

Definitive GuideTM to *Security Intelligence and Analytics*

Find and Stop Attacks Sooner to Prevent
Data Breaches and Minimize Damage



Karen Scarfone, CISSP, ISSAP
Steve Piper, CISSP

FOREWORD BY:

Robert Lentz

Former CISO for the US Department of Defense

Compliments of:

LogRhythm[®]

The Security Intelligence Company

About LogRhythm

LogRhythm, a leader in security intelligence and analytics, empowers organizations around the globe to rapidly detect, respond to, and neutralize damaging cyber threats. The company's award-winning platform unifies next-generation SIEM, log management, network and endpoint monitoring and forensics, and security analytics. In addition to protecting customers from the risks associated with cyber threats, LogRhythm provides innovative compliance automation and assurance, and enhanced IT intelligence. For more information, visit www.logrhythm.com.

Definitive GuideTM to *Security Intelligence and Analytics*

Karen Scarfone, CISSP, ISSAP
Steve Piper, CISSP

Foreword by Robert Lentz, Former CISO
for the U.S. Department of Defense



CYBEREDGE
P R E S S

Definitive Guide™ to Security Intelligence and Analytics

Published by:

CyberEdge Group, LLC

1997 Annapolis Exchange Parkway

Suite 300

Annapolis, MD 21401

(800) 327-8711

www.cyber-edge.com

Copyright © 2016, CyberEdge Group, LLC. All rights reserved. Definitive Guide™ and the CyberEdge Press logo are trademarks of CyberEdge Group, LLC in the United States and other countries. All other trademarks and registered trademarks are the property of their respective owners.

Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without the prior written permission of the publisher. Requests to the publisher for permission should be addressed to Permissions Department, CyberEdge Group, 1997 Annapolis Exchange Parkway, Suite 300, Annapolis, MD, 21401 or transmitted via email to info@cyber-edge.com.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on CyberEdge Group research and marketing consulting services, or to create a custom *Definitive Guide* book for your organization, contact our sales department at 800-327-8711 or info@cyber-edge.com.

ISBN: 978-0-9961827-4-4 (paperback); ISBN: 978-0-9961827-5-1 (eBook)

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgements

CyberEdge Group thanks the following individuals for their respective contributions:

Editor: Susan Shuttleworth

Designer: Debbi Stocco

Production Coordinator: Valerie Lowery

Table of Contents

Foreword	v
Introduction	vii
Chapters at a Glance	vii
Helpful Icons.....	viii
Chapter 1: Surveying the Cyberattack Lifecycle	1
Basic Terminology.....	1
The Cyberattack Lifecycle	2
Phase 1: Reconnaissance	2
Phase 2: Initial compromise	3
Phase 3: Command & control	3
Phase 4: Lateral movement.....	4
Phase 5: Target attainment	4
Phase 6: Exfiltration, corruption, and/or disruption	5
Chapter 2: Understanding Threat Management	7
Threat Management Processes	7
Detection	9
Forensic data collection and processing	9
Discovery through security analytics	10
Qualification	10
Response	11
Investigation	11
Mitigation	12
Recovery	12
Security Intelligence and Analytics Platform	13
The Role of Threat Intelligence	14
Chapter 3: Collecting and Processing Forensic Data	15
Data Generation	16
Enterprise security controls	16
Endpoint software	17
Network flow data	17
Asset data.....	18
Data Transfer	18
Minimization	18
Protection	19
Data Normalization.....	20
Extraction of key data fields	21
Standardization of values.....	21
Timestamp standardization	21
Event classification.....	22
Data Archiving	22
Chapter 4: Automating Discovery through Security Analytics	23
Search Analytics	23
Leverage dashboards.....	24
Drill down for details.....	25
Use search capabilities	25
Use visualization techniques.....	26

Machine Analytics.....	27
Establish baselines	27
Detect threats	28
Prioritize threats.....	30
Chapter 5: Qualifying Security Intelligence.....	33
Alert Analysis	33
Evaluate each alert’s validity.....	33
Improve detection capabilities.....	34
Risk Level Assessment	35
Importance of the target	36
Current attack lifecycle phase	36
Incident Declaration and Prioritization	36
Chapter 6: Streamlining Threat Response Processes	37
Incident Management and Threat Investigation	38
Workflow and collaboration facilitation.....	39
Secure collection of supporting data.....	40
Threat Mitigation.....	41
Common mitigation techniques.....	41
Automated mitigation	42
Chapter 7: Selecting the Right Solution	43
Usability	44
Scalability and Flexibility.....	44
Logging Source Support.....	45
Supplemental Forensic Data Collection	45
Machine Analytics.....	46
Search Analytics	47
Threat Intelligence Service Choices.....	47
Automated Investigation and Mitigation Capabilities.....	48
Customization	48
Technical Support.....	49
Chapter 8: Steps for Successful Implementation	51
Preparation and Planning.....	52
Step 1: Define goals and requirements for threat management.....	52
Step 2: Create and validate policies supporting the requirements	53
Step 3: Prioritize the implementation of threat management	54
Solution Design	54
Step 4: Design a threat management architecture	55
Step 5: Evaluate products and services for the architecture	55
Production Implementation	57
Step 6: Acquire, deploy, and integrate products and services.....	57
Step 7: Gradually transition log sources to the solution.....	58
Step 8: Develop processes and train staff on the solution.....	59
Step 9: Customize dashboards, mitigations, alerting, etc.	59
Maintenance.....	60
Step 10: Refine the solution to improve its performance.....	60
Glossary.....	61

Foreword

In my 10 years as the Chief Information Security Officer (CISO) for the largest information enterprise in the world, the U.S. Department of Defense, we realized after numerous cyber incidents that victim organizations did not possess the tools, processes, staff, or mindset necessary to detect and respond to advanced intruders.

Accordingly, we developed the Cyber Security Maturity Model to create a long-term strategic commitment and an ability to measure tactical performance while institutionalizing a risk management culture. At the heart of that Cyber Security Maturity Model was an intense commitment from every level to continuously improve our ability to detect, respond to, and neutralize cyber threats.

Today, cyber adversaries are more sophisticated, organized, and capable than ever before. Their targets range from nation states and global conglomerates to small manufacturers, regional healthcare organizations, and credit unions. If they want to get in, they will, regardless of the prevention measures put in place to keep them out.

The good news is that a breach of a network does not immediately equate to data loss or service disruption. There are common steps cyber adversaries typically take on the path to achieving their end goal, and there are ways to detect those steps early in the journey.

Organizations that employ a cybersecurity strategy that combines comprehensive visibility, continuous monitoring, advanced analytics and efficient incident response orchestration are well positioned to identify and respond to the early indicators of an intruder, and neutralize the threat before it can result in a material cyber incident.


Security intelligence and analytics platforms offer the ideal centerpiece for a security operation designed to address today's cyber threat landscape. CyberEdge Group's *Definitive Guide to Security Intelligence and Analytics* provides a

concise and useful description of the cyberattack lifecycle and guidance on how organizations can leverage a security intelligence and analytics platform to substantially improve their cybersecurity posture.

Robert Lentz

Former CISO for the U.S. Department of Defense

Introduction



As members of the security community, we're already painfully aware of the sharp increase in major data breaches from glancing at the headlines, getting requests from management to explain the latest security issue, or even receiving notices that our own personal data has been compromised. We don't need to be told that there's a huge problem, and that it's only getting worse.

So what can we do? Many pundits say that compromises are inevitable, so we should shift our focus from prevention to detection. While there's some truth to that, prevention is still incredibly important. Ignoring prevention makes it a breeze for any attacker to succeed.

A saner strategy is to balance controls for prevention and detection. Use preventive controls to stop less-skilled attackers, which reduces the noise, and use detective controls to expedite identification of advanced attackers, preventing them from inflicting major data breaches and causing other significant damage. On top of these controls sits a security intelligence and analytics platform for putting all the pieces together.

This book is intended for anyone with responsibilities related to detecting, responding to, and recovering from major cyberattacks and the cyberthreats behind them.

Chapters at a Glance

Chapter 1, "Surveying the Cyberattack Lifecycle," explains the phases of today's sophisticated cyberattacks as the basis for understanding the need for threat management.

Chapter 2, "Understanding Threat Management," outlines processes for managing the cyberthreats that employ the cyberattack lifecycle against your organization.

Chapter 3, "Collecting and Processing Forensic Data," discusses the actions involved in gathering and standardizing forensic data for cyberthreat management purposes.

Chapter 4, “Automating Discovery through Security Analytics,” explains how security analytics are performed and used in support of cyberthreat management.

Chapter 5, “Qualifying Security Intelligence,” underscores the importance of assessing security intelligence to validate an incident and prioritize the organization’s initial response to it.

Chapter 6, “Streamlining Threat Response Processes,” explores techniques for improving the efficiency of the three major components of threat response: incident management, threat investigation, and threat mitigation.

Chapter 7, “Selecting the Right Solution,” provides recommendations for evaluating security intelligence and analytics platforms for use within your organization.

Chapter 8, “Steps for Successful Implementation,” describes the most important steps for scoping, designing, and deploying a security intelligence and analytics platform.

Glossary defines the key terms (shown in *italics*) used in this book.

Helpful Icons



TIP
Tips provide practical advice that you can apply in your own organization.



DON'T FORGET
When you see this icon, take note as the related content contains key information that you won't want to forget.



CAUTION
Proceed with caution because if you don't it may prove costly to you and your organization.



TECH TALK
Content associated with this icon is more technical in nature and is intended for IT practitioners.



ON THE WEB
Want to learn more? Follow the corresponding URL to discover additional content available on the Web.

Chapter 1

Surveying the Cyberattack Lifecycle

In this chapter

- Define basic terminology related to cyberattacks and cyberthreats
- Understand the phases of the cyberattack lifecycle
- See real-world implications of the cyberattack lifecycle

To understand the insights and recommendations about detecting and stopping cyberattacks made throughout the rest of this book, it's important to first grasp how sophisticated cyberattacks are performed today.

Basic Terminology

A *cyberattack* is an attempt to negatively affect the security of computing resources, and a *cyberthreat* is an entity (individual, group, nation state, etc.) that plans and executes cyberattacks. For brevity, this book drops “cyber” from these terms. A person who performs attacks is known as an *attacker* or a *threat actor*.

The result of a successful attack is a *compromise*: in other words, a loss of confidentiality, integrity, and/or availability of data, systems, networks, or other computing resources. A *data breach* is a compromise that causes a loss of data confidentiality.



Although many compromises are caused by intentional attacks, some occur because of human error. For simplicity, this book uses the term “attacker” regardless of the person’s

intent, because the same tools and techniques are used to detect and respond to both intentional and unintentional attacks.

A final important term to know is *incident*. A compromise indicates an attack has succeeded. An incident encompasses not only a successful attack, but also an attack in progress, reconnaissance activities, and failed attacks of particular concern, such as indications of a new, serious threat against the organization.

The Cyberattack Lifecycle

Cyberattacks can be performed in many ways, but serious attacks designed to breach sensitive data tend to follow the same pattern. This pattern is known as the *cyberattack lifecycle*.

This lifecycle has six phases:

1. Reconnaissance
2. Initial compromise
3. Command & control
4. Lateral movement
5. Target attainment
6. Exfiltration, corruption, and/or disruption

Attacks have evolved this way to overcome the complex, layered defenses protecting stores of sensitive data. Reaching this data often requires a series of compromises over an extended period. Let's look at each of the lifecycle phases to better understand how the entire attack progresses from start to finish.

Phase 1: Reconnaissance

An attack starts with the attacker's choice of what to accomplish, such as financial gain, critical infrastructure disruption, or political message distribution. Next, the attacker identifies an organization and one or more of its systems that can be taken advantage of to achieve this goal, such as gaining access

to credit card numbers held by a retailer, or damaging a company's reputation by revealing the details of its executives' emails. These systems of particular interest to the attacker are called *targets*.

Once a target is selected, the attacker figures out a good way to gain entry to the organization's internal systems or networks. This usually involves conducting research to learn more about the target's environment, such as identifying weaknesses in the organization's Internet-facing systems and perimeter security controls. Such research is better known as *reconnaissance*.

During the reconnaissance phase, the attacker performs any necessary preparation for the initial attack, such as acquiring or writing exploit code, crafting spearphishing emails and associated websites, planning physical theft of equipment, or collaborating with an insider.

Phase 2: Initial compromise

The second phase of the attack lifecycle involves penetration of the organization's perimeter to gain internal access. Attackers most often accomplish this by compromising user credentials, such as acquiring an employee's username and password through a spearphishing attack, or compromising a system with malware or attack tools.

DON'T FORGET



Although user endpoints, such as desktops, laptops, smartphones, and tablets, are frequently the focus of the initial compromise, attackers also look for networked devices that lack robust security controls, including point of sale terminals, medical devices, and printers/copiers.

Phase 3: Command & control

After performing the initial compromise, the attacker takes measures to ensure continued access to the compromised internal system. For example, the attacker might create a new user account to retain access to the system even if the stolen credentials originally used to gain access are changed.

TECH TALK



Similarly, the attacker often installs additional tools on the compromised system to enable direct remote access to it. This gives the attacker easy access to the organization's internal

network because the tools can disguise their communications with the attacker to look like normal user-initiated activity. So the attacker can go right through the perimeter to access the compromised system without raising suspicion.

Phase 4: Lateral movement

After establishing covert remote access to a system on the internal network, the attacker leverages that access to reach other internal systems and attempts to compromise them as well. This process typically involves a chain of compromises, where system A is used to access system B, system B is used to access system C, and so on, until the attacker is in close proximity to the target itself. This chain of compromises is known as *lateral movement*.

Lateral movement is often the most complex and time-consuming phase of the cyberattack lifecycle. The attacker may have to repeatedly find and exploit weaknesses in other systems without being detected, all the while making progress toward the ultimate target.

For each system compromise, the attacker performs actions similar to the command & control phase. The attacker establishes covert remote access to the device by setting up additional user accounts, installing backdoors, etc.

CAUTION



Detection of a single compromise isn't necessarily the end of the game. As long as the attacker has made each compromise look like an isolated incident, the other compromises are unlikely to be detected, and the attacker will still have remote access to all the other compromised systems. Detection of a single compromise within the cyberattack lifecycle is often just a minor setback for the attacker.

Phase 5: Target attainment

In this phase, the attacker makes a final lateral move and reaches the targeted system. The attacker may need to perform additional compromises within that system, such as escalating privileges, to gain access to sensitive data stored on the system or to issue commands with administrator-level privileges.

Phase 6: Exfiltration, corruption, and/or disruption

In the final phase of the cyberattack lifecycle, the attacker performs the ultimate exploitation of the target. All of the attacker's other compromises while moving through the organization have been largely incidental, performed simply to make the final exploitation possible.

The target's exploitation may take many forms. For example, suppose that an attacker is targeting an organization's e-commerce operations. Possible results include the following:

- ✓ **Exfiltration:** A breach of stored credit card information and customer information that enables identity fraud
- ✓ **Corruption:** Alterations to records that allow the attacker to obtain free services or goods
- ✓ **Disruption:** A complete disruption to IT operations, causing the organization to lose revenue

Of these three, exfiltration is by far the most common result. *Exfiltration* is the process of transferring sensitive information from an authorized location (controlled and protected by the organization) to an unauthorized location outside its control.



Attackers can perform exfiltration in many ways, but they often use the covert remote access channels that they've already established within the organization. These channels may conceal their activity from monitoring, such as by encrypting communications. They may also be capable of exfiltrating data slowly, over an extended period, to avoid major changes in bandwidth usage that can set off internal alarms.

We've got good news and bad news

First, the bad news: it's impossible to keep attackers out of your organization's systems and networks. For one thing, insiders perform many data breaches, and they already have system and network access, sometimes even privileged access.

But when it comes to external attackers, the situation isn't much better. An attacker can readily compromise a seemingly unimportant endpoint, such as a random user's laptop or smartphone, through malware, social engineering, or other common techniques. This endpoint may have no direct relationship with the ultimate goal – to steal the contents of the organization's most valuable database – but it gives the attacker a foot in the door.

The cyberattack lifecycle requires the attacker to be highly knowledgeable and patient. Such an individual can perform a series of attacks over a period of months or even years to eventually realize his or her goal. Organizations often fail to detect any of the signs of compromise, also known as

indicators of compromise, from the lifecycle phases preceding the ultimate exploitation of the target.

Even when an organization detects a part of the overall attack, in many cases it won't make the connection with other parts of the attack. Often the attacker has made an effort to isolate each compromise – for example, by frequently switching the IP address used to control the compromised internal systems. This failure to connect the dots allows the overall attack to continue, ultimately resulting in a major data breach, operational disruption, or other highly negative consequence.

Now, the good news: the information in this book can help you stop many of these attackers. There's no such thing as perfect security, so there's no way to stop every attacker. But by focusing more of your organization's efforts and resources on detecting attacks in progress, you're much more likely to prevent serious damage and keep your organization's name out of the headlines.

Chapter 2

Understanding Threat Management

In this chapter

- Learn the basics of threat management processes
- Become familiar with security intelligence and analytics platforms
- Understand the role of threat intelligence

Chapter 1, “Surveying the Cyberattack Lifecycle,” explained the attack lifecycle and its real-world impact. This chapter defines processes for managing, to the degree possible, the threats associated with the attack lifecycle.

Threat Management Processes

Threat management comprises three ongoing processes:

- ✓ Detecting threats targeting the organization
- ✓ Responding to detected threats
- ✓ Recovering from damage caused by threats



These processes, which the rest of this chapter discusses in detail, don't eliminate the need to mitigate vulnerabilities and otherwise defend the organization from exploitation. Rather, they acknowledge that a sophisticated, determined threat will eventually have some degree of success, even against strong defenses. If defenses are weak, threats can easily compromise systems at will.

The goal of threat management is to minimize the damage caused by successful compromises of an organization's data, systems, networks, and other computing resources. Minimization includes preventing threats' ultimate targets from being compromised and limiting incidental damage to intermediate systems.

DON'T FORGET



A determined threat can't be permanently stopped, but it can be slowed and discouraged to the point where it may abandon its efforts to compromise your organization and switch targets instead.

The key to threat management is to stop attacks as early in the attack lifecycle as possible. Obviously, the sooner a threat is detected, the sooner it can be addressed to prevent additional damage. What's important to remember is that many threats go undetected for weeks, months, or even years – if they're ever detected at all.

In far too many cases, the first indication of a major data breach or other compromise comes from outside parties – for example, banks observing fraud from stolen credit card numbers. By the time a report is sent to the organization, some or all traces of the chain of attacks leading to the breach may be gone, making it impossible to determine how the breach occurred. This delay may also thwart identification of which user accounts, systems, etc. were compromised and might still be under an attacker's control.

This scenario shows why threat management is critically important: not only because it detects attacks in progress so they can be stopped before major data breaches and other compromises occur, but also because it's key to identifying and remediating all the other damage to systems, user accounts, etc. If a single backdoor left by a threat isn't found, the threat may be able to re-enter the organization's systems and networks at will, making future data breaches highly likely.

Let's take a closer look at the three components of threat management: threat detection, response, and recovery. Each of these components must function effectively and efficiently in close coordination with the others to minimize the negative impacts of threats.

Detection

The detection component of threat management is more complex than simply finding evidence of a threat acting against the organization. Detection involves the following:

- ✓ Forensic data collection and processing (gathering data on security-relevant events throughout the enterprise, then standardizing the data formats)
- ✓ Discovery through security analytics (analyzing the collected data to identify potential compromises)
- ✓ Qualification (validating the potential compromises)

The outcome of detection is high-quality, actionable information about the most serious threats currently acting against the organization. This information is better known as *security intelligence*.

Forensic data collection and processing

Successful threat detection relies on identifying potential signs of those threats as they perform reconnaissance, penetrate perimeter defenses, compromise systems and accounts, and move from system to system within the organization.



It's important to collect not just obvious indications of successful attacks, but also any other security-related events that could be related in some way to threats. All of this event data is collectively referred to as *forensic data*. For example, suppose that an operating system logs an administrator account being used to copy sensitive data to another system. Is this a legitimate action or an attempt at exfiltration? To answer this question, you need context for the event, such as what happened immediately before it.

Most organizations have huge volumes of security-related event data that need to be analyzed for threat management purposes. This data comes from four categories of sources: enterprise security control logs, endpoint software logs, network flow data, and asset data.

This disparate data is collected in a centralized location, but it's not of much use unless it's converted from the original data formats to a universal format. This conversion involves several processing functions, including extracting key data fields and standardizing values. Once the data is in a universal format, it's an incredibly rich source of raw information regarding threats and their actions.

See Chapter 3, “Collecting and Processing Forensic Data,” for more information on gathering and standardizing forensic data for threat management purposes.

Discovery through security analytics

Organizations should continually use *security analytics* techniques on their forensic data to find the events and sequences of events that are of greatest concern from a security perspective. There are two types of security analytics. *Search analytics* are performed by a person, and *machine analytics* are performed automatically by a system or systems.

Because of the sheer volume of events needing review, organizations must rely heavily on machine analytics, which use a variety of techniques to identify and prioritize suspicious activities. Since they're so labor intensive, search analytics are largely performed on an as-needed basis, such as searching for events with a particular characteristic – for example, a source IP address associated with other attacks. Organizations also often use dashboards to monitor security events at a high level; observing these dashboards and drilling down into events on the dashboards is another form of search analytics.

The output of the discovery step is the generation of security intelligence. For more information on how security analytics are performed and used, see Chapter 4, “Automating Discovery through Security Analytics.”

Qualification

Qualification involves assessing the security intelligence produced by the previous step to confirm its legitimacy and priority – basically verifying that the detected activity necessitates a response.



Qualification is a time-intensive manual process performed by experienced personnel, who must be available around the clock to review the latest security intelligence and confirm that there's a serious security problem that needs to be addressed quickly. Assigning less-knowledgeable people to the qualification role is a recipe for disaster because it will frequently lead to wrong actions or no action at all.

Whenever feasible, organizations should rely on machine analytics instead of manual qualification to conserve personnel resources and to enable faster decision making.

The output of qualification is verified security intelligence indicating that the organization's response capabilities need to address the detected activity. In other words, qualification may result in declaring that an incident has occurred or is about to occur.

More information on qualification is provided by Chapter 5, "Qualifying Security Intelligence."

Response

The detection component of threat management results in identification and characterization of an incident caused by a threat. The next threat management component, response, is the initiation of actions to control and stop the incident as well as thwart the detected threat.

Response actions start with an investigation of the security intelligence associated with an incident, and conclude with mitigation of the threat or threats captured by that security intelligence. All of these response actions are planned and tracked through incident management processes.

Investigation

During an *investigation*, security administrators review the incident's related security intelligence, such as analyzing the alarms triggered by the potential threat, to determine how the threat should be handled. This review often seeks broader patterns that could indicate a wider compromise in progress. For example, an alert for an attacker moving from one system to another might indicate only one in a series of lateral movements through the enterprise.

Chapter 6, “Streamlining Threat Response Processes,” offers additional information on investigation.

Mitigation

In *mitigation*, the organization takes actions to thwart the threat by stopping its in-progress attacks. Examples of mitigation actions include disabling a compromised user account, taking a malware-infected system offline for remediation measures, and blocking all network communications involving a particular IP address (for example, the command and control hosts for a major botnet).

Mitigation is covered in more detail in Chapter 6, “Streamlining Threat Response Processes.”

Recovery

Mitigation is not the last step in threat management. Mitigation stops the current attack, but it doesn’t help prevent a similar future attack, nor does it help the organization recover from the damage caused by the threat. This is where recovery comes in.



Although *recovery* is defined as the third component of threat management, it shouldn’t necessarily be postponed until mitigation has been completed. It’s common for mitigation and recovery actions to go on simultaneously; for example, system administrators patch vulnerable laptops while other administrators collect and rebuild laptops that were already compromised.

Organizations often need to perform administrative measures in addition to technical and operational measures to address damage caused by threats. For example, when a data breach involving customer records occurs, the organization generally needs to notify its customers of the breach and offer them credit monitoring services.



Further discussion of measures for recovery is outside the scope of this book. A great resource is CERT’s incident management site at <https://www.cert.org/incident-management/csirt-development/>.

Time keeps on slippin'

It's hard to quantify the success of an organization's threat management processes. There's no way of knowing what the damage would have been, which additional systems or accounts would have been compromised, etc.

Instead, organizations should focus on measuring the responsiveness of their threat management detection and response components. These measurements are known as mean time to detect (MTTD) and mean time to respond (MTTR). MTTD indicates the time elapsed from the start of an attack or chain of attacks until it was noticed by the organization. MTTR indicates how long it took from the organization's initial detection of the attack to

complete all associated response activities.

Organizations should strive to reduce their MTTD and MTTR values by improving their threat management capabilities. The lower the values of these measurements are, the more effective and efficient the organization is at stopping attacks in progress and limiting their damage.

Typically, the recovery component isn't measured as a whole because it's so different from case to case. However, certain parts of recovery can be measured, such as how long it takes to notify customers that their data has been breached after the attack is discovered.

Security Intelligence and Analytics Platform

So far we've focused on threat management processes, but these processes must be heavily automated to keep up with the volume of security events. A *security intelligence and analytics platform* is the infrastructure, including hardware, software, and services, directly supporting the automation of threat management.

Many organizations do not have a centralized, integrated security intelligence and analytics platform. For example, an organization might have a security information and event management (SIEM) solution for log aggregation and analysis, and a separate incident tracking system, with no direct connection between the two.

Organizations should strive to reduce their MTTD and MTTR by implementing a single, unified security intelligence and analytics platform. This can best be accomplished by adopting

a SIEM solution that offers fully integrated, highly mature incident and threat management capabilities.

The Role of Threat Intelligence

You may be wondering what the difference is between security intelligence and threat intelligence. In this book, we use *security intelligence* to refer to intelligence collected by the organization itself on threats targeting its systems, and *threat intelligence* to refer to intelligence collected by a third party on threats in general.

Threat intelligence plays an incredibly important role in the detection component of threat management. For example, it provides vital information for use in security analytics; this information can help with identifying and prioritizing suspicious activities. It also assists security administrators in qualifying security intelligence by providing insights into the history of particular IP addresses, domain names, etc.

This same information can also be invaluable for the response component of threat management. Security administrators can use threat intelligence to learn more about the nature of a threat they're investigating.

Organizations increasingly use third-party threat intelligence feeds to improve their threat management capabilities, as well as other aspects of their security. For example, one of the most common uses of threat intelligence feeds is to improve the detection and prioritization accuracy of SIEM technologies.



Whether threat intelligence comes into the organization through a SIEM or another route, it's important that it be linked through automated means to the organization's security intelligence and analytics platform. Linkage allows the threat intelligence to be fully integrated with other threat-related information to give organizations better insights into the nature of suspicious activities involving their systems and networks.

Chapter 3

Collecting and Processing Forensic Data

In this chapter

- Understand the need to collect data from a wide variety of sources and the role each data source plays
- Get a glimpse of the behind-the-scenes processes that transform raw data into standardized data and metadata ready for analysis
- Learn important considerations for long-term data archiving

Now that we've completed overviews of the attack lifecycle and threat management, it's time to dig deeper and see how to perform threat management. This chapter addresses the collection and processing of forensic data, which is the first part of the detection component of threat management. The goal of forensic data collection and processing is to establish a centralized source of standardized data for security analytics.

This chapter examines four actions involving forensic data: generation, transfer, normalization, and archiving.



You might remember from Chapter 2, “Understanding Threat Management,” that forensic data is largely composed of security logs. Generation, transfer, normalization, and archiving are all concepts taken from security log management. However, don't think that threat management is only about log management; it's much, much more.

Data Generation

Threat management is used to detect threats targeting systems and networks throughout the organization, so it requires enterprise-wide collection of data regarding security events. Existing enterprise security controls and endpoint software (operating systems and applications) are probably already performing most of this data collection.



Do your homework before reconfiguring security event logging capabilities on enterprise security controls, endpoint software, or other technologies. A single change could greatly increase logging, overwhelming local logs and even the centralized log management infrastructure.

In addition to logs from enterprise security controls and endpoint software, threat management often requires two other types of data: network flow data and asset data. Let's examine all these types of data to better understand what they are and how they complement each other.

Enterprise security controls

The most important source of data for threat management is enterprise security controls. These are the network- and host-based technologies that enforce the organization's security policies, remediate vulnerabilities, and detect and block individual attacks.

There are dozens of categories of enterprise security controls, but here are some of the most important ones for threat management:

- ✓ Vulnerability remediation, such as vulnerability management and patch management software
- ✓ Attack detection, including antivirus software and intrusion prevention systems
- ✓ Network technologies, such as firewalls, virtual private networking, and remote access solutions
- ✓ Identity and access management technologies
- ✓ Data loss prevention (DLP) and other exfiltration detection solutions

Endpoint software

Server and client endpoints have direct knowledge of many security events that occur internally or between the endpoint and another system. While many of these events may already be logged to some extent by enterprise security controls, an endpoint's own operating system and applications may perform more detailed logging and may see events that the enterprise security controls don't. This supplemental information provides additional insights into the nature of individual security events on endpoints.

Endpoint software is also important for providing context for security events. Endpoint software often has the greatest visibility into the endpoint's configuration and use, so endpoint logs may provide context for better understanding the significance of a particular event.

Supplemental forensic information for endpoints may also be recorded by independent endpoint monitoring mechanisms. Examples of this information include process state changes, changes to file integrity, and the creation, use, and termination of network connections.

Network flow data

Network flow data is information collected through network monitoring on the flows of data across networks, including the start and stop time for each flow, the volume of data transmitted in each direction, and the basic nature of that data (i.e., which application protocol it uses). Some organizations already log their network flows, but these logs aren't necessarily being analyzed around the clock.



The analysis of network flow data is an important part of threat management because it can indicate both policy violations and significant deviations from typical patterns. This traffic may come from compromised systems, major data exfiltration attempts, and other serious problems that might otherwise go undetected because of traffic encryption or other ways of avoiding content analysis.

In addition to network flow data, other types of data may need to be collected to supplement the organization's existing security logging capabilities. See Chapter 5, "Qualifying Security

Intelligence,” for more information on supplementing and tuning logging sources.

Asset data

A final important source of data for threat management is information on the organization’s IT assets. For example, knowing the role of each system (server, client, network infrastructure, etc.) may help in prioritizing response efforts. Other asset data that may be useful for each system includes the installed software, the primary user or administrator, and the relative importance of the system.



Some organizations may have more advanced information available regarding their systems, such as an up-to-date list of the current unpatched vulnerabilities on each system. This information can provide valuable insights about whether an attack will succeed or fail.

Data Transfer

The data being generated about ongoing security events needs to be transferred to a central location for threat management purposes. This isn’t a simple matter of replicating all log data. Issues include keeping bandwidth usage at reasonable levels while transferring all necessary information, and ensuring the confidentiality, integrity, and availability of the log data while it’s being transferred.

Organizations handle these issues by establishing a log management infrastructure and automating its operation. Most log infrastructures are based on a SIEM solution. SIEMs harvest log data from various sources, securely transfer that data to a central location, standardize the data to use a consistent and clear format, and then analyze it to identify suspicious events and patterns. SIEMs provide minimization and protection functions to aid with data transfers.

Minimization

Log minimization is the process of removing any unneeded information from a copy of log data to shrink its total size. Techniques for log minimization include the following:

- ✓ *Event aggregation*, which involves replacing a number of related log entries with a single new entry
- ✓ *Reduction*, which is intentional and automatic dropping of unnecessary events or event fields
- ✓ *Compression*, which is storage of the entries from a log in an alternate format that uses less space without any loss of information

DON'T FORGET

Never edit the original logs, only copies of them. In most cases, because of permissions, you won't be able to edit original logs anyway, but even if you can, you shouldn't. Editing a log file compromises its integrity and severely limits its usefulness as forensic evidence.

SIEMs typically use multiple log minimization techniques to help reduce bandwidth usage for log data transfers. These methods provide the additional benefit of reducing storage needs for the SIEM itself.

Protection

Another important function of a SIEM is protection for log data transfers. There are three aspects to this protection:

- ✓ **Confidentiality.** Log data often contains highly sensitive information that must be inaccessible to any unauthorized personnel monitoring network communications.
- ✓ **Integrity.** Attackers would love to modify log data to conceal their nefarious activities.
- ✓ **Availability.** Log data must not be lost; for example, if a network interruption occurs, log data transfers must resume shortly after the interruption ends without losing any of the data awaiting transfer.

TECH TALK

SIEMs protect data transfers through endpoint authentication, to ensure that the source and destination system identities are legitimate; with encryption technologies, to safeguard confidentiality and detect any degradation of integrity; and with reliability methods, to ensure that all

legitimate data being sent by the source is eventually received and processed by the destination (the SIEM).

The Forens-O-Matic

So far in this chapter, we've focused on using forensic data for threat management by having SIEMs centralize and normalize the data. While threat management is incredibly important, forensic data and SIEM solutions can play key roles in helping organizations meet other objectives through security automation.

The prime example is security compliance reporting. Most organizations are subject to one or more laws, regulations, or industry standards for securing their systems, networks, data, and/or other IT resources. Examples include the Health Insurance Portability and Accountability Act

(HIPAA), the Payment Card Industry Data Security Standard (PCI DSS), and the International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) 27001:2013 standard on information security management.

Most security compliance initiatives require organizations to perform extensive security auditing and to report on the results of these audits. SIEMs typically provide built-in support for conducting audits and generating audit reports for all major security compliance initiatives. This can be a big time saver for organizations that have compliance needs.

Data Normalization

Data normalization is a complex process that converts log data from its original format to a descriptive standardized format to facilitate its use with search and machine analytics. Normalization takes away the overhead and errors involved in attempting to interpret thousands of log formats. Normalization also allows SIEMs in different organizations to use the same machine analytics rules, regardless of their IT environments.

It's easier to understand data normalization by looking at examples, so let's consider a few of the ways that a SIEM performs normalization.

Extraction of key data fields

Logs come in many formats, from comma-delimited text files to proprietary binary formats. The first step in normalization is to parse each log, identify the significance of each data field, and extract the values from the data fields that the SIEM needs.

For example, the fourth field in a particular log format may represent a source IP address and the fifth field a timestamp for when an event occurred. Because these are important for threat management, the SIEM would pull the values from the fourth and fifth fields.

Standardization of values

After the SIEM extracts values from logs, these values need to be standardized so they're represented consistently. Suppose that one log format stores source IP addresses as text fields and another stores them as hexadecimal values. Many other representations are also possible.

To facilitate searching, correlation, and other analytics functions, the SIEM converts all these representations to a standard one. This conversion is usually based on the SIEM's built-in knowledge of common log formats, but SIEMs can also be customized to handle proprietary formats.

Timestamp standardization

Timestamps are a special case of value standardization. The SIEM converts them to a single standard representation, just as it does with other values extracted from data fields, but the SIEM also usually performs additional normalization on timestamps to ensure they're accurate and consistent. Examples are ensuring that all timestamps represent the time of day using the same time zone, and correcting for any known inaccuracies in a source's clock. Timestamps must reflect the actual time when each event occurred so they can be put in the proper sequence and analytics performed based on those sequences.



Not all SIEMs perform data normalization well. Inaccuracies in standardized timestamps is a common problem, as is omitting normalization rules for popular log sources and making it

difficult for administrators to create normalization rules for custom log sources. When evaluating SIEM solutions, it's best to ask vendors for detailed information on their data normalization methods.

Event classification

Perhaps the most important aspect of data normalization is event classification. A SIEM can greatly enrich the value of log data by determining what type of event is represented by each log entry or group of entries. This allows the SIEM to better understand the significance of each event and what impact an event or sequence of events may have.

On a host level, an event classification could be anything from a failed authentication attempt or a successful escalation of privileges to a security configuration setting change or a security patch installation and associated reboot. Network-level event classifications could include the establishment of a remote access session or the presence of malware within application communications.

Data Archiving

Data archiving is the process of moving data from the SIEM's primary centralized storage to secondary storage, such as a storage area network (SAN). Unlike data generation, transfer, and normalization, which all usually occur before security analytics, data archiving happens after security analytics have been applied.

Data archiving is necessary for long-term data retention because primary centralized storage is limited and costly when compared to secondary storage. The alternative is data destruction, where old data is purged from the SIEM. Data destruction is typically unacceptable unless the data is so old that it's considered no longer of any value to the organization.

The SIEM should manage all data archiving functions so that data stays linked to the SIEM. For example, the SIEM can transfer the data to any suitable location, but this data must be easily recoverable by the SIEM in case it's needed for long-term analytics. The SIEM may also need to ensure the confidentiality and integrity of the archived data.

Chapter 4

Automating Discovery through Security Analytics

In this chapter

- Learn about the features that security intelligence and analytics platforms provide to aid in performing search analytics
- Understand the processes underlying machine analytics and how they're implemented for threat management

In Chapter 2, “Understanding Threat Management,” we introduced the concept of security analytics as techniques for finding the security events that, if left unchecked, could cause material damage to an organization. The organization applies these techniques to the normalized forensic data, generating security intelligence.



TIP

Because of the need to analyze a large volume of security events continuously, security analytics should be automated as much as possible by heavily relying on machine analytics. In addition, organizations should provide automated tools to assist people in performing search analytics.

This chapter takes a closer look at both search and machine analytics, focusing on the role automation plays in improving their efficiency and accuracy. This chapter also briefly covers generation of alerts, which can communicate and prioritize security intelligence for threat management purposes.

Search Analytics

Although they're performed by people, search analytics can be greatly expedited and improved through the use of automated

tools. Time is of the essence when a person needs to do search analytics, and there are several tool-based capabilities that can help, including leveraging dashboards and using drill-down features, search capabilities, and visualization techniques. Let's look at each of these.

Leverage dashboards

In the context of threat management and security intelligence and analytics, a *dashboard* is a SIEM tool that brings together several security analytics views on one screen. Instead of having to manually run several reports and flip among their results, a person can use a dashboard that automatically runs and refreshes a set of reports, displaying the results in or near real time in a convenient layout. Figure 4-1 shows an example of a security intelligence and analytics dashboard.



Figure 4-1: Sample security intelligence and analytics dashboard.

What makes dashboards so useful for search analytics is the various graphical ways in which they present data. This makes it easy for people to find the information that's important to them, such as anomalies that fall outside the normal patterns for the organization.



Dashboards are highly customizable, not only for the organization as a whole but also for each person performing search analytics. Each person can set up a dashboard to show the

security analytics views that are of greatest value to him or her, then save the dashboard configuration so it's used from then on to display the chosen views.

Drill down for details

Dashboards and other security analytics views are useful in and of themselves, but they're even more valuable when they support drill-down capabilities. These allow a person to click on an element of interest, such as the first bar in a bar graph, and obtain more details about that element.

The details provided through drill down depend on the level of the element. For example, a high-level element might offer drill down to another security analytics view with a narrower scope, while a low-level element might display the header and content of a packet, metadata for a network connection, or a wide variety of information for a particular event.



Drill down allows rapid analysis of a trend or event of interest without having to run additional reports or manually search for more information. And after viewing the details, the person can quickly return to the original view to continue search analytics.

Use search capabilities

It should be no surprise that SIEMs offer a variety of search capabilities for search analytics purposes. For example, a person may want to see all recent activity involving a particular IP address, protocol, website, or other component of network traffic. SIEM capabilities make all that possible through a single search.

SIEM searches can be quite complex, allowing a person to retrieve events meeting a detailed set of criteria. Figure 4-2 shows an example of a SIEM interface for search analytics being populated with criteria. In this instance, the analyst is looking for successful access to certain sensitive information (which includes the terms “confidential” or “forecast”) by a user who's on a list of employees to be terminated next week.




Figure 4-2: Sample interface for search analytics.

Use visualization techniques

All SIEMs support the most basic visualization techniques, graphs, and charts, such as those shown in Figure 4-1. What differentiates SIEMs from each other is their degree of built-in support for more advanced visualization techniques, especially those that not only show the data in a graphically sophisticated way, but also enable interactive manipulation of that graphical representation of the data.

For example, Figure 4-3 shows a dashboard that illustrates network traffic for a typical small business. This dashboard leverages four types of charts showing traffic between systems. Each chart uses a different graphics style to visualize complex network traffic and give the analyst the best possible chance of detecting potential threats or operational anomalies. All four charts are linked so that setting a filter on any chart applies the same filter to all other charts. This capability makes it easier to separate signal from noise and find the interesting network traffic.




Figure 4-3: Sample of network traffic visualization for search analytics.



Ideally, visualization tools should offer drill-down capabilities similar to dashboards so that a person can quickly get more information on an event or pattern of events that visualization techniques indicate is probably anomalous.

Machine Analytics

Machine analytics are the heart of a SIEM. Most security events are analyzed only by machine analytics, because a lack of time restricts in-depth search analytics to a tiny percentage of events, although a larger percentage may be analyzed superficially.



If machine analytics don't detect malicious activity, it's extremely unlikely that a person is going to happen upon that activity through search analytics because search analytics are so labor intensive. So machine analytics must be as accurate and thorough as possible.

Let's take a closer look at the three elements of machine analytics: establishing baselines, detecting threats, and prioritizing threats. Note that these aren't performed sequentially; each is an ongoing process happening at all times.

Establish baselines

To achieve the desired accuracy, machine analytics use a combination of techniques for detecting threats. These techniques complement each other and are intended to collectively cover a wide range of threats. See the "Detect threats" section below for more information.

Some of these techniques rely on establishing and maintaining baselines of normal activity within the monitored environment so that deviations from the norm can be easily detected. For example, a SIEM might study network flows over time to generate a baseline that it can use to find deviations. Each deviation could indicate a malicious activity such as installation of unauthorized servers, participation in botnets, or exfiltration of sensitive data.

**TIP**

Establishing a baseline for normal security events may take a while because different activities occur at different times and on different days, weeks, even months. For example, an organization may perform its preventive IT maintenance on Saturdays. Another example is a retailer that has many more online visits from customers during certain times of the year. An organization may need to rely on a subset of the machine analytics techniques until full baselines are ready to support the other techniques.

Detect threats

Detecting threats is where all the effort put into generating, transferring, and normalizing forensic data pays off. As already mentioned, a SIEM uses a combination of techniques to improve its detection accuracy and speed. The details of these techniques are proprietary, but their main commonalities are discussed below.

Deviations from baselines

The previous “Establish baselines” section provided background on baselines. Baselines are constructed over time by observing normal activity, and they must be maintained over time as well to take into account changes in normal activity. For example, an organization may deploy a new service, causing a significant change in network traffic flows.

**CAUTION**

Attackers are familiar with techniques for detecting deviations from baselines. Some attackers avoid detection by slowly increasing their activities over an extended period so that the SIEM gradually changes its definition of normal to include the attacker’s activity. This is a good example of why detection through baseline deviations may be unreliable.

Suspicious patterns

SIEMs can detect patterns of suspicious activity using multiple techniques. The simplest is to look for a signature that matches known bad activity, such as a sequence of bytes from a particular instance of malware or an attempt to use a default password to log into a system.

SIEMs offer much more sophisticated pattern-based detection methods. For example, a SIEM may analyze an entire application session to determine if the commands within that session are in the proper order and if the replies indicate no unusual activity.

Threat intelligence matches

The last part of Chapter 2, “Understanding Threat Management,” defined threat intelligence and highlighted the valuable role that it plays throughout threat management processes. This role is particularly important when it comes to threat detection.



Nearly all SIEMs support the use of threat intelligence feeds from the SIEM vendor or third parties. These feeds include information on the major characteristics of recent and current threats from all over the world, such as the IP addresses used by systems performing attacks and the URLs used in phishing attempts. Organizations are increasingly finding threat intelligence feeds to be a must-have for their SIEMs.

The SIEM compares the characteristics detailed in a threat intelligence feed to the organization’s aggregated log data to find matches. A match doesn’t always mean that the activity is malicious; for example, many systems may share a single IP address. Likewise, a match to a phishing URL could simply indicate an accidental misspelling in an email message. But a match, at a minimum, indicates that the activity is more likely to be malicious and may merit further investigation.

Correlation

Correlation refers to identifying relationships among security events to bring related information together. For example, information about a single security event may be logged as several separate events by multiple enterprise security controls and endpoint operating systems and applications.

Correlation is one of the most powerful capabilities of a SIEM because it can put together the pieces of a puzzle scattered throughout the enterprise.

Another useful role for correlation is identifying advanced threats. As described in Chapter 1, these threats follow the cyberattack lifecycle, reaching their ultimate target after performing a series of compromises across the enterprise. A SIEM with advanced correlation capabilities may be able to link these compromises together, showing the attacker's progress and indicating what was compromised and how and when each compromise occurred. This functionality is invaluable for stopping and mitigating advanced threats.

Prioritize threats

Although detecting threats is incredibly important, prioritizing them may be just as important, if not more so. Organizations simply do not have the resources to manually act against every detected threat, nor do they need to. Many threats are automatically stopped by other enterprise security controls, for example, so they're basically noise to the SIEM and should generally be considered low priority.



Traditional enterprise security controls are much less likely to stop advanced threats, so these should be the focus of prioritization. Each organization should define its own criteria for threat prioritization and have the SIEM implement them to the extent possible. Below are some of the criteria commonly used for threat prioritization.

Likelihood of success

Generally, the more likely a threat is to succeed, the higher it should be prioritized. Determining the likelihood of success isn't easy, however. In practice, the best measures may be how long the attacker has been inside the perimeter, how far inside the organization's perimeter the attacker has penetrated, and whether the attacker has acquired administrative privileges on valuable hosts.

Potential impact

Sometimes it's easy to estimate the potential impact of a threat. For example, a threat may be repeatedly trying to compromise a database server that contains highly sensitive data;

in that case, it's quite likely that the attacker is attempting to perform a breach of that data.

In other cases, estimating the potential impact is nearly impossible. For example, if an attack is caught in its early stages, the attacker may still be far from the target. Most organizations have multiple targets, including personnel records, customer data, and proprietary information (e.g., trade secrets, strategic plans). In these cases, it may be best to focus on other criteria.

Threat reputation/history

Another potential criterion for prioritization is the current reputation or recent history for a particular threat. Threat intelligence feeds often carry all of this information. Suppose threat intelligence indicates that a particular IP address has been the source of numerous attacks against many organizations. On its own, this information may not characterize the threat. But what if the SIEM showed that this same IP address is being used as the source of a successful remote access session used with an administrator account? This may indicate the presence of an advanced threat with privileges and should be given high priority.



Don't take threat intelligence feeds too literally when prioritizing threats. Just because a threat doesn't appear in the feed doesn't mean that it isn't serious. Any threat, but particularly an advanced threat, may change its characteristics at any time by requesting a different IP address, switching the system it uses to launch its attacks, or otherwise altering its appearance. Therefore, it's best to use threat intelligence to raise priority (for example, because of confirming that an IP address is a known threat) but not to lower priority (for example, because of noting that an IP address is absent from the feed).

Alert! Alert!

The security analytics process ultimately results in the generation of alerts. Each alert indicates the detection of potentially serious activity and assigns it a priority. SIEM dashboards can display the latest alerts to prompt human review. An alert itself is concise, but a SIEM dashboard allows people to drill down through an alert to access all the associated information. See Chapter 5, “Qualifying Security Intelligence,” for more information on human analysis of alerts.

Although alerts are designed to be readable by people, they’re often used to initiate automated responses. For example, suppose

an alert indicates that an administrator account has been compromised and is actively being used by an attacker. The SIEM could be configured to react to such an alert by terminating the attacker’s existing administrative sessions and disabling the administrator account in question to prevent further use. The SIEM can trigger these responses in a fraction of a second, immediately preventing further damage, where a human response would take far longer.

For more information on automated responses to threats, see Chapter 6, “Streamlining Threat Response Processes.”

Chapter 5

Qualifying Security Intelligence

In this chapter

- Understand the importance of alert analysis
- Learn the key facets of doing a risk level assessment
- Review criteria for incident declaration and prioritization

In Chapter 4, “Automating Discovery through Security Analytics,” we looked at security analytics processes and described how they create security intelligence in the form of alerts. The vast majority of this security intelligence is generated automatically by the SIEM. In many cases, a person should review security intelligence to ensure it truly indicates malicious activity and it’s properly prioritized for incident response purposes. This review process is better known as *qualification*.

Alert Analysis

Qualification begins with a security administrator’s review and analysis of the alerts produced by the SIEM. Let’s look at the major elements of alert analysis: evaluating the validity of each alert and improving detection capabilities.

Evaluate each alert’s validity

A security administrator should take reasonable measures to determine if an alert is valid. Sometimes this is fast, such as immediately seeing clear evidence of a major attack succeeding. In other cases, alert analysis is considerably more

involved. A security administrator may need to review supporting data held by the SIEM, and even reach back to the original sources of that data for additional information.



TIP Explaining the art and science of in-depth alert analysis is a whole publication in itself, so we don't discuss it further here. Ask your SIEM vendor for additional resources to help you better understand analysis techniques that are relevant to its product.

No one is to blame

In conjunction with determining if an alert is valid, a security administrator may need to find out whether it's being triggered by an actual security incident or an operational problem. For example, suppose that a server is suffering from a denial of service. This could be caused by an attack, but it could also be the result of sudden interest in a particular product generated by a video that has "gone viral."

Differentiating the former from the latter is important because each case calls for a different response strategy. Stopping a denial of service caused by an attack may require the incident response team to coordinate traffic filtering activities with the Internet service

provider. In the other case, an organization generally will ask its operational teams to increase capacity to accommodate the spike in customer demand.

In some cases, however, immediately determining whether an activity is malicious or benign isn't possible because it requires a deeper investigation than can be performed during security intelligence qualification. In these cases, it's often best to err on the side of caution and treat the activity as malicious, handing it off to an incident response team that can engage operational staff for support as needed. The key is handling serious activities quickly, rather than wasting time trying to figure out the intent behind them.

Improve detection capabilities

The goal of alert validation isn't only to confirm the alert, but also to identify ways to improve threat detection capabilities, especially by reducing *false positives* – instances where benign activity was misclassified as malicious. The main methods of doing this are:

- ✓ **Tune logging sources, especially enterprise security controls.** Examples are reconfiguring an intrusion prevention system (IPS) to stop reporting certain events as attacks or to use a different threshold for declaring activity to be malicious.
- ✓ **Supplement existing logging sources.** If logging capabilities aren't robust enough, they may need to be enhanced. For example, some SIEM solutions offer software that can be installed and configured on endpoints to collect data on a wider range of events and to collect more detailed information on each event.
- ✓ **Tune the SIEM.** The SIEM itself may need an adjustment to take into account the unique characteristics of the environment or to compensate for quirks in logging sources that can't otherwise be addressed. For example, the SIEM could be reconfigured to ignore certain events or assign them a lower priority.



It's important for the incident response team to initiate after-the-fact analysis if an incident went undetected for an extended period of time. This is a likely indicator of *false negatives* by the logging sources, which are instances where the sources incorrectly categorized malicious activity as being benign. False negatives need to be addressed so similar incidents can be detected more rapidly in the future.

Risk Level Assessment

Once an alert has been validated, the security administrator needs to assess the associated threat's relative level of risk to the organization. As discussed in Chapter 4, "Automating Discovery through Security Analytics," SIEMs should perform much of the threat assessment and prioritization, such as determining the likelihood of success, the potential impact, and the reputation or history of the threat.

However, a security administrator may need to adjust the SIEM's assessment and prioritization based on factors not necessarily available to the SIEM. Possible additional factors

for human consideration include the importance of the target and the current attack lifecycle phase.

Importance of the target

It's obvious that if the intended target is particularly important, such as files holding the organization's trade secrets, threats against it should be given a high priority. Some targets may be unexpectedly important because of their context, such as a user endpoint that's assigned to a domain administrator.



The organization should ensure that security analysts assessing the risk from threats have access to accurate and complete information on the relative importance of each system, user account, and other major IT attributes.

Current attack lifecycle phase

If feasible, the security analyst should determine the threat's position in the attack lifecycle. As previously discussed, threats should be stopped as early in the attack lifecycle as possible to minimize damage. The corollary is that the farther the threat has progressed, the more likely that a major breach will occur soon, and the faster a response should be initiated.

Incident Declaration and Prioritization

The last step in qualification is incident declaration and prioritization. At this point, the security analyst has validated the SIEM alert and assessed the risk posed by the associated threat. It's now time to determine if an incident should be declared and what priority it should be assigned.

Not every valid SIEM alert necessitates human involvement. For example, a low-priority alert for unusual activity, with no evidence of a successful compromise, may trigger increased monitoring. Another alert indicating an external system is attempting to infect an internal system may merit temporarily blocking the external system and reporting it to the ISP.

Chapter 6

Streamlining Threat Response Processes

In this chapter

- Understand the importance of automating incident management and threat investigation processes
- Review common threat mitigation techniques and see how they can be applied through automated means

Way back in Chapter 2, “Understanding Threat Management,” we outlined the three processes that comprise threat management: detection, response, and recovery. In Chapters 3 through 5, we examined the components of detection in detail because detection is the most complex part of threat management.

Now it’s time to turn our attention to response. Responding to verified threats is a huge topic, so this chapter focuses on one important aspect that’s relevant for all organizations: expediting threat management by streamlining threat response processes through automation. Automation is particularly helpful for improving three components of response: incident management, threat investigation, and threat mitigation.

DON'T FORGET



There’s a subtle distinction between threat response and incident response. Think of incident response as a subset of threat response. While incident response is focused on handling a particular attack or chain of attacks, threat response does all that, and also addresses the handling of the threat itself. Examples of threat response actions include identifying individuals or groups posing specific threats to the organization, and sharing threat information both within the

organization and (after sanitization) with other organizations to help expedite future detection of these threats.

Incident Management and Threat Investigation

The transition from the detection component to the response component of threat management is indicated by the declaration of an incident, as described at the end of Chapter 5, “Qualifying Security Intelligence.” So the response efforts start with new incidents to be managed.

Incident management is a highly complex undertaking. Imagine the myriad pieces that make up the management of a single incident: all the people with roles to play, all the data and metadata that’s collected and generated, and all the manual and automated actions that have to be taken to investigate and mitigate the incident, as well as recover from it.

Now consider all of the incidents currently being managed throughout the organization. Add to those any recently resolved incidents. All of these incidents need to be tracked to correlate old attacks with new attacks, identify recurrences of threats, discover previously undetected damage, and perform other actions that require access to previous incident data. The sheer amount of information to track and safeguard, as well as the number of people who may need access to that information, often under emergency conditions, are overwhelming.



To maintain effective control over incident management, an organization needs a *case management system*, also known as an *incident management system*. This system provides a secure, centralized home for storing, accessing, and analyzing all information being tracked related to the management of an organization’s incidents.

A case management system provides huge benefits for security staff and other IT personnel who participate in incident investigations and mitigations. For the individuals responsible for incident declaration, such a system enables easy, immediate creation and prioritization of new incidents. A case management system also expedites secure collection and centralized

storage of data related to an incident. This includes controlling and auditing all access to incident data, as well as enabling authorized personnel to readily access the appropriate data and to add notes and other supporting information.

A case management system also provides a single place for incident management oversight, including the following:

- ✓ Review the status of all current incidents to reprioritize response efforts.
- ✓ Identify issues to escalate to management.
- ✓ Ensure that investigations are progressing.
- ✓ Determine that multiple incidents are actually different views of the same larger incident.



Organizations with more mature incident management capabilities may also find case management systems invaluable in helping to generate metrics, such as the mean time to detect (MTTD) and mean time to respond (MTTR) measurements discussed in Chapter 2, “Understanding Threat Management.” Metrics such as these allow an organization to assess its response processes over time and set goals for future improvements.

Case management systems offer dozens of features that can help an organization in many ways. Here are two examples of particularly important features.

Workflow and collaboration facilitation

As already mentioned, handling an incident can involve many people, ranging from security, system, and network analysts to other IT professionals, as well as IT and organizational management and, potentially, human resources (for an internal threat), public relations (for public notification), facilities management (for physical security breaches), and other departments.

A case management system helps to ensure that the right people get the necessary information, such as task assignments and status updates, as quickly as possible, and that incident

workflow is managed effectively and efficiently. For example, it might be necessary to obtain management approval before using certain mitigation techniques. The process of coordinating people and tasks and providing the people with the necessary information is better known as *incident response orchestration*.

Fostering collaboration is key, especially under conditions where every minute counts, because it provides a way to share information about an incident and related response efforts. Collaboration features in a case management system should be fully integrated with workflow features.



An organization's SIEM and case management system functions, including workflow management and incident response orchestration, should be provided by a single security intelligence and analytics platform. An integrated platform provides an optimal solution in terms of efficiency and effectiveness, making a real difference in stopping incidents sooner and avoiding major breaches altogether.

Secure collection of supporting data

Although an organization's SIEM already centralizes secure collection of much of its security event information, additional information is often needed after the incident is declared. For example, an incident handler may use various tools to collect information from a compromised host's hard drive. This information may require further analysis by other incident handlers, and it may also need to be shared with system administrators so they can look for similar changes to other systems.

In addition, the organization may also need to preserve this information as evidence for future use in disciplinary proceedings or legal actions. Further, the organization may need to audit the storage of the evidence and all access to it to ensure that the evidence has not been tampered with.

Without a case management system, an organization will be hard pressed to meet its incident handling needs. Just transferring information from an affected system to a centralized location may be a challenge because of reliance on *ad hoc* methods, including plaintext email messages and physical movement of removable media. A robust case management

system helps automate and secure information collection, storage, protection, and auditing processes.

Threat Mitigation

Threat mitigation is an important part of threat response processes. In some cases, every second counts when it comes to stopping an active attack that's damaging the organization and putting valuable assets at risk.

Common mitigation techniques

There are many techniques for mitigation, and for any given incident there may be several options. Here are some examples of commonly used mitigation techniques:



- ✓ Terminate malicious network connections.
- ✓ Reconfigure network-based security controls, such as next-generation firewalls or intrusion prevention systems, to block all network connections with particular attributes (such as a source IP address associated with malicious activity).
- ✓ Disable all user accounts that the threat is utilizing.
- ✓ Kill unauthorized or compromised processes running on hosts.
- ✓ Disable or block access to a vulnerable service.
- ✓ Quarantine a targeted host (such as on a remediation virtual local area network (VLAN)) or disable its network access altogether.
- ✓ Remotely wipe a lost or stolen laptop or mobile device.

Each technique differs in its effectiveness and its potential impact to the operational environment. Each organization must set its own standards regarding the technique or techniques to use under particular circumstances.

Automated mitigation

Mitigation used to be a solely manual process, with the incident response team asking security, system, and network administrators to perform necessary actions. Today SIEMs offer robust, automated capabilities that greatly speed mitigation.

Automated mitigation can be customized based on the characteristics of each situation. For example, a SIEM may automatically initiate mitigation actions when a particularly valuable host is under attack. On the other hand, if the risk of disrupting a server's operations is too high, the case management system might require management approval before tasking an administrator to take the server offline.

Sub-Zero Group Keeps Its Cool

Sub-Zero Group, Inc., an appliance manufacturing company based in Wisconsin, needed to improve the efficiency of its security monitoring and incident response-related processes. As the company grew, it became more difficult for its administrators to keep up with monitoring the security events occurring on its numerous systems and networks. Monitoring individual security logs took far too much time, and accessing the information needed to investigate a suspicious event or generate reports for management caused unacceptable delays.

To streamline its threat response processes, Sub-Zero Group looked for a security solution offering the following:

- Centralized storage for and access to all security log data
- Strong capabilities for correlating security events across logs to bring the pieces of individual events and series together

- An easy-to-use yet powerful interface for administrators conducting searches, investigating potential incidents, and preparing reports

After reviewing the top 10 SIEM products from the latest Gartner Magic Quadrant Report, Sub-Zero Group's team selected LogRhythm and implemented it in the company's enterprise environment.

The LogRhythm solution has made a major positive impact on Sub-Zero Group by making security operations much more efficient. Administrators can now act quickly to identify and stop threats, thus preventing damage, and can use their time and expertise more effectively by focusing on strategic projects.

For more information on Sub-Zero Group and LogRhythm, visit <https://logrhythm.com/resources/sub-zero/>.

Chapter 7

Selecting the Right Solution

In this chapter

- Understand important technical considerations to include in your security intelligence and analytics platform evaluation
- Learn what other operational attributes to look for during solution evaluation

Selecting the right security intelligence and analytics platform for your organization is vitally important. It can make the difference between detecting and stopping a threat early in the attack lifecycle and finding out about a major data breach after the fact.

This chapter describes 10 criteria that should be included in any security intelligence and analytics platform evaluation. Several of the criteria are fairly technical, while others relate to how smoothly the solution will operate and how well it will integrate and interoperate with other enterprise security controls.



TIP This chapter always refers to security intelligence and analytics platforms, not SIEMs. While today's best security intelligence and analytics platforms are all SIEMs, many of today's SIEMs aren't full-fledged security intelligence and analytics platforms, so these terms aren't interchangeable. When you read "security intelligence and analytics platform" in this chapter, think of the subset of SIEMs that have broad and robust threat management capabilities.

Usability

Evaluating the usability of a security intelligence and analytics platform is an important step that's often overlooked. At least two major aspects of usability should be evaluated:

- ✓ **The learning curve.** The platform's interface should be intuitive and easy for users in all roles to learn. Options should include hands-on training and detailed, reader-friendly documentation.
- ✓ **Day-to-day use.** Everyday use of the platform should involve minimal overhead; for example, users should be able to extensively customize their interface with the platform to automatically conform to their preferences. It's also important to consider each way in which people will use the platform, such as for search analytics and incident management.

Scalability and Flexibility

Consider both existing needs and likely future needs when selecting a security intelligence and analytics platform. Products are available in several forms, including hardware appliances, virtual appliances, server-based software, and cloud-based services.

CAUTION



Cloud-based services offer the greatest scalability and flexibility, but they also involve significant latency if large volumes of log data have to be transferred from the enterprise facilities to the cloud provider. Many organizations also have serious concerns about storing their sensitive log information in a public cloud.

Of the non-cloud-based solutions, some offer all-in-one products, while others support distributed architectures – for example, having one system dedicated to log collection and processing, while another system does all the security analytics functions. Smaller organizations with low scalability and flexibility needs are likely to find an all-in-one product easier to deploy and manage. For larger organizations, which typically place more weight on scalability and flexibility, it's generally best to err on the side of caution and go with a more modular distributed solution.

Logging Source Support

No security intelligence and analytics platform has built-in support for every possible source of security-related data. For example, many organizations have custom-written applications that use proprietary log formats. Utilizing these logs will require either customizing the security intelligence and analytics platform to understand the log formats, or modifying the custom-written applications to use a different log format. In many cases, the latter isn't feasible, so organizations should ensure that they can easily configure a prospective security intelligence and analytics platform to accommodate unusual log sources.

Every security intelligence and analytics platform should have built-in support for logs from all major enterprise security controls, operating systems (with the exception of mobile device operating systems, which rarely support security logging), and applications with significant security logging capabilities (e.g., databases). This support will make solution deployment faster, and in many cases it'll also provide superior log processing, in part because it's been vetted by many other organizations.



Other important aspects of a security intelligence and analytics platform's logging source support should also be evaluated, including how accurately the product extracts and normalizes the necessary information from log fields, and how well it utilizes the fields from each logging source. For example, does the product simply report that an error code 301 was observed, or does it have information regarding the significance of that error code?

Supplemental Forensic Data Collection

We've already made several references to the need to collect additional forensic data. This often necessitates changing the logging configuration on the log source itself, but sometimes a log source simply doesn't have the necessary features. Perhaps it fails to log enough details or to log some important events altogether.

To address this concern, some security intelligence and analytics platforms can perform their own security monitoring and logging on behalf of operating systems, applications, and other log sources with insufficient capabilities. This may supplement existing logging or replace it. For example, a security intelligence and analytics platform could monitor networks to capture forensic data on connections and traffic over those connections that wouldn't otherwise be available.

TECH TALK

Some security intelligence and analytics platforms also provide built-in support for the use of forensic data collected by an organization's honeypots. A *honeypot* is a specialized device that exists solely to attract threats and monitor their actions. Honeypots provide an effective way of detecting threats, but more importantly, they give an organization an opportunity to observe a threat in action and study its behavior.

Complying with compliance

Chapter 3, "Collecting and Processing Forensic Data," introduced the concept that security intelligence and analytics platforms are helpful not just for threat management purposes, but also for security compliance. A security intelligence and analytics platform can bring together security event data from throughout the enterprise, so an organization only has to produce a single report for each compliance initiative instead of a separate report for each relevant system. Obviously, this could be a huge timesaver.

When performing a security intelligence and analytics platform evaluation, determine whether the product already supports reporting for all the laws, regulations, and security frameworks that currently apply to your organization. It's also important to make sure that it offers robust report customization capabilities to meet possible future needs. Your organization could someday be subject to a regulation that the security intelligence and analytics platform vendor can't support.

Machine Analytics

As Chapter 4, "Automating Discovery through Security Analytics," explained, machine analytics comprise the vast majority of security analytics workloads because they're highly automated. However, even though they're critically important to threat detection, they're also very difficult to evaluate. That's because every security intelligence and analytics

platform needs time to establish baselines and needs tuning to take into account the unique and unusual characteristics of the environment.



If at all possible, conduct hands-on testing of machine analytics capabilities in your environment before acquiring a security intelligence and analytics platform. Although this takes additional time, it'll result in a better decision and will also speed the official production deployment process.

Search Analytics

The dashboard is the core of the search analytics interface. Any evaluation of a security intelligence and analytics platform should include hands-on testing of the dashboard by some of the people who'll actually be using it in production. The dashboard should provide all the necessary up-to-date information on demand, making a security administrator's life easier.

As discussed in Chapter 4, "Automating Discovery through Security Analytics," the dashboard should provide robust drill-down capabilities and search functions for security administrators. Visualization techniques are also quite useful for representing events in different ways and highlighting unusual activity.

Threat Intelligence Service Choices

There's no doubt that threat intelligence improves threat detection. It's become an absolute must to utilize a threat intelligence service with any security intelligence and analytics platform.

Unfortunately, most security intelligence and analytics platforms restrict organizations in their choice of threat intelligence services. In fact, it's common practice for a platform to support only one particular service. This can lock in an organization to a threat intelligence service that may be less than optimal for its needs. To avoid this situation, organizations should favor security intelligence and analytics platforms that offer a choice of more than one threat intelligence service.

Some platforms enable the use of nearly any commercial or open source threat intelligence service feed. This flexibility allows an organization to change or add threat intelligence

service usage on demand, which is an incredibly valuable capability that could be a game changer at any time.

TECH TALK

Some organizations may benefit from using two or more threat intelligence feeds simultaneously. This approach should provide a somewhat more comprehensive snapshot of current threats, especially if different feeds don't provide exactly the same types of metadata for each threat. Multiple feeds also help with search analytics because security administrators can see if the information on a particular threat is consistent from one feed to another. Consistency increases confidence in the nature of the threat.

Automated Investigation and Mitigation Capabilities

A security intelligence and analytics platform with highly automated capabilities that support investigation and mitigation enables an organization to respond more quickly to threats and attacks. For example, a platform may be able to automatically link the pieces of a complex attack to each other, showing the security analyst the threat's path through the organization and providing a list of all the systems and user accounts it's compromised.

Chapter 6, "Streamlining Threat Response Processes," listed several examples of common mitigation techniques. All of these can be automated through one mechanism or another. For example, a security intelligence and analytics platform might have built-in support for interactions with other enterprise security controls to initiate certain mitigations. Other mitigations could be performed through automated means by having the platform launch a custom script written by the organization's security or system administrators.

TIP

Organizations needing automated mitigation capabilities should carefully study how security intelligence and analytics platforms could interact most effectively with their existing security infrastructure to stop attacks in progress.

Customization

A recurring topic in this book is the need to customize a security intelligence and analytics platform. This isn't a shortcoming of the technology; rather, it's a reflection that

every environment is unique in its IT usage, security requirements, and risk profile. Customization is a positive thing, and organizations should look for security intelligence and analytics platforms that offer extensive customization capabilities. Such capabilities should be available for the following, at a minimum:

- ✓ Dashboards: both the layout of each user's dashboard and the parameters used to generate each element
- ✓ Alerting and prioritization: such as setting thresholds for alert generation, disabling generation of unnecessary alerts, and specifying how to weigh various factors when setting alert priorities
- ✓ Mitigation: including defining what circumstances should trigger automatic mitigation actions and tailoring mitigation actions to take advantage of other enterprise security controls

Reporting customization is also important if an organization wants to create its own reports for internal use or for external auditors, for example.



Don't fall into the trap of perceiving extensive security intelligence and analytics platform customization capabilities as being more important than the platform's out-of-the-box capabilities. Some platforms rely on the organizations adopting them to devote countless hours to customization before they're able to be of value. All evaluations should factor in key out-of-the-box capabilities, such as built-in rules for log normalization, analysis, and correlation; default searches, queries, and reports supporting common scenarios; and scripts for initiating automated mitigation actions.

Technical Support

A final factor to consider during evaluations is the quality of the platform vendor's technical support. An organization that conducts hands-on testing of prospective solutions has a perfect pretext to contact technical support at each vendor and see how quickly and accurately reps respond to typical support questions. Vendors should also have extensive, well-maintained online knowledge bases to expedite troubleshooting.

In addition, it's important to know how technical support will function during emergencies. Questions to ask vendors include the following:

- ✓ Is support available at all times, even on holidays?
- ✓ What is the guaranteed maximum response time for technical support inquiries?
- ✓ If the platform is hardware-based, how quickly will hardware be replaced if there's a serious defect or malfunction during the warranty period?

Real-world evaluation considerations

LogRhythm's website includes case studies describing how a wide variety of organizations have selected, implemented, and used LogRhythm solutions. The case study on Redcats USA showcases a retailer with nearly 10,000 employees and over 400 physical locations. Its IT infrastructure included a wide variety of devices, operating systems, and applications that generated security event log entries. Redcats needed to ensure that all these logs were monitored, collected, and analyzed on an ongoing basis for incident detection and PCI compliance.

To achieve this, Redcats evaluated security intelligence and analytics platforms based on several criteria, the most notable of which were usability and automated investigation and mitigation capabilities. Redcats recognized that these characteristics would make it easier for employees to do their jobs and save them a great deal of time, as well as ensure that incidents would be detected and stopped much more quickly to prevent damage. The company also looked for solutions offering robust scalability/flexibility, logging source support, and machine analytics.

Redcats chose LogRhythm to meet its particular needs and quickly reaped benefits. Its staff became much more efficient in identifying and addressing security issues and operational problems. Administrators became more proactive, anticipating potential failures and acting to prevent them instead of reacting after the fact. Redcats was able to accomplish far more with its existing IT staff than was possible before, thanks to LogRhythm.

Visit <https://logrhythm.com/resources/redcats/> to read the full Redcats USA case study. Additional case studies illustrating real-world considerations for product evaluations include:

- ALPS, <https://logrhythm.com/resources/alps-funds-services/>
- Fortis Bank, <https://logrhythm.com/resources/fortis-bank/>
- Phoenix Suns, <https://logrhythm.com/resources/phoenix-suns/>
- Ventura, <https://logrhythm.com/resources/ventura/>

Chapter 8

Steps for Successful Implementation

In this chapter

- Understand the preparatory and planning actions that are crucial to success
- Get insights into what's involved in designing a security intelligence and analytics platform architecture
- Learn tips and tricks for smoothly implementing and integrating a security intelligence and analytics platform
- Review maintenance actions that must be performed on an ongoing basis to keep the platform effective and efficient

There's a lot of hard work involved in a successful security intelligence and analytics platform implementation, but in the end it's worth it. It's hard to quantify the value of its benefits – for example, you'll never know what damage would have resulted from incidents that never happened – but better threat management has positive effects across the enterprise in reducing incidents and achieving compliance with external security requirements.

CAUTION



Threat management is much more than selecting a best-of-breed SIEM platform. It's also about establishing policies, making sure that systems log all the necessary security event information, and implementing sound log management practices throughout the enterprise. Without all these other pieces in place, threat management will have much less of a positive impact on the organization.

Let's break the 10 steps of the security intelligence and analytics platform implementation process into four phases and examine each phase in more detail:

- ✓ Preparation and planning
- ✓ Solution design
- ✓ Production implementation
- ✓ Maintenance

Preparation and Planning

In the preparation and planning phase, the organization is performing actions that lay the foundation for designing and implementing a security intelligence and analytics platform.



It's only natural to want to jump in and start trying out products, but this may be premature. Greater emphasis on preparation and planning up front should help expedite and smooth the rest of the phases.

Step 1: Define goals and requirements for threat management

The first step is to define the organization's goals and requirements for threat management. Examples of common goals include the following:

- ✓ Improve the accuracy and/or speed (MTTD) of threat detection
- ✓ Decrease threat response time (MTTR)
- ✓ Expedite enterprise recovery processes
- ✓ Document evidence of compliance with external security requirements

Once there's consensus about the goals and the relative importance of each one, it's time to define the corresponding requirements. Goals are high-level statements of what matters to the organization, while requirements are concrete statements supporting one or more of the goals.

An example of a requirement is that the security intelligence and analytics platform must support the collection and processing of logs from all enterprise security controls, as well as desktop, server, and network operating systems. Another example is that throughout the enterprise, all logs containing security event information need to be retained for a minimum of 30 days.

An organization might also have a wish list of desirable but not absolutely necessary characteristics for its threat management capability.

Step 2: Create and validate policies supporting the requirements

After defining the requirements, the organization should alter existing policies and create new ones as needed to support threat management. For example, the hypothetical requirement to retain all security-related logs for at least 30 days would likely be added to an existing enterprise data retention policy.

Many organizations find it valuable to create a policy specifically for threat management, to encompass security log management. Such a policy should include the following:

- ✓ Definitions of all threat management-related roles and responsibilities throughout the enterprise
- ✓ Operational, security, and privacy requirements for each phase of the security event log management lifecycle, including log generation, collection, transfer, normalization, archiving, and destruction
- ✓ Requirements or guidelines for how often logs should be analyzed through machine or manual means
- ✓ Lists of which general mitigation techniques are preferred, permitted, and prohibited
- ✓ Requirements and/or guidelines for the actions to be taken under various circumstances, such as when it's okay to use fully automatic mitigation and when management approval is required before mitigation

CAUTION

It's critically important to ensure that policies are updated to support threat management, and that they're carefully vetted and validated by affected personnel throughout the organization. A new policy statement that sounds perfectly reasonable on the surface might be impractical or even impossible for others to comply with at this time because of budget shortfalls, technical limitations, and other reasons.

Step 3: Prioritize the implementation of threat management

Threat management is complex, involving many technologies, people, and processes. Implementation isn't going to happen overnight. Some organizations may take years to acquire and implement all the components, particularly if there's a need to stagger technology purchases and address other major security goals and requirements at the same time. Other organizations recognize an immediate need for improving threat management and will implement a security intelligence and analytics platform much more quickly.

CAUTION

Perhaps the most important part of planning threat management implementation is changing the common assumption that security controls will stop all threats. People need to understand that some threats will succeed no matter what security controls are in place. You need to show all levels of the organization that there's a serious problem that's getting worse, and that threat management is the best weapon against it. Make threat management a priority and people should be more supportive of it.

Solution Design

After the preparation and planning phase is complete, solution design begins. In this phase, the organization designs the threat management architecture, then evaluates products and services needed to establish that architecture. At the end of this phase, the organization should know exactly what hardware, software, services, etc. it will be acquiring and deploying in support of threat management.

Step 4: Design a threat management architecture

Organizations should design a complete threat management architecture, and not just an architecture for a security intelligence and analytics platform. Suppose that an organization plans on selecting a leading SIEM product that includes a full-fledged case management system. It's still overwhelmingly likely that other threat management components will be needed, ranging from log rotation utilities for system administrators to advanced search analytics add-ons for the identification of particular types of threats. All such components should be included in the architecture design.



In terms of the security intelligence and analytics platform itself, important characteristics to consider include the following:

- ✓ Which solution form or combination of forms is best (hardware appliance, virtual appliance, server-based software, or cloud-based service)
- ✓ Whether an all-in-one or distributed approach is best
- ✓ How well the platform can handle expected peak logging volume, and how easily the platform can be expanded as volume increases
- ✓ What degree of fault tolerance and redundancy is necessary for the platform
- ✓ Whether log collection should be agent based and/or agentless
- ✓ How data retention and archiving should be handled

Step 5: Evaluate products and services for the architecture

We've already dedicated Chapter 7, "Selecting the Right Solution," to explaining a set of 10 criteria for evaluating security intelligence and analytics platforms. And as mentioned in step 4, a threat management architecture may include various

additional components for system administrators, security analysts, and other staff involved in threat management.

We still need to highlight one critical part of the threat management architecture: the intake of one or more threat intelligence feeds. As the “Threat Intelligence Service Choices” section in Chapter 7 explored, some threat management platforms provide considerable flexibility as to which threat intelligence feeds each organization can use.



If there's a choice to be made, then by all means the organization should carefully evaluate each candidate threat intelligence service to see which feed or combination of feeds would best meet its needs. And if a platform provides no choice of feeds, the organization should still evaluate the feed it supports to ensure that it's of sufficiently high quality to meet its needs. If it isn't, consider that a showstopper and evaluate other platforms instead.

Production Implementation

During the production implementation phase, the components of the architecture are acquired, deployed, integrated, configured, and customized. Other important elements of this phase include developing processes related to the solution and training administrators and others to use the solution.

Testing, testing, 1...2...3...

Any organization that's serious about improving its threat management performance should consider setting up a permanent test environment. During solution design, a test environment is obviously beneficial because the organization can bring in trials of prospective products and evaluate them hands-on. However, a test environment has several beneficial uses even after a solution is selected and acquired, including the following:

- Testing updates and upgrades before deploying them in production
- Developing and testing customizations before duplicating them in the production environment (especially important for automated mitigations)
- Training new administrators, analysts, and other users on the solution itself and on threat management principles and techniques in general

CAUTION

The length of this phase may vary greatly based on the strength and maturity of the organization's existing security technologies components and processes. The timing is dependent on many environment-specific factors, but one thing is true for nearly every organization: implementation is best performed gradually. Installing a security intelligence and analytics platform and trying to get it to consume, analyze, and report on 10,000 new logging sources all at once is almost certainly going to cause chaos. It's much more productive to focus on achieving the most important use cases first, thus getting real value from the platform quickly, and adding the other use cases over time.

The steps in this phase, and potentially other phases as well, shouldn't necessarily be performed in a serial fashion, completing one step before starting the next. Indeed, it may make sense to change their order, especially if parts of the implementation are already in operation.

Step 6: Acquire, deploy, and integrate products and services

Organizations should acquire and deploy the components of the threat management solution following their usual processes for new enterprise security technologies. However, one thing in particular differentiates security intelligence and analytics platforms: the platform shouldn't affect production when it's initially deployed and as threat management tools are first integrated with it. Once it's receiving data from logging sources (see step 7 below), that's another story.

One of the most important parts of threat management implementation is enabling automated mitigation actions. When you're working with a new security intelligence and analytics platform, you should definitely begin integrating it with the enterprise security controls that it can manipulate through mitigation actions.

CAUTION

However, it's extremely unwise to enable automatic mitigation until the platform has been in full production use for a bit. Activity needs to be monitored for days or weeks (even a few months in some environments) to identify false positives and other conditions that could trigger automatic mitigations, causing denials of service to customers and other legitimate users.

Step 7: Gradually transition log sources to the solution

In this step, the organization configures all of its logging sources to supply data to the security intelligence and analytics platform. This step can occur concurrently with step 6, but it's listed separately because it's often still going on long after step 6 has ended.

This step is generally the most labor intensive of all, because it may necessitate changing the configuration of every enterprise security control, operating system, database, and other critical enterprise applications so that they all log the necessary security event information locally and participate in transferring it to the centralized solution.



There are many options for getting log data from the log sources to the security intelligence and analytics platform. For example, some solutions pull data from the sources, while others have the sources push their data to the platform. Most platforms support the use of both agent-based and agentless technologies. The primary disadvantage of agent-based technology is that software has to be installed on each system and granted privileges to access the security logs. On the other hand, this degree of access also enables the agent software to gather additional information that the system's built-in logging capabilities can't record.

As new log sources start feeding logs to the platform, it's likely that additional changes will have to be made to most of these sources. A common example is reducing false positives by changing the source's logging configuration, such as raising thresholds for generating certain log entries or suppressing some log entries altogether. This illustrates why a gradual approach is strongly recommended; lessons learned from integrating one log source can be applied to integrate similar log sources much more smoothly.

Step 8: Develop processes and train staff on the solution

Although some process development and staff training can happen during steps 6 and 7, the majority can't be done effectively until there's a significant volume of security event data

from production being processed and analyzed by the security intelligence and analytics platform. At that point, security administrators can try out all the search analytics tools and techniques on actual data sets.

Processes can be developed, documented, and refined in accordance with how the solution actually behaves in the production environment, such as what volume of logs the platform itself generates on a daily basis.

Step 9: Customize dashboards, mitigations, alerting, etc.

Customization is an important part of solution implementation, and for the sake of efficiency, universal customizations should be performed before individual users implement their own. For example, the organization might develop a few dashboard templates showing the information that's most valuable for meeting its primary requirements. Individual users can select the template that's closest to what they want and, over time, customize their copy.



Other customizations, such as those for mitigations, alerting, and reporting, typically apply to the entire solution and don't provide customization capabilities for individual users. These customizations should be tightly controlled by a small number of authorized administrators using a formal change management process, because they can affect all users of the solution and the organization as a whole. An example is accidentally disabling the wrong alert, which allows serious attacks to go undetected.

Maintenance

The final phase of threat management implementation is maintenance. All sorts of maintenance actions are needed over time, from testing and installing updates to replacing faulty hardware and adding more storage for data archiving.

Instead of rehashing all the typical maintenance duties that apply to any enterprise security control, let's focus on one particular to threat management: fine-tuning the solution.

Step 10: Refine the solution to improve its performance

Within the description of the production implementation phase, there are several references to making adjustments – changing what log sources record, supplementing built-in operating system and application logging capabilities with agent-provided logging, changing alert thresholds and disabling unneeded alerts, etc.



One of the biggest mistakes organizations make time and time again with enterprise security controls, and especially security intelligence and analytics platforms, is underestimating how much time they'll need for ongoing maintenance. Threats and vulnerabilities are changing all the time, but so are technologies. There's always a new application being deployed, a new operating system coming out soon, or a new form of malware being detected by enterprise security controls.

Organizations need to understand that the dynamic nature of IT in general, and security in particular, necessitates making changes to the log sources and the security intelligence and analytics platform itself on an as-needed basis. Ongoing maintenance keeps the threat management solution functioning efficiently and effectively: stopping threats early in the attack lifecycle and preventing data breaches and other damaging events.

Glossary

attacker: A person who performs cyberattacks. Also known as a *threat actor* or a *cyberattacker*.

attribution: The process of determining who's responsible for causing an incident. In other words, attribution is the discovery of the identity of a threat.

case management system: A system that provides a secure, centralized home for storing, accessing, and analyzing all information being tracked related to the management of an organization's incidents. Case management systems also facilitate efficient and effective incident response orchestration. Also known as an *incident management system*.

compromise: The result of a successful attack. A compromise occurs when there's a loss of confidentiality, integrity, and/or availability of data, systems, networks, or other computing resources.

correlation: Identifying relationships among security events to bring related information together.

cyberattack: An attempt to negatively affect the security of computing resources. Also known as an *attack*.

cyberattack lifecycle: The pattern that serious cyberattacks tend to follow for breaching sensitive data. The six phases of the cyberattack lifecycle are reconnaissance; initial compromise; command and control; lateral movement; target attainment; and exfiltration, corruption, and/or disruption. Also known as the *attack lifecycle*.

cyberthreat: An entity (individual, group, nation state, etc.) that plans and executes cyberattacks. Also known as a *threat*.

dashboard: A SIEM interface that brings together several security analytics views on one screen.

data breach: A compromise that causes a loss of data confidentiality.

data normalization: The process of taking log data from its original format and converting it to a descriptive, standardized format to facilitate its use with security analytics.

exfiltration: The process of transferring sensitive information from an authorized location (controlled and protected by the organization) to an unauthorized location outside the organization's control.

false negative: An instance where security controls failed to detect the presence of malicious activity.

false positive: An instance where security controls incorrectly categorized benign activity as malicious.

forensic data: All of the security-related event data being collected by an organization. Forensic data comes from four categories of sources: enterprise security control logs, endpoint software logs, network flow data, and asset data.

honeypot: A specialized device that exists solely to attract attackers and monitor their actions. An organization can use honeypots as a source of security intelligence.

incident: The occurrence of security events of particular concern to an organization. An incident may be declared when an organization detects a successful attack, an attack in progress, or indications of a new, serious threat, such as unusual reconnaissance actions or failed attacks.

incident response: The process of handling a particular attack or chain of attacks. Incident response is a subset of threat response.

incident response orchestration: The process of coordinating people and tasks involved in incident response and providing the people with the necessary information.

indicators of compromise: The signs of a compromise. An organization with knowledge of indicators of compromise can look for the presence of those indicators in security logs, file systems, and other locations to identify additional systems that have likely been compromised.

investigation: The process of security analysts reviewing security intelligence to determine how a potential threat should be handled. Investigation may also look for broader patterns that could indicate a wider compromise in progress. Investigation is a major part of the response component of *threat management*.

lateral movement: The act of repeatedly leveraging a compromise of one internal device to compromise another internal device, so as to move through an organization and reach a target.

log minimization: The process of removing unneeded information from a copy of log data to shrink the total size of the data. Techniques for log minimization include event aggregation, reduction, and compression.

machine analytics: Security analytics performed automatically by a system or systems.

mean time to detect (MTTD): A measure of the average elapsed time from the start of an attack or chain of attacks to the detection of the activity.

mean time to respond (MTTR): A measure of the average elapsed time from the detection of an attack to the completion of all response activities.

qualification: The process of assessing security intelligence to confirm its legitimacy and priority. The purpose of qualification is to verify that the detected activity necessitates a response.

reconnaissance: Research conducted by an attacker to learn more about its target's environment.

search analytics: Security analytics performed by a person.

security analytics: Techniques used on aggregated forensic data to find the events and sequences of events that are of greatest concern from a security perspective.

security information and event management (SIEM): A security control designed to centrally store, normalize, and analyze security log data gathered throughout an enterprise.

Some SIEMs also offer incident and threat management capabilities.

security intelligence: High-quality, actionable information about the most serious threats currently acting against an organization. Security intelligence is collected within an organization on the threats against it.

security intelligence and analytics platform: The infrastructure, including hardware, software, and services, directly supporting an organization's automation of threat management.

target: A system of particular interest to an attacker in achieving a goal, such as breaching certain data.

threat intelligence: Information collected by a third party on threats in general.

threat management: The processes for managing the threats that use the cyberattack lifecycle. Threat management comprises three ongoing processes: detecting threats targeting the organization, responding to detected threats, and recovering from damage caused by threats.

threat mitigation: The process of thwarting a threat by stopping its in-progress attacks. Threat mitigation is a major part of the response component of threat management.

threat response: The process of performing incident response and handling the threat behind the incident.



They will get in. They can be stopped.



We can help. LogRhythm's next-generation security intelligence and analytics platform identifies high-impact threats and neutralizes them before they can result in a material breach. It uniquely unifies SIEM and log management with network and endpoint forensics and advanced security analytics to provide comprehensive threat life cycle management and the ideal foundation for today's cybersecurity operations.

Start monitoring your network for threats today: logrhythm.com/freemium

 **LogRhythm**
The Security Intelligence Company

Learn how security intelligence technologies can thwart today's advanced threats and stop attacks before data breaches and other major damage occur.

Advanced threats are methodically evading enterprise security controls and causing major data breaches. These threats may target any organization, so no one is immune. Fortunately, you can use security intelligence and analytics technologies to greatly improve detection, especially to stop threats early before major damage has been done. If you are involved in threat or attack detection, or if you play a role in incident response, this book is for you.

- **Understanding attacks and threats** — learn the basics of the attack lifecycle and threat management processes
- **Improving detection** — see how a security intelligence and analytics platform can automatically discover attacks and threats, then help prioritize responses
- **Streamlining response processes** — learn how to use a security intelligence and analytics platform for incident management and threat investigation and mitigation
- **Selecting the right solution** — review what to look for when evaluating security intelligence and analytics platforms
- **Deploying a solution** — know tips and tricks for designing and deploying a security intelligence and analytics platform

About the Authors

Karen Scarfone is an accomplished freelance security writer and a recognized technical expert in the field. Read her works at www.scarfonecybersecurity.com.

Steve Piper is a high-tech veteran with over 20 years of experience. Steve has authored more than a dozen books on IT security, networking, and Big Data. Learn more at www.stevepiper.com.



CYBEREDGE
PRESS

Not for resale

ISBN 978-0-9961827-4-4



9 780996 182744 >