devices will not only assist in short-term troubleshooting, but also long-term design planning and purchasing decisions.

1 "Information Classification—Who, Why, and How", SANS Institute InfoSec Reading Room.

2 <https://sourceforge.net/projects/netdisco/>

Chapter 3. Policies

A NOTE FOR EARLY RELEASE READERS

With Early Release ebooks, you get books in their earliest form—the authors’ raw and unedited content as they write—so you can take advantage of these technologies long before the official release of these titles.

This will be the 3rd chapter of the final book. Please note that the GitHub repo will be made active later on.

If you have comments about how we might improve the content and/or examples in this book, or if you notice missing material within this chapter, please reach out to the editor at sevans@oreilly.com.

Policies are one of the less glamorous areas of information security. They are, however, very useful and can be used to form the cornerstone of security improvement work in your organization. In this chapter we will discuss why writing policies is a good idea, what they should contain, and the choice of language to use.

Why are policies so important? There are a range of reasons:

Consistency

Concerns about inconsistent approaches from day to day or between members of staff should be vastly reduced in the wake of decent policies. A written set of policies reduces the need to make a judgment call, which can lead to inconsistent application of rules.

Distribution of knowledge

It is all well and good for *you* to know what the policy is with regards to not sharing passwords with others, but if the entire organization is

unaware, then it is not providing you much benefit. Policy documents disseminate information for others to consume.

Setting expectations

Policies set rules and boundaries; by having clearly defined rules, it becomes equally clear when someone breaks those rules. This enables appropriate action to be taken. Departments like human resources find it difficult to reprimand someone because it “feels like” they may have done something wrong. A clear contravention of a rule is easier to enforce.

Regulatory compliance and audit

Many industries are regulated or pseudo-regulated, and many have auditors. A criteria common amongst nearly every regulatory compliance or auditing scheme is the existence of policies. By having a set of policies, you have already ticked a box on the regulatory compliance or audit checklist.

Sets the tone

The policy set can be used to set the overall tone of a company’s security posture. Even if not explicitly laid out, the policy set gives an overall feel as an organization’s approach to security.

Management endorsement

A management-endorsed policy, published within an organization’s official document library, lends credibility to the policy set itself and by extension to the security team as well.

Policies are living documents—they should grow with an organization and reflect its current state. Making changes to policy should not be frowned upon; evolution of both the policies themselves and the associated documentation is a positive change. A scheduled annual review and

approval process of policies will allow you to ensure that they remain aligned with business objectives and the current environment.

Language

Policies should lay out *what* you, as an organization, wish to achieve in a series of policy statements. Detail as to specifically *how* this is achieved is outlined in procedure and standards documentation. For this reason there is no need to get caught up with complexity and detail. Policy statements should be fairly simple, clear, and use words like “do,” “will,” “must,” and “shall.” They should not be ambiguous or use words and phrases such as “should,” “try,” and “mostly.”

For example, a good policy will use statements such as:

A unique User ID shall be assigned to every user.

As opposed to

A unique User ID should be assigned to a user.

The use of “should” as opposed to “shall” gives the impression that this is a “nice to have,” not a rule. If there are times when a policy can be overridden, then this should be stated as part of the policy statement. This is often achieved by using phrases such as “unless authorized by a manager.” Care should be taken not to introduce ambiguity with such statements, however; for example, it must be clear what constitutes “a manager” in this case.

Documents should be designed to be read. There is no need to fill documents with excessively wordy statements or some kind of confusing legalese. Each policy statement can be only a few sentences, often only one, in a bullet point format.

Document Contents

Policy documents should contain a few key features:

Revision control

At the very least, this should include a version number and an effective date for the document. This allows a user in possession of two versions of a document to quickly establish which is the current version and which is out of date and no longer applicable.

Revision detail

A brief summary of what has changed since the last revision allows approvers and those already familiar with the policy to quickly understand changes and the new content.

Owner/approver

Being clear as to who owns and approves any particular document is useful not only for demonstrating that it has been accepted and agreed upon by the appropriate level of management, but it also serves to facilitate feedback and suggestions for updates in future revisions.

Roles and responsibilities

Defining whose responsibility it is to implement, monitor, abide by, and update policies ensures that there is little room for ambiguity with regard to roles.

Executive sign-off

By ensuring that executive sign-off is clearly marked on each document it is clear to the reader that it is endorsed at the highest level and approved for immediate use.

Purpose/overview

This provides a brief overview as to what the policy document covers. This is typically only a paragraph and is intended to allow the readers to gauge if they are looking at the correct policy document before they get to the point of reading every policy statement.

Scope

In all likelihood, the scope section will only be a couple of sentences and will be the same for most policy documents. This explains who the policy document applies to; for example, “*This policy applies to all <Company Name> full-time employees, part-time employees, contractors, agents, and affiliates.*” Of course, there could be policies that only apply to a particular subset of readers for some reason, and the scope can be adjusted accordingly.

Policy statements

As discussed earlier, these are the guts of the document—they are the policies themselves.

Consistent naming convention

Consistent naming conventions not only for the documents themselves, but also for artifacts they reference, ensure that they are easy to understand and can be applied consistently across the organization.

Related documents

Cross references to other relevant documents such as standards, policies, and processes allow the reader to quickly locate related information.

For ease of reference during an audit, it is prudent to also include references to sections of any relevant regulatory compliance, standards, and legal requirements.

Topics

For ease of reading, updating, and overall management it is probably easier to produce a set of policy documents rather than a single monolithic document.

Selecting how the policies are broken up is, of course, a matter of determining what is most appropriate for your organization. You may have a favorite security framework, such as ISO 27002, for example, from which you can draw inspiration. Similarly, aligning policy topics with a particular regulatory compliance regime may be more aligned with your organization's objectives. In reality, there are many high-level similarities between many of the frameworks.

SANS, for example, publishes a list of **template policies** that you can edit for your own needs. At the time of writing, its list of topics are:

- Acceptable Encryption Policy
- Acceptable Use Policy
- Acquisition Assessment Policy
- Analog/ISDN Line Security Policy
- Anti-Virus Guidelines
- Automatically Forwarded Email Policy
- Bluetooth Baseline Requirements Policy
- Clean Desk Policy
- Communications Equipment Policy
- Data Breach Response Policy
- Database Credentials Policy
- Dial In Access Policy
- Digital Signature Acceptance Policy
- Disaster Recovery Plan Policy
- DMZ Lab Security Policy
- Email Policy

- Email Retention Policy
- Employee Internet Use Monitoring and Filtering Policy
- End User Encryption Key Protection Plan
- Ethics Policy
- Extranet Policy
- Incident Handling - Chain Of Custody Form
- Incident Handling Forms
 - Incident Communication Log
 - Incident Contacts List
 - Incident Containment
 - Incident Eradication
 - Incident Identification
 - Incident Survey
- Information Logging Standard
- Intellectual Property Incident Handling Forms
 - Incident Communication Log
 - Incident Contacts
 - Incident Containment
 - Incident Eradication
 - Incident Form Checklist
 - Incident Identification
- Internet DMZ Equipment Policy

- Internet Usage Policy
- Lab Anti Virus Policy
- Lab Security Policy
- Mobile Device Encryption Policy
- Mobile Employee Endpoint Responsibility Policy
- Pandemic Response Planning Policy
- Password Construction Guidelines
- Password Protection Policy
- Personal Communication Devices and Voicemail Policy
- Remote Access Mobile Computing Storage
- Remote Access Policy
- Remote Access Tools Policy
- Removable Media Policy
- Risk Assessment Policy
- Router and Switch Security Policy
- Security Response Plan Policy
- Server Audit Policy
- Server Malware Protection Policy
- Server Security Policy
- Social Engineering Awareness Policy
- Software Installation Policy
- Technology Equipment Disposal Policy
- Virtual Private Network Policy

- Web Application Security Policy
- Wireless Communication Policy
- Wireless Communication Standard
- Workstation Security (For HIPAA) Policy

This is not an atypical list; however, many of the policies listed will not apply to your organization. This is completely fine.

Storage and Communication

The nature of policies and procedures is meant to lend as much standard communication as possible to the organization as a whole. To do this, policies must be easily accessible. There are many software packages that can not only provide a web interface for policies, but also have built-in review, revision control, and approval processes. Software with these features makes it much easier when there are a multitude of people and departments creating, editing, and approving policies.

Another good rule of thumb is to, at least once per review process, have two copies of all policies printed out. As the majority of them will be used in digital format, there will be many policies that refer to and are in direct relation to downtime or disaster recovery procedures. In cases such as these, they may not be accessible via digital media so having a backup in physical form is best.

Conclusion

Policies are important tools used to express the direction of an organization from a security perspective, clearly articulating expectations and providing a level of consistency. They can also be used to explicitly state and enforce rules that have previously been ambiguous or inferred.

Policies are not set in stone forever—they are living documents that can grow and change in line with your organization.

About the Authors

Lee Brotherston is a senior security advisor, providing information security consulting services to a range of clients. Having spent nearly two decades in information security, Lee has worked as an internal security resource across many verticals, including finance, telecommunications, hospitality, entertainment, and government, in roles ranging from engineer to IT security manager.

Amanda Berlin is an Information Security Architect, co-hosts the “Brakeing Down Security” podcast, and writes for several blogs.