

▶▶ BRILYANT

THE
DEFENDER'S
[**ADVANTAGE**]

A GUIDE TO ACTIVATING CYBER DEFENSE

MANDIANT

THE
DEFENDER'S
[ADVANTAGE]

A GUIDE TO ACTIVATING CYBER DEFENSE

MANDIANT

CONTENTS

Foreword	6
Introduction	9
What is Cyber Defense?	12
Intelligence is the Guiding Light	17
What is to be accomplished with the intelligence?	20
Who will consume the intelligence?	20
How will intelligence be communicated?	22
How will feedback be gathered and consumption measured?	26
What information sources are needed?	26
Threat Intelligence Activation	28
Maintaining the Mission with Command and Control	31
Information Transfer	32
Major Incident Management	33
Authority to Act	34
Hunting for Active Threats	37
Goals of Threat Hunting	39
Automated Hunting	43
Hunting Process	45
Detecting and Investigating Malicious Activity	51
Threat Intelligence in the SOC	53
A Nuanced Approach	54
Detection Development	54
The Challenged SOC	56
The Optimized SOC	59
Activating the SOC	60
Using Automated Defense Technologies	62

Responding to Compromise	65
Initial Triage	66
Investigation Lifecycle	72
Investigation Accelerators	76
Incident Remediation	80
Remediation Accelerators	88
Lessons Learned	90
Targeted Testing and Validation of Controls and Operations	95
Understanding the Attack Surface	96
Beyond Breach and Attack Simulation	98
Validating the Effectiveness of Your Staff	102
Validating the Design of Incident Response and Remediation Plans	102
Activating Cyber Defense	105
Stakeholder Buy-in	105
Staffing Considerations	106
Leveraging Accelerators	106
Engaging Managed Services	108
Flexible Consumption Models	108
Conclusion	109
Appendix A	111
Multifaceted Extortion	111
Appendix B	115
Investigative Theory	115
Appendix C	119
SOC Ransomware Procedure	119

FOREWORD

In today's world, cyber threats, physical systems, and geopolitical issues intersect on the battlefield where the cyber security war is being fought. Adversaries are leveraging ransomware and multifaceted extortion campaigns with unprecedented frequency. It can seem daunting, and pervasive attacks do not require sophisticated, coordinated efforts. Organizations often succumb to phishing and password reuse that leads to initial compromise and ultimately the deployment of ransomware.

In my role as Senior Vice President of Mandiant Services in EMEA, I often see the lack of confidence organizations exhibit in their ability to thwart attacks and ready their defenses against these attacks. This is rarely due to a scarcity of tools. It is more likely due to improper deployment of capabilities, lack of properly trained forces with ineffective automation to support them, a deficiency in understanding the threats being faced and poor application of defenses against them.

We must remember that the battle is not being fought on the adversaries' turf. We own the battlefield—the Defender's Advantage. This provides opportunities to do better. We have the capabilities. We need to activate them and bring them to the battle.

To ready the battlefield and prepare our forces to fight the adversaries, we must:

- **Use intelligence to guide all actions within Cyber Defense with centralized command.** Organizations often have an abundance of intelligence coming in but don't understand how to verify its credibility, applicability, or how to action it. Intelligence should provide situational awareness of the cyber threats an organization faces and feed a command and control system to orchestrate each Cyber Defense function. This intelligence-driven approach offers a Cyber Defense battle plan to reduce systemic risk and provide a common front against the evil being faced.

- **Apply intelligence to activity seen in the environment to provide the right information to enable teams to fight.** More data seems great, but when analysts are presented with too much data, they can miss critical events. Intelligence should be applied to events BEFORE they are presented to the analyst to prioritize investigation efforts and reduce the noise that can distract them.
- **Continually assess defenses against active threats and stay nimble as the battlefield changes and our enemies evolve.** A critical failure of Cyber Defense organizations is to quickly set up controls with the intent of circling back to optimize the deployment. As businesses progress and adversaries change tactics, defenses quickly become outdated and ineffective. Continuous validation of the effectiveness of security controls against the latest threat-actor tactics, techniques, and procedures (not just alert signatures) is required to reduce the security deficit.

With proper preparation, I believe we can change the course of the battle. Now is the time to activate our cyber defenses against the adversaries and fight. Gloves off!

S. MCKENZIE

Stuart McKenzie

Senior Vice President of
EMEA Services, Mandiant

INTRODUCTION

Prominent attacks dominate the headlines and have security leaders scrambling for solutions, legislators imposing new cyber security requirements, and businesses demanding answers from their security groups. Ransomware and multifaceted extortion are just some of the threats organizations must defend against. Insider threats and the consolidation of risk in the cloud are also top of mind for security leaders.

The **Defender's Advantage** is the concept that organizations are defending against attacks in their own environment. This provides a fundamental advantage arising from the fact that they have control over the landscape where they will meet their adversaries. **Organizations struggle to capitalize on this advantage.**

Establishing and orchestrating robust cyber defenses help organizations take command and galvanize their defender's advantage. It allows organizations to prepare their environment to identify malicious activity, detect and respond to compromise and validate the effectiveness of controls and operations against active threats. Once established, security organizations must activate their cyber defenses, advancing capabilities from a prepared state to active duty. Threat intelligence guides this activation.

With effective use of threat intelligence, organizations can understand who is targeting them, what threat actors are after and if they can be compromised.

Threat intelligence is leveraged to:

- Trigger hunt activities through the use of information about active advanced persistent threat (APT) groups and the latest relevant attacks to **identify active or past compromise**.
- Prioritize vulnerabilities based on the likelihood and impact of compromise. IT and Security groups use this to **inform patch and upgrade priorities**.
- Inform security engineering teams of the monitoring required to **alert on activities tied to active APT groups**.
- Prompt security operations groups to **refresh playbooks to reflect shifts in adversary tradecraft**.
- Provide context around breaches so that incident responders can **scope, rapidly contain a breach and avoid repeat compromise**.
- Update validation efforts on the latest TTPs to continually assess the controls and operations' ability to **prevent or reduce the impact of an attack**.

The functions of cyber defense, as described in this book, are rarely built and resourced entirely inside an organization. A strategy of task and process automation is key to maintaining consistent quality and amplification of existing expertise. To accelerate achieving Cyber Defense capabilities, organizations leverage strategically selected managed services to provide full Cyber Defense coverage, microservices for targeted needs and expert resources for in-house deployment and operations development. It is also critical to ensure that all capabilities are continuously validated to ensure Cyber Defense performance is meeting expectations.

Capitalizing on The Defender's Advantage is achievable by operationalizing intelligence, applying effective automation, leveraging services to fill in capability gaps and having a rich understanding of the environment (i.e., the battlefield). This book provides information about what Cyber Defense functions make up mature security organizations and the activation of the capabilities within each function.

With multifaceted extortion and ransomware, a successful network intrusion precedes the manual ransomware/malware deployment. Organizations need to shift left and catch the attack in its earliest stages. It's less about catching the attacker's malicious payload, and more about catching the intrusion that precedes the payload deployment.

Steve Ledzian, VP, CTO – APAC, Mandiant




Learn about multifaceted extortion in Appendix A.

WHAT IS CYBER DEFENSE?

Cyber Defense is actively resisting attacks and minimizing the impact of a compromise. It is one of the four domains of Information Security with the other domains being Security Governance, Security Architecture and Security Risk Management. A successful Cyber Defense organization seamlessly integrates with the other information security domains to create a resilient security program.

Figure 1. Four domains of Information Security



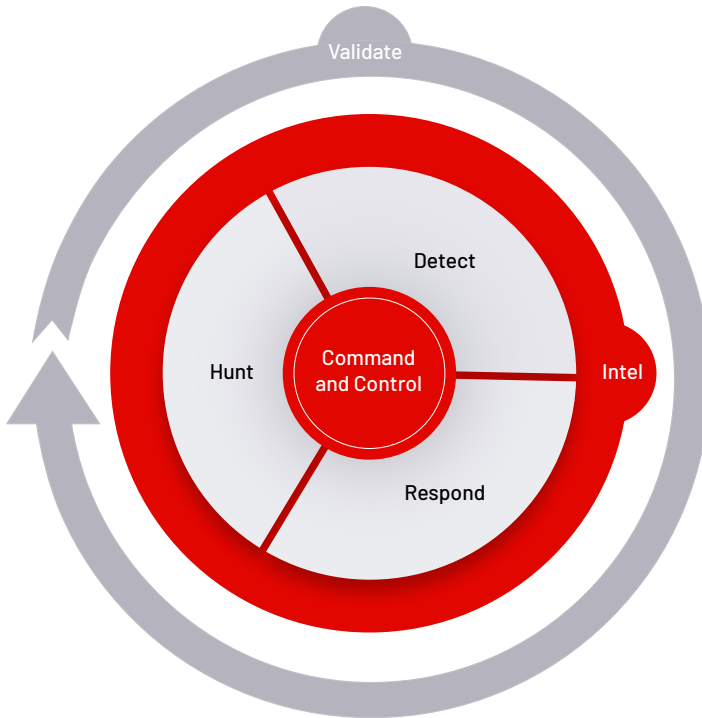
The Cyber Defense domain is made up of six functions to achieve its mission of identifying and responding to threats to the business or organization. **The service provided by a Cyber Defense organization is to allow the organization to continue to operate in the face of threats.** These functions work together to provide a common front against attackers.

Figure 2. Functions of Cyber Defense



Each of these functions focuses on a unique piece of the Cyber Defense mission, and feeds into each other, allowing each function to benefit from the capabilities of the other functions, focused on different goals. The functions of the Cyber Defense domain are Intelligence, Command and Control, Hunt, Detect, Respond and Validate. Each of the functions are associated with different activities, actions, or responsibilities, but they all represent core strengths used collectively to improve cyber defenses.

Figure 3. Functions of Cyber Defense in Action



Intelligence is the lifeblood of Cyber Defense as it directly feeds into every other function. From providing indicators of compromise (IOCs) that can be used to develop use cases within the Detect function, providing guidance to build mission-driven Hunt activities, or developing adversary emulation to test security controls within the Validate function, threat intelligence is critical to every element of the ecosystem.

The **Hunt** function provides much of the proactive capabilities of Cyber Defense to identify active threats within the environment. It includes advanced capabilities such as insider threat identification, deception tactics, and threat modeling exercises.

Resources (time, people, and money) are always constrained. Using intel to focus your efforts to the most critical areas helps make the best of those precious resources.

Alex Wood, Head of Enterprise Security,
The Anschutz Corporation

Many of the traditional elements seen in security operations are performed in the **Detect** function. The function also includes enhancing contextualization, providing detection analytics, increasing visibility to give organizations a clearer picture of threats to the environment and providing a more comprehensive view of the environment itself.

The **Respond** function focuses on capabilities such as response and containment, and includes automation and orchestration, which drive faster remediation of incidents to minimize impact.

Ensuring that security controls are providing value is part of the **Validate** function. The Validate function provides assurance that the security control ecosystem is operating as designed throughout changes to the environment, as well as identifying any vulnerabilities in the environment. This function is not limited to technical controls, but also validates that response procedures remain effective to the changing threat landscape.

The final piece is the **Command and Control** function, which manages the Cyber Defense functions to ensure they are operating in an effective manner to accomplish their mission. This function is focused on Cyber Defense program management and establishes formal processes for resources management, communications, metrics, and crisis management. This program management ensures that the Cyber Defense capabilities remain resilient to changes within the organization and threat landscape.





INTELLIGENCE IS THE GUIDING LIGHT

Organizations subscribe to an average of

7.5

Threat Intelligence feeds*

66.5%

still disseminate CTI through Email, PPT, Spreadsheets, Documents**

Only

43%

has documented CTI requirements**

Intelligence is a differentiator for Cyber Defense and guides actions within the Cyber Defense organization. Most security organizations subscribe to threat intelligence feeds, but they struggle to operationalize the intelligence and use it to protect the business. They focus on quantity of IOCs provided over the quality of IOCs. Furthermore, organizations fail to create plans for utilizing threat intelligence and often dive straight into intelligence gathering, which can lead to wasted time going after the wrong sources, dealing with information overload, or acquiring information that is not actually needed. While there is no method to fully automate every intelligence need, there are ample opportunities to leverage strong actionable intelligence more efficiently with key targeted automation.

*Forrester Wave ETIS Q1, 2021

**SANS CTI Survey 2021

Cyber Threat Intelligence (CTI), when implemented effectively, helps prioritize response actions, supports strategic risk-management decisions, and guides other Cyber Defense actions. CTI is more than a feed of IOCs or common vulnerabilities and exposures (CVEs), which may or may not contain surrounding context.

CTI is the summation of observations from multiple data sources, enriched by context from commercial data designed to support strategic, operational, and tactical decisions. It creates a tailored picture of an organization's exposure based on known exploitable technology and whether the adversary groups leveraging these IOCs or CVEs have historically targeted the organization or industry peers. This data builds the foundations of a cyber threat profile.

The intelligence provided inside of the threat profile should be at a level to guide security monitoring operations, provide visibility and awareness of cyber threats, and prioritize detection of those threats within the enterprise.



Creating actionable intelligence starts with creating a threat intelligence plan. A basic threat intelligence use case plan needs to answer the following questions:

1. What is to be accomplished?
2. Who will consume the intelligence output?
3. How will the intelligence output be communicated?
4. How will feedback on the output be gathered and consumption measured?
5. What additional intelligence sources are needed to fill existing gaps?

The Cyber Threat Profile is arguably the most important document for a cyber intelligence program. And most programs either don't have one or aren't using it to drive their operations.

Andrew Close, Principal Consultant, Intelligence Capability Development, Mandiant

What is to be accomplished with the intelligence?

The goals of intelligence should be to:

- Gain insight into attackers and how attacks are being executed.
- Determine attackers' motivations and targets.
- Find vulnerabilities that exist in the environment.
- Explore the likelihood the organization will be targeted by a threat given their profile (relevance).
- Ascertain the impact to the organization should a compromise occur.

Each one of these intelligence components may have a specific value to the organization and dictate the best format and means of consumption by the intel recipient.

Who will consume the intelligence?

Organizations can purchase dozens of the "best" intelligence subscriptions; however, even the best intelligence is wasted if it is not utilized properly. It is important to understand who will consume the intelligence and what they will do with it.

An incident response (IR) team may be most interested in IOCs, attacker infrastructure and typical attacker tools and techniques which can act as investigation accelerators. Easily identifiable attacker behavior can expedite the investigation, allowing incident responders to focus on areas that are most important based upon attacker motivations. This type of raw intel can be programmatically leveraged to highlight potential signals in the noise, giving incident responders pivot points during their investigation.

Even within the Cyber Defense program, different teams may require different types of information. The threat hunter, for example, may want the same raw intel as responders during a hunt. In addition, they might also require finished intelligence and the cyber threat profile to aid in the development of a hypothesis for hunt missions and new detection rules for security information event management (SIEM).

Threat detection might benefit most from well-curated feeds to contextualize the alerts and prioritize the team's actions related to alerts.

	Tactical Level	Operational Level	Strategic Level
Security Roles	Security Operations Center Network Operations Center Vulnerability and Patch Management Team	Incident Response Team Forensics Team Red Team/Pen Testing	Chief Information Security Officer Risk Management Security Management
Tasks	Indicators to security tools Patch systems Monitor, triage, and escalate alerts	Determine attack vectors Remediate Hunt for breaches Emulate adversaries	Allocate resources Communicate with executives
Problems	False positives Difficult to prioritize patches Alert overload	Event reconstruction is tedious Difficult to identify damage	No clear investment priorities Executives are not technical
Value of CTI	Validate and prioritize indicators Prioritize patches Prioritize alerts	Add context to reconstruction Focus in on potential targets	Demystify threats Prioritize based on business risk

Table 1: Value of CTI per Security Role

How will intelligence be communicated?

Preferences vary among organizational stakeholders as to how they prefer to consume intelligence. Strategic consumers will often prefer higher-level reports that contextualize how changes in adversary operations are impacting the broader industry threat landscape, whereas more tactical consumers will often prefer more technically dense reports that contain indicators they can immediately use to scan the enterprise. A threat intelligence plan should document how and when to communicate for each consumer type.

Humans aren't great at evaluating risk. We tend to catastrophize, focusing on the worst-case scenario, and then start to view this worst case as the most likely case—even if there is no evidence or reason for us to do so. Once we're aware of threats, we come to see them as more likely to happen. And this is where cyber threat intelligence can play a big role in keeping our evaluations of cyber risk honest.

Mark Owens, Principal, Intelligence Capability Development, Mandiant

- **Security Engineering** will likely require detailed information about the vulnerabilities being exploited—down to the affected software version. This can be communicated via email or (preferably) through a ticketing system so both the vulnerability and the prevention status can be tracked. Automating the contextualization of external vulnerability intelligence with attack surface and organizational vulnerability details allows for more rapid and meaningful prioritization of patching, remediation and monitoring.



- **Incident Responders** will want a system to quickly process logs and telemetry data for quick intel evaluation of attacker infrastructure, file signatures, IP addresses and other IOCs. This can be automated with custom scripts and tools, making it easy and fast to get initial triage done. This can even help find early indicators based on high-confidence IOCs to determine potential adversary attribution. This can greatly accelerate response and containment activities.
- **SecOps** will confirm the signatures that identify the malicious behavior is active or updated in the controls, such as endpoint agents, next generation firewalls (NGFWs), or other defense systems. This requires specific information about the threats.
- **Hunt teams** will want bigger picture information such as how a specific attack plays into known APTs and uncategorized (UNCs) group clusters, so they can hunt for all raw indicators and other activity from the attack group. This is typically more of a “pull” of threat briefing or profile information from a threat intelligence platform. Threat hunters can sweep for lower fidelity indicators during a known APT or UNC attack. These types of indicators would not be deployed to security technologies for day-to-day operations due to the noise. However, during an attack, the hunt team can correlate many low fidelity indicators to individual endpoints which can help illuminate attacks.
- **Executives and Board Members** will likely seek short updates focused on business impact, not on technical specifics.

I use CTI risk ratings to create vulnerability situation reports. The report provides CISOs information about critical vulnerabilities, so they have immediate and actionable intelligence, often before I am asked for it. The automated report correlates data between our vulnerability vendor and our CTI vendor to show overall company exposure, the vulnerability severity rating, the CTI risk rating, the exposure broken out for each line of business or subsidiary and a write up on how the vulnerability works and key links for more information.

Gibby McCaleb, Director of Security Operations at Sony Pictures Entertainment

How will feedback be gathered and consumption measured?

An intelligence feedback loop provides organizational stakeholders with an opportunity to convey whether the content, structure and message delivered in an intelligence report was useful; whether it sparked follow-up questions; or if the product missed the mark. They can then offer suggestions for refinement in subsequent products. Barring feedback, intelligence producers will continue to write products they believe are helpful to specific audiences. While metrics describing how certain intelligence products help drive cyber security decisions within an organization, intelligence producers grade themselves on whether they have been able to satisfy existing requirements, and feedback is key to informing this process.

What information sources are needed?

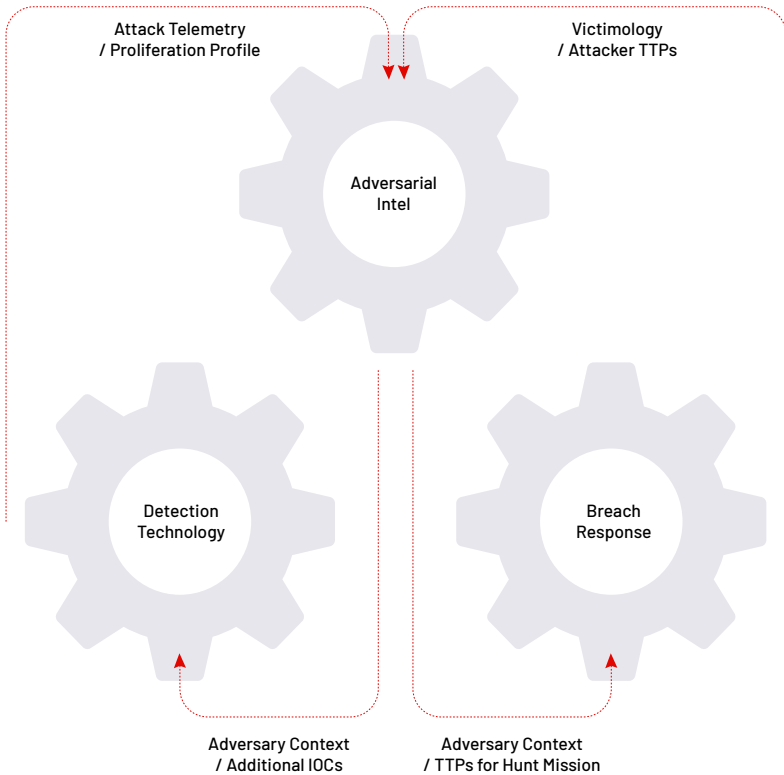
Information sources are the final consideration after the previous four questions are sorted, as these sources identify what information is needed.

Common information sources include:

- Threat intelligence platforms (TIPs): A central repository of analytic assessments made about incidents and threat actor capabilities.
- Open-source intelligence (OSINT): Data and information available to the general public for free. This includes almost everything found on the internet that is not behind a paid firewall. OSINTs include a vast amount of raw data but also requires a lot more digging for high quality, relevant data.
- Information sharing communities (ISACs/ISAOs): Industry-tailored intelligence that has usually been vetted for accuracy.

A common standard for reliability and credibility is the NATO-devised Admiralty Code which ranks reliability from 'A' through 'F'. 'A' means a source is extremely reliable. 'F' means the exact opposite.

Figure 4. Flow of Information to Refine CTI



Diversity in sources is necessary because they all hold a different piece of the puzzle to be solved. However, all sources are not equal. Both the reliability and credibility of sources must be considered. Reliability and credibility include fidelity, currency, visibility and verification. Maintaining diverse and extensive intelligence collections, and continually expanding and curating these is an ideal use for strategic partnerships, thus limiting the requirements individually to just those that are closely held and directly sourced intelligence collections.

Threat Intelligence Activation

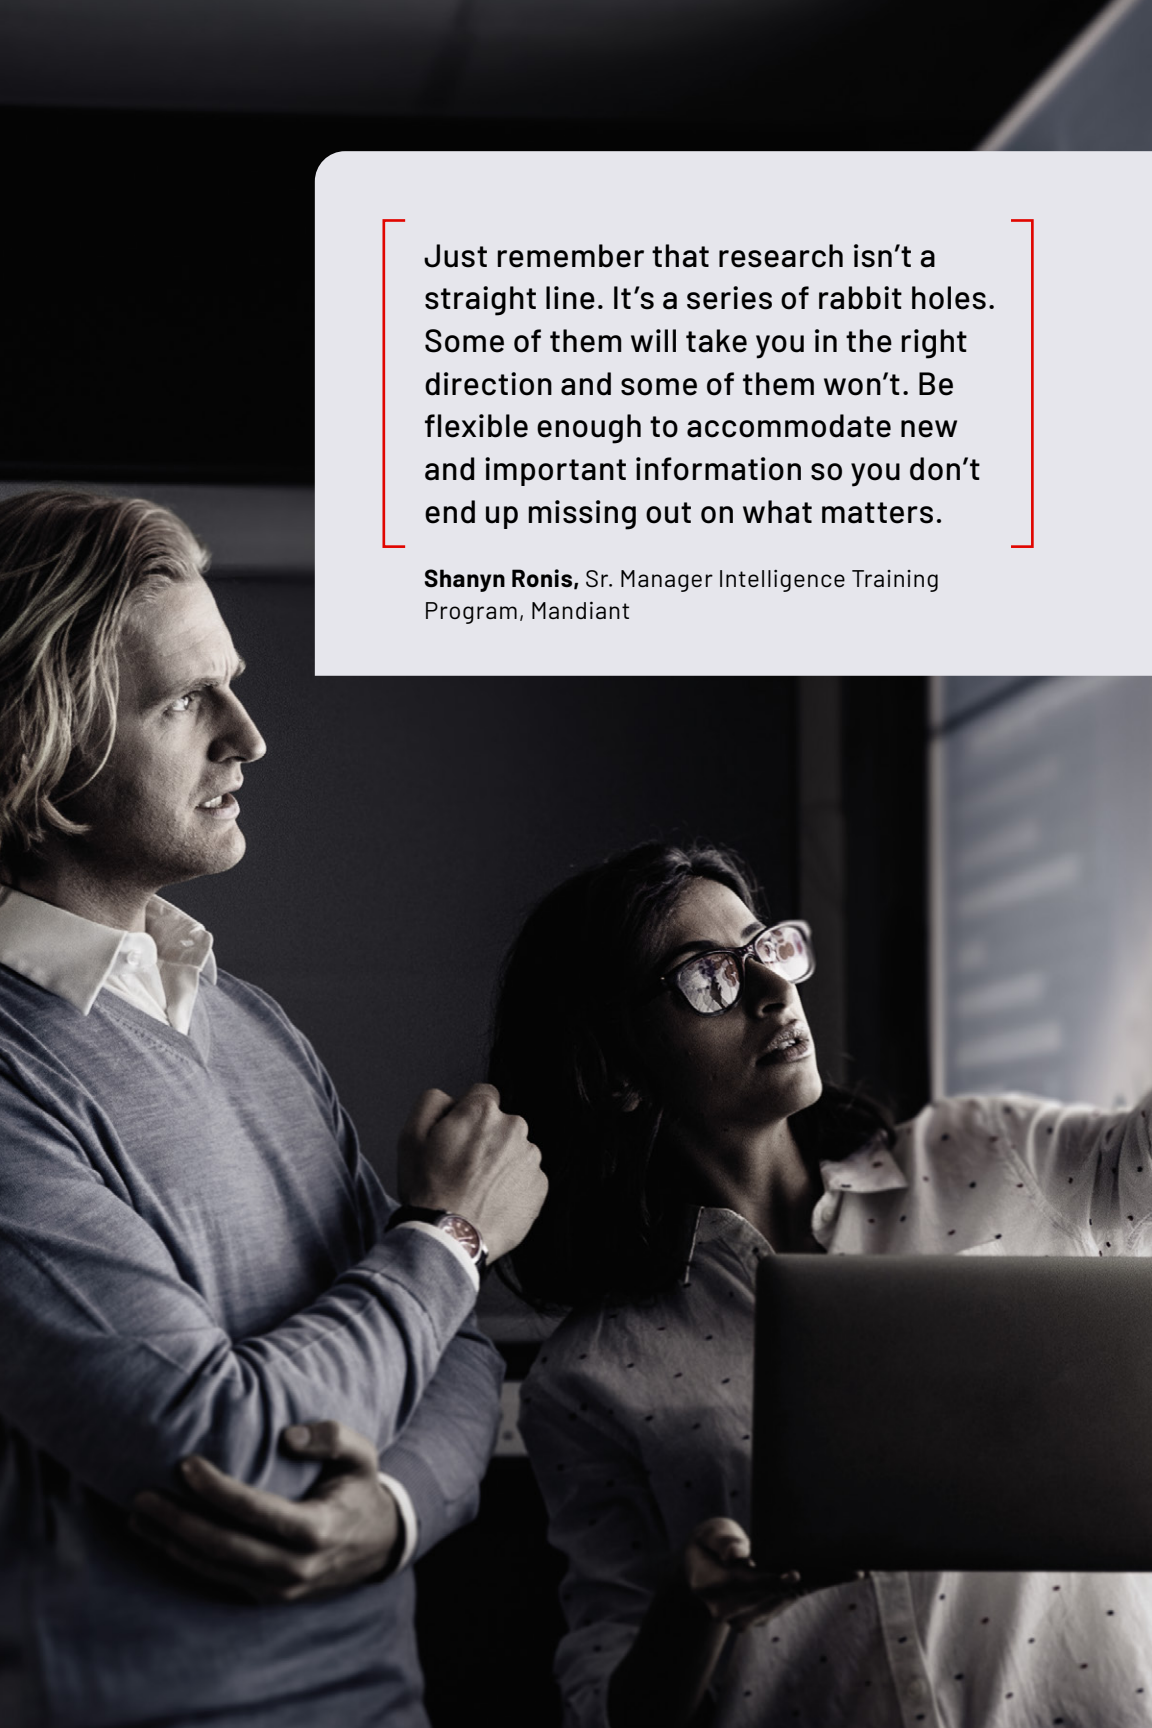
Activating threat intelligence triggers the action of all other functions within a Cyber Defense organization.

To initiate threat intelligence, take the following steps:

- 1. Create a threat intelligence plan**
 - Document the different consumers of threat intelligence in the organization and their requirements.
 - Compare requirements against collection sources and document any gaps.
 - List what information will be communicated to each consumer and how it will be used.
 - Identify the format the intelligence will be delivered in and how often it will be shared.
 - Socialize and broker buy-in from stakeholders on the plan and determine how feedback on utilization of the intelligence will be collected.
- 2. Map out the threat landscape and cyber threat profile**
 - Identify threat actors in the current cyber threat landscape and their TTPs.
 - Identify adversaries likely to target the organization based on high-value targets or crown jewels and their methodologies.
 - Determine the likelihood of compromise by identified threats and the impact on the organization if the attacks are successfully executed.
- 3. Prioritize vulnerabilities**
 - Identify vulnerabilities actively or likely to be exploited by relevant threat groups against implemented technologies.
 - Help the information technology groups properly order their patching and upgrading efforts.
 - Indicate where security controls should be added or updated to detect malicious behavior.


These actions will be continuously evolving as the threat landscape changes and business needs evolve. The goal is to inform decisions to reduce cyber risk.




A man with long blonde hair and a woman with dark hair and glasses are looking at a laptop screen in a dark office setting. The man is on the left, wearing a grey sweater over a white collared shirt. The woman is on the right, wearing a light-colored patterned shirt. The background is dark with some blurred light sources.

Just remember that research isn't a straight line. It's a series of rabbit holes. Some of them will take you in the right direction and some of them won't. Be flexible enough to accommodate new and important information so you don't end up missing out on what matters.

Shanyn Ronis, Sr. Manager Intelligence Training Program, Mandiant





MAINTAINING THE MISSION WITH COMMAND AND CONTROL

While the other functions of the Cyber Defense domain establish capabilities to identify, mitigate, and respond to threats, the Command and Control function keeps these capabilities aligned to the Cyber Defense mission. This management function is referred to as Command and Control because it establishes authority and direction for the Cyber Defense functions. In collaboration with the larger information security program, the Command and Control function prioritizes the Cyber Defense resources to protect the organization from threats with the highest likelihood of impact for the organization based on intelligence, as well as prioritizing resources to protect the crown jewels of the organization.

The Command and Control function focuses heavily on people and processes within the Cyber Defense domain. Mature process maintenance, proper resource training and appropriate resource allocation can help ensure that the established Cyber Defense capabilities remain effective. In many cases, organizations build strong detection capabilities and tools, but due to poorly established and documented processes, incidents are not responded to appropriately. This can lead to an increased impact to the organization.

Information Transfer

An important role of the Command and Control function is to facilitate information transfer between the functional components of a Cyber Defense organization. This function does not need to be part of every conversation or email that passes between groups. Instead, it needs to be aware of relevant information or efforts that would impact the security of the organization. This group must maintain a keen awareness of the organization's security posture to identify the impact of business decisions, active threats or organizational changes in real-time.

Example

The Hunt function develops and executes a campaign that uncovers activity matching known IOCs tied to a potential high-risk exploit that has recently been made public within the industry. It is the Command and Control's responsibility to ensure that the Detect function develops detection criteria for use cases to identify and alert on potential activity tied to the IOCs. In addition, it is the responsibility of the group to ensure that the Respond function has the appropriate response procedures documented to react to a compromise as quickly and efficiently as possible. These efforts should be identified and tracked to completion to ensure that the organization is properly prepared for the threat before the environment is impacted.

Command and Control is often an overloaded concept in cyber security between people and systems. It is important to understand it boils down to leadership intent (i.e., prioritized outcomes) integrated with alignment of the team (i.e., who is doing what, when, and where). Done well, the right people, processes and technology are at the right place at the right time to have the right things happen.

Josh O'Sullivan, Chief Technology Officer,
Ardalyst

Major Incident Management

The Command and Control function is responsible for the facilitation of communication during a major incident, as well as the coordination of different groups for investigation and remediation. It also acts as the primary function responsible for identifying and routing resources to focus on attack mitigation, remediation, and post-incident analysis to drive preparation activities to prevent future, similar attacks. Some organizations place this responsibility on the members of a Security Operations Center (SOC) or a computer security incident response team (CSIRT). These stakeholders remain an integral part of these proceedings; however, the management of the incident should be relinquished to the Command and Control team.


The Command and Control team is typically composed of individuals that have multiple roles for which they are responsible. Various stakeholders from other functions of the Cyber Defense organization will likely be involved in the Command and Control function. Members from multiple functions must contribute to this group for it to be effective.

Authority to Act


Another critical element of the Command and Control function is the authority to act. Much like a CSIRT model, this group must have the authority granted by executive leadership to drive communication efforts between teams and act efficiently, unhindered by traditional operational challenges. This group shares additional similarities to a CSIRT in that their mission objective extends to day-to-day operations rather than sole activation during an incident. However, this group must also have the backing of leadership to drive efforts in a quick and efficient manner to keep the other functions communicating and accountable throughout their individual efforts.


A common challenge that Cyber Defense organizations face is that each group is acting independently of one another and lacks coherence. There is no authoritative body that has the oversight into each group that can combine the efforts together into something larger than these disparate efforts together into something greater and more cohesive. The Command and Control function addresses this challenge.

The Command and Control function is critical for efficient operations. A Cyber Defense organization's success is often limited by the unwillingness to build this oversight team with members of the various core functions and to grant them the authority and resources necessary to ensure every team is working towards a unified goal driven by the same intelligence.









HUNTING FOR ACTIVE THREATS

Threat Hunting proactively uses intelligence about an adversary and their operations to search the enterprise for active or previous compromise. An ancillary goal of the threat hunting process is to identify whether any gaps exist in security controls that would otherwise facilitate detection of a compromise.

A key element of Threat Hunting is ascribing to the adversary mindset, thinking through what would be of interest to specific adversary groups, which may vary from what an organization has identified as its crown jewels. Hunt teams also understand business operations processes, system configurations, naming conventions and security controls as designed. Using these insights into how an adversary may be able to compromise the environment and remain undetected becomes a key step in developing an initial hypothesis. Adopting the mindset that assumes an adversary is already present or capable of bypassing one or more security controls opens the organization's eyes to detect key indicators of an attack.

Analysts have devious imaginations. We understand the realm of possible for how threat actors can compromise systems and routinely ask why we haven't observed a certain actor group using certain methods. Threat hunting allows organizations to test these types of hypotheses.

John Doyle, Principal Consultant, Global Government Operations, Mandiant

When developing a hypothesis, there are several ways to guide the development to achieve higher relevancy for organizations. The use of past incident data or red team assessments can provide indications of previous weaknesses or previous critical controls, that had these controls been bypassed, would have allowed more significant exposure for the organization. This form of data has a high fidelity as it has already been demonstrated and observed in the environment.

Further hypothesis development can be guided using internal and external CTI, especially when combined with local knowledge of an organization's environment and critical assets. Hypothesis creation can be aligned to and further refined with threat trending and attacker TTPs. Ideally, these would be mapped to a common framework such as the Mandiant Attack Lifecycle, the cyber kill chain, or MITRE ATT&CK[®] framework.

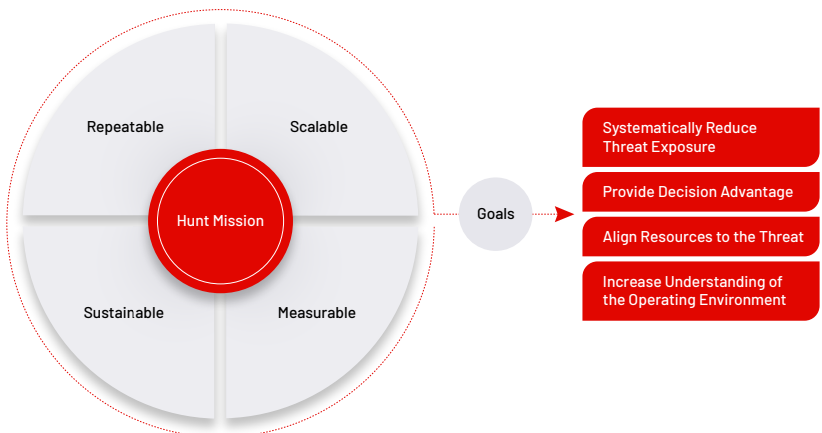
Goals of Threat Hunting

A goal of the Hunt function is to identify potential evidence of compromise and escalate to the Incident Response team.


Other goals include:

- **Systematically reduce threat exposure** through proactive detection, which reduces dwell time and facilitates rapid identification of new tactics, techniques and procedures used by attackers and the creative use of existing data assets to detect malicious activity.
- **Provide decision-making advantage** to bridge the gap between automated, computer-driven detection and human analysis.
- **Align resources to combat current and future threats** through awareness of existing security control gaps.
- **Increase understanding of the operating environment** and baseline knowledge of “good” behavior to make detection of anomalies and “evil” more apparent.

Figure 5. Goals of Threat Hunting



When developing a threat hunting program, ample consideration must be given to the qualified resources available to complete hunt missions. It is equally important to determine the frequency in which threat hunt missions will be executed and what detection methods already exist in a more traditional alerting and detection use case approach.




Considerations when developing mission-based hunts include:

- **Threat scope**
 - What threats should the business focus on while hunting? With hunting, you assess the threat, in part, based on who is targeting you, their intent/objective, their level of sophistication, and the possible impact to your organization if an attacker were to be successful.
 - What are critical controls and systems likely to be subverted or leveraged to accomplish the objective of an attack?
 - What user behaviors are considered anomalous?
 - What is the time frame in which these threat actors typically complete their objective from initial contact/compromise, or from the first opportunity of detection to final mission completion?
 - How many and what are the other detection mechanisms that would be associated with this typical attack sequence? Threat hunting will generally provide less value if the time required to complete the threat hunting exceeds the time for an attacker to complete their objective.

- **Identifying security posture issues**
 - How does hunting help identify issues with your security posture?
 - Is there enough visibility in the environment to conduct hunt missions?
 - Pinpoint an area of the network that has been overlooked and is not being monitored or additional logs which should be collected from existing sources.
 - What are the logging retention practices?
 - Logging may identify the need to hunt back further than the data that has been collected.
- **Executive decision-making**
 - How can hunting enable executive decisions?
Hunting helps to advocate for resources, such as additional monitoring and new tools to best protect the network, personnel, or partnering needs.
- **Intra-team communication**
 - How are hunting results communicated out to other teams?
Hunting results can inform other teams and help prioritize security operations, reduce the time to detect by IR, generate alerts, increase situational awareness, streamline vulnerability management and encourage future intelligence team support. Hunters and pentesters should share information bidirectionally to validate exploitation can occur based upon hunt findings and ensure exploitation hasn't already occurred based upon pentest findings.
 - How do we apply lessons learned to inform future hunts?
A finding from one hunt may generate a hypothesis to trigger a new hunt. For example, looking for commonality and trends on the type of job roles targeted using certain spearphishing lures.
- **Threat Hunting Platform, technology and skill set**
 - As the hunt mission complexity and sophistication increase, the capabilities of a platform threat hunters need will increase. Migration from simple hunts leveraging a well built query of existing data to more data science driven methods and ultimately methods leveraging ML across big data will become more critical as common hunts become more efficient and new hunts are being designed. Additionally, the relevant skills may move from those more closely associated with a security analyst or incident responder to a data scientist or developer.





Automated Hunting

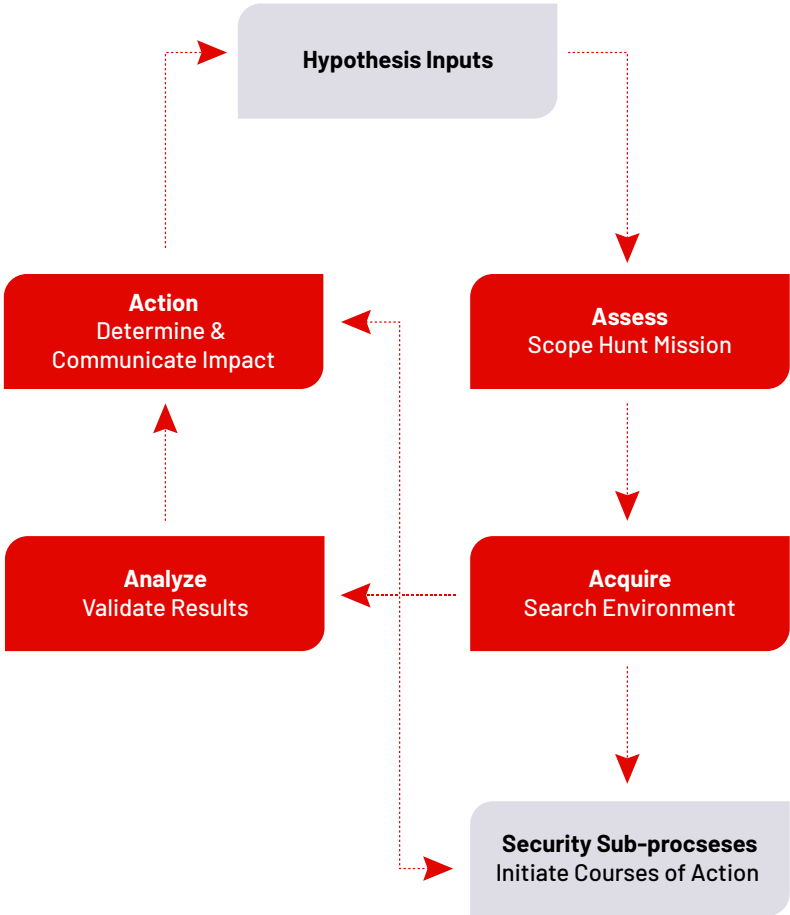
Hunting needs to be scalable with the existing resources or extended with partners—and enabled by technology. Activities should be flexible to adapt to the changing threat landscape and documented to provide consistency and follow a repeatable process. The actions and outcomes of the Hunt function must be measurable to show effectiveness and be used to guide future business decisions.

While threat hunting itself cannot be fully automated, the act of threat hunting begins where automation ends, and it benefits heavily from automation. The collection and data preparation steps being automated can be a key time savings allowing for more hunting in the environment and less time on those steps. **CTI is key to guide Hunt teams by focusing on areas targeted by adversaries, increasing hunt accuracy.**

Threat hunting should not be siloed. Integration leads to enhanced response capabilities, security improvements, enriched internal threat intelligence, and new detection methods.

Numerous threat hunting activities can be aided by strategic partnerships, especially those hunts that are tied closely to well-known methodologies associated with external threat actors. This allows for more frequent hunt mission execution and allows for more internally driven business specific hunts to be conducted by internal resources. It is easier for internal resources to understand how to hunt for insider threats, intellectual property and business rule violation related hunts than for a typical external partner. If that partner can hunt for well-known methodologies and techniques used by the most likely threat actors, then internal resources can focus on those items nearly exclusive to their environment.

Figure 6: Threat Hunting Process



Hunting Process

Successful hunting teams have formalized, repeatable processes that remain flexible to adapt to change. The four pieces of the process are Assess, Acquire, Analyze, and Action. The pieces include a security sub-process to activate upon discovery of a certain triggering action. The security sub-process could include activating an incident response plan, developing new threat detections, performing a security architecture review, or incorporating lessons learned into various Cyber Defense functions.

Assess (Scope the hunt mission)

Attack modeling identifies what you have that the attacker wants and intelligence provides possible motivations. Use these to identify specific adversaries that attack similar organizations and identify their TTPs.

- Develop and propose the hypothesis.
- State the adversary's technique that was likely used to enter or that which persists in the environment.
- Identify forms of targeted data, set the hunt timeframe and potential cost limits. Hunting can become compute intensive which could impact direct costs if compute is cloud based, and indirect costs if some manner of productivity is impacted.
- Determine the data required to carry out the hunt. Understand what type of visibility, collection and search capabilities exist for this data.
- Assess hunt value compared to more automated detection methods. There is variation in cost associated with each designed hunt. Typically, searches against existing data sources for previously overlooked suspicious activity is a lower cost of execution than pulling net new information across the enterprise. The latter of these, however, with the proper expertise, can give visibility to previously undetected or detectable indicators and novel techniques. Generally, the lower fidelity the indicators the more expertise required to find suspicious activity.

Acquire (Search the environment and gather data)

This is the data-gathering phase. From the outputs of the first phase, a purpose for the hunt is defined, and collection methods built to search for how that activity would look in the environment. An atomic IOC is needed for the search, as well as a behavior or pattern and a segment of the environments to search. This yields potential matches.

To search the environment and gather data:

- Identify access requirements and tools
- Initiate data collection and search
- Validate the completion of searches
- Perform initial analysis, inclusive of stacking and frequency analysis
- Escalate high-impact threats by activating the security sub-process

Analyze (Validate results)

Output needs to be validated before continuing, to make sure results match what was expected. The outputs are logical conclusions; the judgements and facts are based on analysis. Assessment of the hunt outcome will serve as input to the ultimate action of what is communicated: recommendations, threat summary, additional CTI product, additional hunts needed, or security content for detection.

To validate results:

- Evaluate target matches
- Correlate, sort, and link data then prioritize
- Pivot to related/new data
- Perform inferential analysis
- Determine attack vectors and TTPs
- Determine control effectiveness
- Identify hunt limitations and constraints





Action (Communicate impact)

The intelligence product and the associated recommendations need to be disseminated to key stakeholders. These communications may be a combination of a high-level strategic overview along with very tactical (e.g., patch management) recommendations such as a business decision around budgeting.

To communicate impact:

- Determine overall impact
- Develop threat summary
- Form strategic outlook
- Identify gaps in process
- Identify data to block or alert on
- Deliver report and obtain feedback

Security Sub-process (Initiate courses of action)

The security sub-process is initiated based on the results of the hunt.

To initiate courses of action:


- Activate incident response plan
- Develop new threat detection content
- Perform a security architecture review
- Resource allocation assignments
- Incorporate lessons learned into Cyber Defense functions

Pivoting

Pivoting helps threat hunters push through when a hunt stalls. This provides a chance to be creative in finding unique information to use for pivoting.

Pivoting sources can include:


- **Previous hunts**
 - Has activity like this been seen before or something like it? What was done then?
 - Can the same analysis/actions be repeated? This saves time and effort.
 - If it was insufficient then, now you have identified another place to look.
- **Intelligence (internal and external)**
 - What does the intel say about this indicator?
 - What kind of activity is it tied to?
 - What else should be searched for?
- **Automated tools (sandboxes, webpage scanners)**
 - Does this file appear malicious?
 - What other indicators can be acquired from it?
- **Community sources (blogs, etc.)**
 - What is the community saying about this? Have others seen similar activities?
- **Indicator investigation (whois, pDNS)**
 - Are the indicators tied to high-risk infrastructure (IPs in unexpected countries' domain registrars)? Are they tied to other suspicious elements (other domains with same registrant)?



Threat hunting helps shore up defenses by unearthing previously unknown compromises or gaps in security controls. Intelligence-driven hypotheses are predicated on an understanding of the cyber threat landscape, common attacker behavior, the organization's security posture, business processes and the attack surface to guide hunt efforts. Threat hunters must stay nimble and maintain awareness of shifts in organizational decisions or commentary from senior leadership—such as announcements of potential mergers or acquisitions, expansion into different geographic locales, or other public announcements—and cyber threat actors' motivations and targeting.








DETECTING AND INVESTIGATING MALICIOUS ACTIVITY

The Detect function identifies malicious behavior based on activity seen in an environment. Preparation is the initial, and arguably most important, phase of the Detect function. Preparation brings insight, forethought, planning, justification, capability, training, prioritization, and coordination. Each of these elements are vital in the burgeoning field of threat detection engineering—the practice of identifying a relevant threat, measuring affectations to a system that indicate the threat’s presence, and informing operators so that they can select an effective course of action.



As this is an engineering discipline, practitioners execute in an iterative fashion:

1. Identify the organization's adversarial value.
2. Select a known threat actor likely to target the organization.
3. Profile the behavior of the threat actor.
4. Identify the threat actor's TTPs.
5. Determine measurable impact of each TTP to the system.
6. Implement visibility to showcase the measure.
7. Qualify the measurement's fidelity.
8. Set alert thresholds.
9. Design and build the alert.
10. Test, validate, and integrate the alert into the production environment.
11. Review triggered alerts for accuracy and calculate its value.
12. Sustain, enhance, or dispose of the alert based on the review findings.

While there appears to be a glut of threat actors, there is a numerable quantity of tactics and techniques used by these groups. Consequently, there is significant overlap between threat actors and TTPs.



Threat Intelligence in the SOC

Threat intelligence provides the core of effective detection. Intelligence makes up the first four steps of the threat detection engineering process. These steps, actioned by a threat intelligence analyst, determine the Who, How, and Why of cyber-attacks. By contextualizing and categorizing incidents, reports, and data feeds, the intelligence analyst generates insights that define the attacker's objective and methodology. These insights inform the selection of threat actors and their techniques for detection prioritization. A detection engineer then works with system engineers, administrators, and developers to characterize and measure the attacker's actions.

In some attacks, the attacker's impact is well understood and easily measured; visibility can be achieved by placing an IDS in front of a server or an EDR agent on a user workstation. With other attacks, the effects are more nebulous—for instance, a unique privilege escalation or remote code execution (RCE) in a custom web application. Perhaps the most common attack pattern, however, is when the attacker leverages compromised, but legitimate credentials and abuses the environment in non-sophisticated ways—such as logging into a remote desktop server using stolen compromised admin credentials.

The detection engineer works through these various threat models and develops system visibility requirements, which define system logging and or additional security sensors that are used to measure the presence of an adversary's tool or execution of an adversary's technique.

A Nuanced Approach

The approach to gathering visibility requirements then defining system logging and additional sensor needs is often opposite of the approach taken by many organizations. The typical approach has the SOC rely on vendor-defined logging supplemented with non-validated security appliances, all logging to a centralized SIEM. Then, the SIEM engineer or SOC analyst is told to monitor logs and escalate when there is a problem. However, this approach is neither driven by threat intelligence nor is it prioritized for loss prevention. This gluttonous approach to logging regularly leads to SOC operational and capital expenses ballooning over time with little to no increased value to the business.

To avoid this high-cost, low-value situation, a methodical and intentional approach should be taken for log and sensor selection. For the most efficient execution of security monitoring, at the time of ingestion, logs should be known to be relevant to a detection use case, or investigation forensics. System administrators and other stakeholders need to be included in the security engineering process.

Detection Development

Once visibility requirements are established, detection engineers continue their work contextualizing and affirming the signals they receive from the monitored systems. For instance, deletion of volume shadow copies may be highly indicative of incoming ransomware for some organizations. For others, this technique may also be used to clear free space on a server's taxed hard drives. This contextualization and measurement of signal actionability is unique to every organization.

Identity technologies are important to continuously answer: Do I know you? Do I trust you? How much access will I give you? In order to do this well, it is important to gather context. This is a first line of defense, and in-line defense to stop inhuman and fraudulent access attempts.

Mary Writz, VP Product Management, ForgeRock

Signal actionability measures how effective a signal is at determining malicious intent. This is similar to, but distinct from, the standard binary classification terms such as “false positive” and “true negative.” Extending the previous example, a system log may accurately identify the deletion of a volume shadow copy but does not verify if the deletion was malicious. The detection engineer is concerned both with the accuracy of the signal and the value of the signal to the SOC.

The detection engineer has multiple tools available to enrich and contextualize signals into actionable alerts. Some signal types may require baselining and setting an alerting threshold. For instance, a threat actor may intend to disrupt business by purposely failing logins and locking out a large user base. While a single account lockout may not be actionable, multiple account lockouts, from a single source, well above the baseline, is actionable.

Additionally, the detection engineer may seek to combine various signals into a single alert. The engineer may determine that system administrators do commonly delete volume shadow copies, but they never do so by using PS Exec. Consequently, an alert may look for the presence of both the PS Exec service and volume shadow deletion. These nuances of alerting are driven intelligence observed, reported, and shared.

Every detection use case must be managed through a detection lifecycle. Adversaries change, their capabilities and manipulations to the victim environments change. As such, the detections must also change. To maintain the health of each alert, the methodology continues into testing, validation, continuous evaluation, and eventually disposition.

The Challenged SOC

When SOC's face overwhelming alert volumes it means their detection use cases have not been managed, evaluated, and prioritized. Alerts pouring into the SOC are a symptom of the "log-it-all-and-let-the-analyst-sort-it-out" approach. Even when these alerts are enriched by SIEM datasets and SOAR lookups, they still require additional analysis to validate and initiate investigations.

The classic SOC organizes the SOC analysts into tiers like a help desk. The lowest tier follows step-by-step scripts to "manually automate" certain tasks. In other words, while certain tasks are repetitive, they may not be rudimentary, or the interface may not be programmatically accessible. In that case, it is more efficient to throw people at it than to automate using a SOAR-like platform.

While executing these scripted procedures, a flow chart makes the Tier-1 analyst's decisions. The analyst does not make the decisions. If the flow chart is insufficient, the incident is escalated to a Tier-2 analyst who may be empowered to think more critically.



A man with a beard, wearing a grey suit jacket, is seen from the side, looking towards a server rack in a data center. The server racks are illuminated with blue and white lights, creating a high-tech atmosphere. The man's hand is resting on a ledge in front of the server racks.

Example

During a casual investigation of an IDS alert, an analyst processed it according to the documented procedure. It was a high-severity IDS event from a user's workstation to an internal web server. However, the web server wasn't vulnerable to the attack, so the investigation was closed. The SOC manager did a review of high severity alerts later and recognized that the attack was launched from an internal asset against another internal asset. This meant there was an attacker controlling a user's workstation. The analyst had followed the written procedures for IDS events but failed in this case because the procedures were written with externally sourced attacks in mind.



SOCs also commonly have their analysts specialize in alert sources. Some analysts may be focused on email alerts, others focused on EDR, etc. The issue is that an effective threat narrative may cross multiple alerts. For example, a suspicious email alert that highlighted a suspicious subject line, such as “Package delivered,” may contain a payload that, when opened triggers an IDS alert to a suspicious user agent. When persistence is established, the EDR tool may alert to a suspicious service creation. A subsequent crypto miner may have been downloaded and blocked by the endpoint antivirus.

If this attack story is evaluated from a threat narrative approach, this is clearly a single incident that happens to have multiple detections as the attack progresses. Yet, separating analyst responsibility based on signal type gives them only a small piece of the story.

When analyzing the suspicious “Package delivered” email, the email may be crafted to resemble a reply to an existing thread with a known third-party business, leading the email analyst to dismiss the threat. The user agent highlighted by the IDS alert may look very similar to common Mozilla user agents already used in the environment. Since antivirus blocked the crypto miner, that may not even be investigated by the SOC.

In this example, while the EDR alert may trigger a response from the SOC, the investigation must still uncover the point of entry of this attack, even though this alert was already reviewed by the SOC.

Situational awareness becomes increasingly difficult as SOCs onboard additional telemetry. As security and operational logs are ingested by the SIEM, more alerts are triggered, and more noise generated, for analysts to sift through. This in turn distracts the SOC from identifying real threats. Most alerts are benign, so analysts spend a lot of time sifting through non-actionable alerts. The most capable SOC analysts are required to extract value from the low fidelity content. It is tedious work and leaves analysts burnt out.

At the same time, organizations are reasonably concerned with the possibility of a false negative; when the attacker is present, but security does not respond. So even though a SOC is overwhelmed, the organization hesitates to tune, reprioritize, or even silence false positive noise in fear of missing out. However, an uncomfortable truth hides behind the dashboards: alert fatigue means false positives create false negatives.

The Optimized SOC

Instead of attempting to investigate each signal that may or may not indicate malicious behavior, an effective SOC merges multiple related alerts into a unified investigation. Ultimately, the goal of the SOC is to determine whether a system or an account has been compromised. To this end, SOC analysts must gather all contextual clues about potential victims.

Instead of examining each alert separately, analysts should investigate the victim, addressing the relevant alerts together into a cohesive investigation. This methodology alters the skill set of the SOC analyst.

Unlike the traditional SOC where analysts are dedicated to log sources and flow charts, the optimized SOC takes advantage of their analysts' training, experience, and critical thinking. This SOC analyst operates like a detective, leveraging the clues, evidence, and forensic artifacts to uncover the story behind the incident. This analyst is savvy with each of the relevant log sources and security appliances, using every available tool to triage, classify, scope, and contain the threat.

Activating the SOC

Investigations typically start with either external notifications or an alert to the SOC. These alerts may be generated by security tools such as a network intrusion detection system, or antivirus, or the alerts may be triggered by logic applied to logs through a SIEM, or human notification of suspicious activity. Regardless, this alert gives the analyst the primary starting point.

SOC analysts do not need to see every security event, instead they should be presented with high-fidelity alerts which warrant further investigation. Overloading the analysts with irrelevant security events can blind them to higher-priority security events which may require access to additional information to perform investigations to determine compromise versus threats to be mitigated. Security systems should not be streaming endless events down the screen.

More visibility requires more technologies, which creates more information for analysts to sift through. This requires extra tooling and automation to be put in place to make sure the additional visibility does not blind the analysts.

Saskia Hoffmann, Security Operations Manager at eHealth

41%

of incidents Mandiant investigated in 2020 stemmed from external notification*

Intelligence helps analysts focus on what is important; however, without the proper mechanisms in place to operationalize that intelligence, they often don't know what to do with it and it is ignored. Automating the use of intelligence in a SIEM or automated defense technology utilized by the SOC produces a significant benefit to analysts. Proper implementation of the intelligence would allow SOC technologies to automatically prioritize events tied to widely exploited vulnerabilities with critical impact to the organization. Tracing these events back to common TTPs of known attack groups or clusters can also assist in the investigation to accelerate analysis and playbook selection.

A lack of intelligence means the analyst must rely on their training, experience and intuition alone to track the threat actor.

Daniel Nutting, Manager for Cyber Defense Operations Consulting, Mandiant

How the analyst proceeds and the general objectives the analyst pursues are defined by SOC processes and coordinating playbooks.

*Mandiant M-Trends 2021



Using Automated Defense Technologies

Recent developments in automated defense technologies mitigate the risk of procedural and human error in the investigation and evidence gathering process. Automated defense technologies emulate the reasoning and decision-making of expert security analysts and incorporate data from security tools across the infrastructure to enrich alerts and determine if they are malicious and actionable, or benign.

Automated defense technologies augment the role and responsibilities of security analysts. By providing security analysts with scoped and prioritized investigative cases from which to quickly choose the appropriate incident response path, the security analysts can focus more on identifying adversarial intent versus vetting false positives.

By providing analysts with prioritized and enriched data, automated defense technologies remove many of the initial triage steps when investigating suspicious activity. This allows the analysts to spend their time analyzing and responding to activity rather than validating alerts. Automated defense technologies will never replace a human analyst's ability to understand the context of the data; however, it is a highly useful tool in an analyst's toolkit. It should be the goal of all SOC teams to present analysts with high fidelity alerts that contain actionable data so that skilled analysts can reduce the dwell time of attackers and prevent catastrophic impact.

Automated defense technologies scope all related systems and activity for the duration of an attack. The incident may span a few seconds or many days. The technology prioritizes the incident investigation, factoring in the scope, asset criticality, attack stage and confidence in the escalation.

The prioritized incident is then presented to the analyst with supporting evidence including:


- The identified malicious behaviors and signatures
- An event timeline (a series of events from various security tools over time)
- The internal systems and assets impacted
- Attributed threat intelligence data
- Attack stage progression mapped to the MITRE ATT&CK[®] framework

Having an in-depth understanding of TTPs based on reliable threat intelligence is critical for SOC analysts. Understanding past breaches aids analysts in predicting future attacker activity to a single system and an enterprise-wide incident.

David Lindquist, Managed Defense Operations Manager, Mandiant

Revisiting the high severity IDS alert example above, automated defense technology would have provided the following:

"High severity IDS event from user's workstation to an internal web server. Identified the IDS signature involved the execution of Empire PowerShell Payload. The same payload signature is observed executing from several other internal workstations, aimed at multiple critical servers. Investigation escalated to the incident response team."





RESPONDING TO COMPROMISE

The Respond function of the Cyber Defense domain is responsible for the response and potential remediation of suspicious activity in an enterprise environment. As the Detect and Hunt functions identify suspicious activity, the Respond function confirms if that activity is malicious. It then takes actions to understand the full extent of compromise, minimize business impact, return computing services to normal operations, eliminate the threat from the environment, and enable compliance to applicable regulations and standards. To prevent repeat incidents, the Respond function is responsible for identifying lessons learned from the incident and directing tactical and strategic enhancements through the Command and Control function. The Respond function is also responsible for feeding observations back to the Intelligence function.

Initial Triage

The Respond function begins when evidence of potential unauthorized activity is escalated for further investigation. This initiates the Triage phase, which guides future phases of the investigation based on the following:

- Confirm the accuracy of information provided by the alert.
- Determine if the alert is actionable.
- Determine immediate remediation steps if any.
- Prioritize incident queue based on other active incidents.

The Triage phase matures and qualifies events highlighted by the Detect and Hunt functions into alerts that indicate malicious activity. The analysts that drive this function may collect additional evidence to help understand and document the context of the suspicious artifacts. Through this process, if the analyst determines that the evidence indicates a breach, they may declare an incident and initiate a full investigation.

The first step of the triage analysis is to determine if the signal received from the Detect or Hunt function is a true positive or a false positive. A true positive occurs when the detection matches the intent of the use case. A false positive occurs when the detection does not match the intent of the use case. Similarly, a false negative is when the implementation of the use case fails to detect activity that would match its intent.

For instance, a SOC may code a SIEM use case to alert to any invocation of the “whoami” command. Attackers and penetration testers both commonly execute this to verify their control of the victim platform. This can indicate malicious presence, but legitimate users also run this command. If the SIEM accurately detects the execution of “whoami” by a legitimate user, is this a false positive? It depends on the intent of the use case. If this alert is meant to highlight an event of interest that does not alone necessitate an investigation, it is a true positive. However, if the use case is intended to detect attackers invoking “whoami,” then it is a false positive.



A true positive may be benign, such as when an administrator executes “whoami.” Alerts that are prone to benign closures should be considered for informational use only. Leverage these activities to contextualize other alerts against the same endpoint or account. False positives necessitate tuning: changing the security or operational configuration to prevent detections that do not match the use case intent. Tuning may involve things like finessing a Snort rule, disabling a specific antivirus signature, or ignoring a known good file hash.

If the detection is determined to be a true positive, then the analyst collects enough information to determine the most appropriate next steps. Triage involves collecting data, assessing the risk, and determining the type of incident that has occurred.

Triage consists of the following general steps:

- Collect and analyze information
- Identify decision points and determine next steps
- Review playbooks

Data Collection and Analysis

Another critical responsibility during the triage phase is for the responder to focus on the collection and analysis of data surrounding the potential incident. Information collection involves gathering both technical and non-technical information surrounding the facts of the incident. The extent of data collection should be driven by the facts surrounding the potential incident, including the criticality of the incident.

When collecting technical data, consider the extent to which the integrity of the target system must be preserved. For example, if it is likely that the potential incident would result in litigation, it is important to document each interaction with the system and to maintain chain of custody documentation.

Non-technical information may come in from a self-report from an end user or a suspicious activity report where access to raw data may not be immediately accessible. If an incident or other form of malicious activity is reported by a user, it may be necessary to contact the person who reported the incident directly. This could be to obtain details about how the system was being used prior to the suspicious event being identified. For example, if there are signs of malware or attacker tools on the system, ask questions to determine whether the user is aware of these artifacts. This can help differentiate a misuse of resources situation from a critical APT incident. Obtain as many dates and times as the user can provide, such as when the user first noticed the issue, and what steps the user took to try to resolve the issue on their own.

Technical information about the environment can be collected from system administrators or monitoring tools that will provide useful context for understanding the situation. For example, an analyst may obtain hostnames and IP addresses of systems that are associated with an alert. They may then obtain information such as the operating systems, program affiliation, purpose of the system, and the functional titles of users who interact with the system regularly.

Important evidence can be lost if it is not collected immediately, therefore, the analyst should consider that when prioritizing evidence collection. If the response activities are likely to impact system performance, consult the appropriate IT and Infrastructure stakeholders in advance.

Decision Points and Next Steps

Depending on the nature of the activity identified, the analyst may determine that there is enough context known to immediately direct containment activities and close the incident. Mature incident playbooks will help the analysts make this decision. If the analyst determines that this cannot be resolved immediately, then evidence is passed to the Investigation Lifecycle phase to continue the investigation.

At this point, the analyst should work through the Command and Control functions to bring in other stakeholders as needed and make key decisions—such as whether to utilize cyber insurance or involve legal counsel. This can be especially important if an organization is bound by regulatory compliance concerning a breach. Some regulations require notification of data theft or ransomware infection, and the clock often begins ticking for that notification as soon as the compromise is identified. Legal counsel will layout the notification requirements and provide priorities to meet the requirements.

Breach notification obligations often force decisions to be made in hours or even minutes. Because these decisions are often made on incomplete intelligence and even less certain law, having counsel familiar with the business, available on speed dial is critical to effective response.

Gerry Stegmaier, Law Partner, Reed Smith LLP





Playbook Review

The analyst driving the triage process is responsible for maintaining playbooks for common incident response scenarios. As part of this maintenance, these playbooks should be reviewed and revised often to ensure they are kept up to date as there are few things more challenging than handling an incident using an outdated guide that refers to outdated processes or infrastructure.

Playbooks that are regularly revised and updated are clearly being used and provide value to the security operations team. But as soon as the playbook goes stale, it no longer acts as a reliable resource to analysts. Analysts should be empowered to make immediate updates to their playbooks as soon as they identify discrepancies.

Similarly, the SOC should develop metrics for playbook usability. How often are playbooks used? Which playbooks are never used? What search terms do analysts use to find playbooks? Do they need to execute multiple searches to find the right playbook? When was the last time a playbook was updated, reviewed, and executed? Questions like this help to nurture a healthy knowledge base for the SOC.

Investigation Lifecycle

The goal of the investigation lifecycle phase is to answer key questions about the attack. This context will inform other stakeholders to allow them to make strategic decisions about legal and regulatory obligations, communications to employees and customers, and other critical business decisions. The findings from this phase will also inform the incident containment and attacker eradication plans.

The activities of the investigation phase typically include:

- Determining the scope of the intrusion and whether it is ongoing
- Determining the earliest date of compromise and cause of intrusion
- Determining the type and extent of data exposed to the attacker
- Identifying the threat actor and motives
- Providing context to drive the incident containment and attacker eradication plans

Comprehensive investigations of attacker activity depend on a cyclical process called the Investigation Lifecycle. This process starts with an initial lead from the triage phase of the investigation. This could be any forensic artifact that indicates unauthorized activity from a threat actor. For example, Windows event log entries on a server indicating a successful login by the attacker from a workstation involved in a successful phishing attack.

Investigators can then preserve relevant evidence and conduct deeper analysis of forensic artifacts.

This could include:

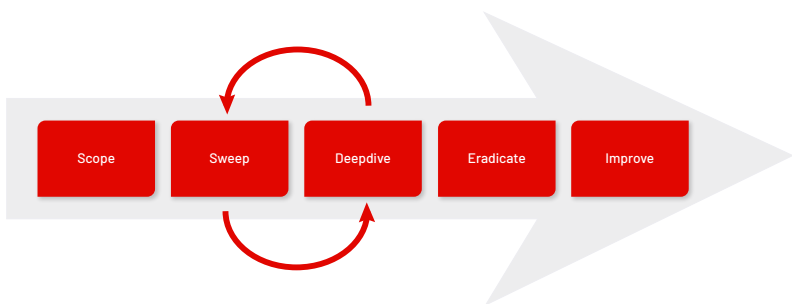
- Live Response triage of impacted systems
- Analysis of forensic images of impacted systems
- Malware analysis
- Log analysis
- Network traffic analysis
- Intel queries

The output of these analyses will potentially feed two workstreams:

- 1. Identify Additional Leads** – The investigator develops findings and timelines of attacker activity associated with the analyzed forensic artifacts. This leads to additional leads that restart the analysis cycle.
- 2. Inform Environment Sweeps** – As investigators learn more about the attacker and the attack, they codify that information into what are called Indicators of Compromise (IOCs). Investigators then conduct IOC sweeps, where they utilize enterprise response tools to search sources of evidence across their environment for similar artifacts. IOC sweeps are meant to be paired with manual analysis of the results, which allows IOCs to be more flexible than typical signatures. Investigators may conduct IOC sweeps for anything from very specific artifacts (for example, artifacts associated with a specific malware family) to much broader patterns (for example, extracting items from the WMI repository on Windows endpoints associated with persistence to look for anomalies). Investigators then analyze the results to identify additional leads and restart the analysis cycle.

Throughout this iterative process, investigators create a comprehensive timeline of attacker activity. This timeline drives the answers to the key investigative questions and allows investigators to fully scope the incident.

Figure 7: Investigation Lifecycle



In a recent investigation, Mandiant helped thwart an adversary who used compromised credentials against the organization's VPN access. The organization detected account abuse when the adversary emailed the help desk asking for a password reset, the help desk responded, and the actual user reported the impersonation.

With this alert, the initial evidence was limited to the IP address and the account used by the adversary.

The investigation team began with three questions:

1. What other activity has been seen from this IP address?
2. What sensitive systems and data did the compromised account have access to?
3. What did the adversary do with this account?

In asking these questions, the investigation team also developed an investigative plan. These questions inform the analysts what artifacts and evidence they need to seek out.

1. By examining VPN logs and firewall logs they found other compromised accounts and access to the organization's web facing email.
2. By examining security groups and access control lists, they determined the first compromised account had access to a sensitive file server mounted as a network share to one of the accessed endpoints.
3. By leveraging user activity artifacts, they found the adversary broadly searching for documents using queries such as *.pdf.

As the investigation continued, the adversarial intent became clearer. In particular, the endpoint artifacts showcased the adversary searching for broad,

Example

untargeted queries, like *.pdf, *.xls, *password*. At this stage, the adversary exhibited only reconnaissance activities. Furthermore, the attacker did not use keywords and phrases relevant to the organization's name, sensitive projects, or key personnel. The untargeted queries intimated that the attackers were still in an information gathering phase or that this was possibly an opportunistic attack. Opportunistic attacks like these are more likely to be monetized by various means, such as selling access to another attacker or ransomware.

By profiling the attacker, the observed tooling and the adversary's current control, ransomware did not appear likely. Although the attacker had been tracked across dozens of other intrusions and participated in data theft, they had never been observed deploying ransomware to victim environments. In addition, this threat actor did not immediately pursue domain admin permissions or seek other means for distributing ransomware across the enterprise. They kept the systems they accessed focused on low-privileged users who had access to file shares.

When specific intelligence is available, the next steps can be quickly developed based on the threat actor's habits and practices. Otherwise, defenders must rely on generic intelligence, surmising the adversary's objectives along the way.


A critical question that must be asked regularly is, "Does this make sense?"



Investigation Accelerators

Strong attacker intelligence can significantly increase the efficiency of an investigation. As investigators collect IOCs and learn more about an attacker, they can leverage CTI to attribute this activity to known attack groups. This information can help investigators understand the potential motives of the attacker, what the attackers have done in previous campaigns, and how they have done it. They can then use this information to prioritize data collection and analysis tasks, and possibly containment and eradication tasks (discussed in the next section).

This can have several significant business impacts. First, it maximizes the efficiency of valuable human capital, and frees up important resources to complete other tasks. Second, every collection action that investigators take has a non-zero resource cost associated with it. By prioritizing collection actions based on CTI, investigators can limit the impact to the environment they are investigating. Finally, investigations are often taking place while the intrusion is being investigated—positioning the investigators in a race against the threat actors to enumerate their access to the environment and eradicate it before the attacker completes their mission. By leveraging CTI, investigators can get answers faster—and reduce the likelihood that attackers complete their mission before eradication steps are complete.





Example

During an investigation of what initially appeared to be commodity phishing and malware, analysis identified a specific TLS certificate within the payload uniquely used by FIN8. Based on this attribution, the incident responder targeted subsequent analysis to known persistence and privilege escalation techniques used by the adversary, which identified relevant forensic artifacts in mere minutes, compared to what would have otherwise taken hours or days to comprehensively review.



Microservices

The investigation phase of the Respond function includes some extremely specialized technical tasks. For most organizations it is not cost effective to develop in-house skills, such as malware analysis. Even when organizations do develop internal teams for tasks like malware analysis, it is difficult and expensive to retain that talent. Nevertheless, these tasks can be critical to the success of an enterprise scale intrusion investigation.

Rather than attempt the expensive and challenging task of developing these skill sets in-house, organizations should consider establishing relationships with specialized consulting firms that offer microservices. Microservices can allow for outsourcing individual analysis tasks, rather than full investigations, extending the ability of the organization's Respond function, as if there was a team devoted to these specialized tasks.

Tasks that organizations should consider outsourcing through microservices include:

- Malware analysis
- Forensic analysis
- Intelligence gathering

IOC Hunting Automation

As described in the Investigation Lifecycle section, investigators codify IOCs and drive IOC sweeps throughout an investigation. For simple IOCs (IP addresses, domain names, and file hashes), there are usually well-defined and consistent sources of evidence that investigators leverage for sweeping activities. There are also typically repeatable processes for searching those sources of evidence to identify new leads.

This consistency and repeatability create an opportunity to increase efficiency through automation and orchestration tools. Organizations can create automated routines to search relevant evidence sources for predefined categories of IOCs. This reduces the amount of valuable human capital dedicated to mundane IOC sweeps during an intrusion investigation, reserving it for more complex IOC sweeps or other critical investigative tasks.

IR Retainers

Most organizations do not regularly respond to large-scale intrusions. They often lack the experience necessary to conduct a comprehensive investigation for significant or complicated intrusions. To account for this gap, organizations should obtain incident response retainers (IRR) from IR providers. Due to the importance of speed in an investigation and response – organizations should consider establishing service level agreements (SLAs) with their IR provider. If an IRR is established without an SLA, organizations should consider establishing IRR agreements with multiple vendors to mitigate the risk of a provider being busy and unable to provide the service.

Incident Remediation

The goals of the Remediation phase are to remove the threat from the environment and restore systems to normal operational conditions. In addition, the Remediation phase can inform the Lessons Learned phase to direct enhancements to the security posture of the organization.

The scope and scale of the remediation phase can vary widely based on the investigative findings from the Triage and Investigation Lifecycle phases, as well as the size and complexity of the impacted environment. For most incidents, Remediation activities will involve taking simple tactical steps to remove an attacker's access. These could include quickly disconnecting a compromised system from the network and reimaging it, disabling access from impactful accounts, changing passwords of impacted accounts, or blocking access to known command and control channels.

In some circumstances, comprehensive planning may be required to successfully remediate the incident.

Examples of circumstances that may call for a more thorough remediation plan include:

- Incidents where an attacker has access to a large scope of systems in the environment
- Incidents where an attacker may have multiple unknown entry points to an environment
- Incidents with a long dwell time
- Situations where remediation activities necessary to respond to a less severe threat are complex, or require coordination between several internal groups to implement

In these circumstances, organizations should follow a 2-part, 4-stage process that takes place in parallel with the Investigation Lifecycle:

Part 1: Remediate the current incident

Containment - Take actions to disrupt attacker activities, monitor, harden and remove the attacker from a sensitive system or network segment to regain control of the affected environment.

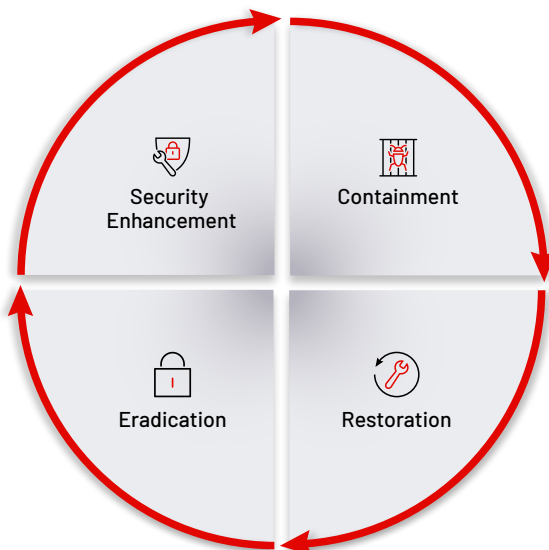
Restoration - Take actions to restore environments impacted by destructive attacks. For example, restore encrypted endpoints, applications, and services, to re-establish business functionalities.

Eradication - Remove an attacker from the environment and implement security improvements to inhibit the attacker from quickly regaining access to the environment. Should be performed in a concise and coordinated manner.

Part 2: Improve the organization's security posture

Security Enhancement - Inform the Lessons Learned phase to enhance the security posture of the organization (e.g., process improvements, privileged account management, network re-architecture, etc.)

Figure 8. Incident remediation flow



The execution of these stages often requires communication and coordination with stakeholders outside of the Cyber Defense function. It also involves decision making that can have a significant impact on business operations or create reputational damage for the company. Therefore, the actions described in each of these phases are coordinated with stakeholders through the Command and Control function.

These stakeholders could include:

- Executive Leadership
- Legal
- Compliance
- Internal and External Communication Teams
- Information Technology
- Application and System Owners
- Human Resources

Remediation plans must be customized based on the details of an intrusion, as well as the victim organization's unique operational complexities and business needs. What worked well for one incident may not be advisable for another. While remediation playbooks should be robust, they should also provide room for flexibility and decision making—and the Cyber Defense team needs to be empowered to make those decisions.

One of the key considerations in the Remediation phase is the timing of containment and eradication activities. The more that is understood about an attack increases the likelihood that containment and eradication actions will be successful. Hasty actions can be counterproductive, prolonging the intrusion and increasing the organization's risk of negative impact.

In other circumstances, organizations can significantly reduce the impact of a breach by acting quickly, even when they have incomplete information about attacker activity. For example, if the attack is still in a preliminary phase and the attacker does not have multiple ways to enter an environment, quickly eliminating their known access could end the threat entirely. Even when an attack is further along, there are circumstances when immediate containment actions should be taken. For example, if there is a reason to believe the attacker is going to initiate a destructive attack, like a large scale ransomware deployment, then the organization may be willing to tolerate immediate and aggressive containment actions that have a large business impact, simply because they will have less of an impact than a successful destructive attack.

Organizations should empower analysts to make these decisions with robust playbooks, training, and experience in simulation exercises.

It is very common for organizations performing their own incident response to panic and attempt a premature remediation. They often jump to remediation efforts and introduce changes that complicate the investigation. This whack-a-mole approach will lengthen the investigation, cause incomplete remediation efforts and can lead to repeat attacks.

Eric Scales, Vice President, Mandiant

Containment

The goal of the Containment stage is to limit an attacker's access to an environment and support the investigation while an eradication plan is staged and executed.

During the Containment stage, the remediation team may direct short-term tactical actions that limit an attacker's access to data and systems or disrupts their activities. In addition, the remediation team may direct actions to increase visibility in the environment or harden the environment to prevent re-compromise after eradication.

Actions taken during the Containment stage commonly include:

- Enhancing logging and monitoring
- Patching and mitigating vulnerabilities exploited by an attacker
- Hardening system-to-system communications and endpoint controls
- Limiting exposure of credentials on endpoints
- Reducing the scope of privileged accounts
- Hardening local administrative accounts and permissions
- Reviewing and hardening remote access methods or access to cloud systems
- Temporarily revoking access to systems with critical data or taking them offline

During the Containment stage, the remediation team prepares for the eradication event. This can entail discovering and documenting important information that will enable eradication activities, including:

- All existing backend authentication mechanisms
- Organizational unit (OU) structure for housing accounts in Active Directory
- Privileged accounts across the enterprise and cloud platforms
- Egress paths and technology in place to restrict egress traffic
- Application and business unit owners
- Remote access technologies
- Remotely accessible SaaS platforms



Restoration

The goal of the Restoration stage is to recover business operations for systems directly impacted by the incident. This stage is usually only needed in destructive attacks, such as when attackers deploy ransomware.

In the case of destructive attacks, the victim organization may find themselves locked out from accessing key systems and services or that critical business processes are severely disrupted. In these cases, the organization must execute a plan to recover and reconstitute their environment in parallel to planning and executing containment and eradication. It is critical to coordinate these actions to regain control of the environment to support recovery and reconstitution efforts.

Common actions include:

- Creating a communications plan
- Establishing a secure VLAN for recovery and reconstitution
- Restoring business services
- Rebuilding Active Directory or other identified platforms
- Restoring systems from backup
- Hardening rebuilt endpoints and services



Eradication

The goal of the Eradication stage is to eliminate unauthorized access to and regain control of the impacted environment. In some circumstances, containment and eradication actions are performed simultaneously. In other cases such as active data exfiltration, it is necessary to first disrupt the attacker's activity, while planning for a more extensive eradication event.

Typically, the eradication plan will call for the remediation team to coordinate a series of ordered tactical actions in a short period of time.

These typically include:

- Implementing network blocks and DNS sinkholes
- Disabling compromised accounts
- Removing infected systems
- Implementing privileged account security plan
- Conducting enterprise-password reset
- Rotating local administrator passwords
- Replacing compromised systems



Security Enhancement

The final stage of Remediation is Security Enhancement. The goal of this stage is to inform the Lessons Learned phase to reduce the likelihood of future breaches. Attackers regularly retarget victims post-eradication. It is critical to aggressively mitigate risk to an environment after the successful remediation of a breach.

Ideally, all incident investigations should determine root cause. This is important to improve the organization's ability to prevent such incidents from occurring in the future, and to enhance monitoring mechanisms to detect IOCs more effectively.

Through the process of executing the previous three stages of remediation, the remediation team often uncovers weaknesses in the impacted environment and identifies areas for additional assessment. During this phase, the remediation team documents those lessons and generates recommendations to the organization. Those recommendations are fed into the Lessons Learned phase for action.

Remediation Accelerators

Similar to the investigation process, automation can be used during the remediation process to maximize the efficiency of an incident response team. Some of the individual tactical steps that make up containment and eradication plans are common across different incidents and are enforced through consistent and predictable technologies. That consistency and predictability creates an opportunity to leverage automation and orchestration tools.

Some organizations have taken this a step further, allowing detection toolsets to automatically trigger remediation actions without relying on human intervention. This frees up valuable resources for other tasks and can significantly increase the speed of common containment actions, greatly reducing the likelihood of an attacker completing their mission.

For example, many organizations do not assess antivirus alerts if the file in question was blocked or quarantined. But these attacks, although mitigated, can be a treasure trove of valuable information, containing details about live threats that have impacted the environment.


Organizations can utilize automation and orchestration tools to execute some of the information gathering from the attack to automatically identify and block attacker infrastructure.

A playbook can be built that:

1. Identifies a quarantined antivirus event
2. Retrieves the quarantined file
3. Submits the file to a malware sandbox
4. Pulls IOCs from the sandbox (IPs, URLs, file hashes)
5. Pushes IOCs to detection and blocking platforms (IPs -> Firewall, URLs -> Proxy, Hashes -> Application Control)
6. Runs searches for other endpoints matching on IOCs
7. Contains endpoints and pulls triage forensics on matches

Attacker Intelligence

Robust CTI can help responders infer the objectives of the threat actors and tailor their remediation plan accordingly. Tailoring that plan may include protecting certain types of data or systems likely to be targeted by the attacker. It could also help determine the timing and aggressiveness of countermeasures.



For example, by leveraging CTI, investigators can determine that a threat actor with unauthorized access to their environment is likely to deploy ransomware to monetize their breach. In this scenario, businesses may be more willing to accept the risk that countermeasures cause short term disruption to business because the disruption caused by a widespread ransomware deployment could be much greater. They are more likely to deploy aggressive and immediate containment activities that severely limit an attacker's ability to deploy malware at scale.

Examples of these could include:

- Rapidly de-privileging of legitimate accounts
- Rotating passwords
- Using host-based firewalls to block administrative ports and significantly restrict lateral movement
- Disabling remote access technologies

Lessons Learned

The goal of the final phase of the Respond function is to digest lessons learned from the investigation and remediation phases and drive changes to improve the organization's security posture. Recommendations developed in this phase should flow through the Command and Control function to be evaluated and executed by other stakeholders in the organization.

This process should include analysis to identify root cause, not just for the initial infection vector, but for the attacker's success across the entire attacker lifecycle. Recommendations should encompass changes that could inhibit future attacks, allow the organization to better detect malicious activity, or help the organization respond to future breaches.

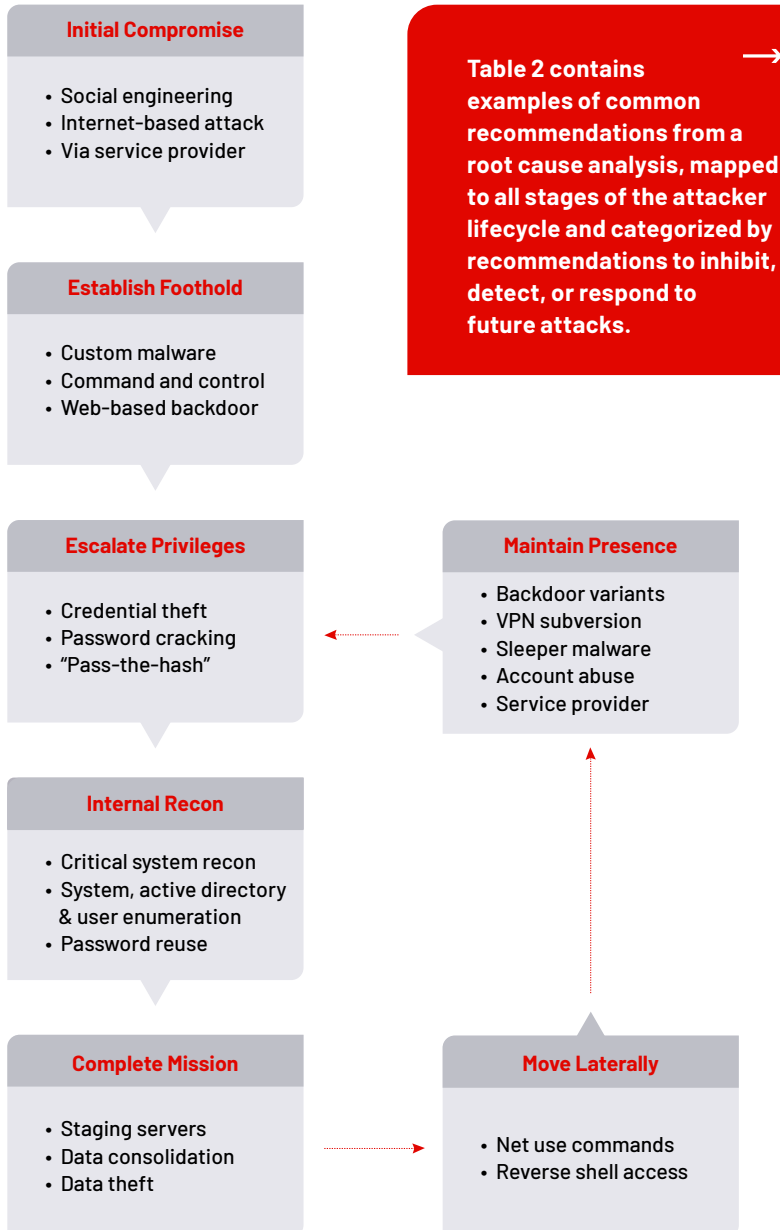
In some cases, this process may identify areas for future analysis outside the scope of the Lessons Learned phase. For example, if an attacker was able to quickly escalate their privileges in an Active Directory Forest – the root cause analysis may identify specific weaknesses that the attacker was able to exploit and recommendations to address those weaknesses. The recommendations may also include the need for a comprehensive assessment of the Active Directory environment performed by subject matter experts in Active Directory security.

The Lessons Learned phase should also include a review of the incident documentation. In most cases, incident documentation can be driven through the organization's incident management systems. For more complex or impactful incidents, there may be a need to create a formal Incident Report. This decision should be routed through the Command and Control function.

In addition to incident information, appropriate threat intelligence should also be collected, compiled, and operationalized throughout the environment. Victim organizations are often retargeted after a successful eradication event. It is important to leverage the intelligence gained about the threat actor through the Respond function to empower more effective detection of similar activity in the future.

The following figure illustrates the lifecycle common to targeted attacks.


Figure 9: Targeted Attack Lifecycle




	Initial Compromise	Establish Foothold / Maintain Presence	Escalate Privileges
Inhibit	<ul style="list-style-type: none"> Enhance phishing prevention program Launch user awareness trainings Conduct an external penetration test Block legacy authentication for Microsoft 365 	<ul style="list-style-type: none"> Restrict end-user privilege Implement host-based security tools 	<ul style="list-style-type: none"> Review and mitigate accounts with SPNs Reduce privileged accounts footprint Perform comprehensive Active Directory security assessment
Detect	<ul style="list-style-type: none"> Implement Network IDS at critical choke points Enable Microsoft 365 to forward alerts to a centralized log management. Build new detection routines to identify password spreads and credential stuffing 	<ul style="list-style-type: none"> Drive hunting exercises for persistence mechanism Ensure an EDR and AV solution is deployed to all servers and endpoints 	<ul style="list-style-type: none"> Document domain based privileged accounts Develop custom alerting for high privileged accounts
Respond	<ul style="list-style-type: none"> Develop user password reset playbook Hire additional Detection and Response personnel 	<ul style="list-style-type: none"> Develop rapid containment playbooks Develop playbooks for endpoint isolation and empower Cyber Defense team with authority to isolate endpoints 	<ul style="list-style-type: none"> Establish playbooks for account disabling or privilege reduction Establish contact points for owners of privileged accounts to validate anomalous behavior

Table 2: Mitigating the Targeted Attack Lifecycle

Internal Recon	Move Laterally	Complete Mission	
<ul style="list-style-type: none"> Restrict availability of IT information internally Develop a program to identify and eliminate credentials for public cloud environment in internal code repositories 	<ul style="list-style-type: none"> Evaluate network segmentation strategy for crown jewel assets Block workstation-to-workstation communication Restrict ability for highly privileged accounts to authenticate to unprivileged systems 	<ul style="list-style-type: none"> Restrict servers initiating outbound connections Implement a data leakage prevention solution 	Inhibit
<ul style="list-style-type: none"> Validate detection capabilities for common AD enumeration tools Ensure logging of command line process creation events 	<ul style="list-style-type: none"> Develop detections based on authentication anomalies Enhance east-west visibility Enable PowerShell, Windows Remote Management, and WMI auditing 	<ul style="list-style-type: none"> Establish alerting for anomalous data transfers Develop detection routines for common ransomware deployment methods 	Detect
<ul style="list-style-type: none"> Establish playbooks for remediating documents revealing sensitive IT information 	<ul style="list-style-type: none"> Develop a playbook to revoke all "active sessions" used by the IAM Role of compromised instance in AWS Develop capability to track internal logins for individual users 	<ul style="list-style-type: none"> Develop and practice ransomware recovery procedures Ensure response playbook account for regulatory and contractual disclosure requirements 	Respond





TARGETED TESTING AND VALIDATION OF CONTROLS AND OPERATIONS

Targeted testing, mission based or objective based testing, and continuous controls validation are an important function of Cyber Defense to prove that security controls are protecting critical assets as expected. Additionally, validation data can be used to prove security investments are providing the expected value. Targeted testing enables reporting on the status of an organization's security posture, if risk is being reduced, and how effective the organization is at identifying compromise. Security validation is best not just as a one point in time report, but rather continuous effort. Continuous and automated controls validation allow security teams to identify gaps, redundancies, areas for improvement and the ability to measure improvements over time.

Security validation provides quantitative data to guide business decision-making, where to invest, and opportunities for cost-savings while maintaining the appropriate level of risk for the business. Validation bridges the gap between controls as designed and inherent risk to actual effectiveness and residual risk. Validation also exercises people and processes with simulated adversary activity. To this point, targeted testing and continuous controls validation has become a business imperative.

Understanding the Attack Surface

In order to properly establish what controls need to be tested and crown jewels assessed, organizations must understand their current attack surface. Traditionally, an organization's understanding of an attack surface would be to gather a list of external-facing assets, such as IP ranges/Netblocks, web applications, VPN servers, and external remote management services. These lists are often manually gathered and exist in spreadsheets.

Additional set of critical assets typically overlooked or left off the list of assets within the attack surface are:

- External-facing database and remote access services
- Developer accounts on sites such as Github or Gitlab
- Staging and QA environments
- External-facing buckets or blob storage within a cloud environment
- Service accounts used for externally facing systems
- More esoteric application software or and network services exposed to the Internet
- Secondary email systems that can be used to deliver payloads w/o content filtering

Discovery of the above assets can occur through understanding the environment, third-party asset management tools, and scripts executed against the organization. Attack surface management does not replace vulnerability management or remove the need for penetration testing but does provide scope and context for those efforts.





The mere-exposure effect creates a cognitive bias that can cause leaders to prioritize controls they are most familiar with over the ones that are most needed. It is vital to use validation techniques that don't reinforce your assumptions, but allow you to make objective, data-driven decisions.

Andrew Roths, Director of Security Engineering at Amazon

Beyond Breach and Attack Simulation

Targeted attack testing goes beyond breach and attack simulation (BAS) to incorporate the latest threat intelligence to perform authentic, active attacks against an organization's defenses. The purpose of targeted attack testing is to evaluate the effectiveness of an organization's current security controls and operations. These technical testing types include penetration testing, red teams, blue teams, and purple teams.

Penetration testing is the systematic testing of defenses and critical assets to pinpoint and reduce vulnerabilities and security system misconfiguration. Penetration testers use real-world attacker TTPs against systems, applications, embedded devices, industrial control systems, and even against people using social engineering. The purpose of penetration testing is to determine if critical assets are at risk and to identify complex security vulnerabilities.



Our adversary simulation tests are almost always successful in identifying vulnerabilities and gaps in security configurations. Finding these risks is not a bad thing, doing nothing about it is. It is critical to use the latest intelligence when performing these tests to ensure we find the gaps before the attackers do and that is key to a security organization's value.

Evan Peña, Director of Proactive Services, Mandiant

Red Teams test security effectiveness to gain an understanding of where an organization's weaknesses exist. Red teams provide an objective based approach to testing by leveraging current attack techniques to accomplish a specific mission. These activities use highly skilled practitioners attempting to complete the objective while avoiding detection. This provides the organization with an excellent prospective on what an attack might be able to achieve. The usefulness of Red Teams relies on the skillfulness of their methods and the currency of the intelligence on active attacker techniques. By utilizing highly skilled Red Teams to perform unannounced exercises, the organization can identify gaps in team member skill sets, Cyber Defense processes, and toolsets.

Blue Teams attempt to detect and prevent the actions of a Red Team and when they are unsuccessful in doing so, take the data provided by Red Teams and remediate where needed to optimize security effectiveness. The Blue Team relies on the Red Team's findings to tune controls and address gaps and vulnerabilities. Red and Blue Teams typically perform their functions in an asymmetric mode of operation.

Purple Teams bring Red and Blue Teams together to work in a more symbiotic fashion. They often leverage automation of validation tests and integrated threat intelligence. This lets Red Teamers test controls with multiple step-by-step scenarios to demonstrate how the security technologies and the Blue Team perform against the threats most likely targeting the organization. For Blue Teams, the automation delivers prescriptive analytics that allow metrics showing improvement in the effectiveness of controls and operations over time while still having meaningful Red Team curated tests executed.

Targeted attack testing and continuous security validation should be performed to:

- **Prioritize the most relevant threats by leveraging threat intelligence**

Prioritizing how and what to test requires active adversary intelligence about what threats are most relevant to the company. Security teams should not make threat intelligence analysis retrospectively, but rather utilize current data that informs what attackers are likely to do next, who they will target and what methods they may use. As a first step in the validation process, threat intelligence can identify the threats that matter and drive a validation strategy. This insight enables security teams to execute relevant validation content and attacker TTPs to challenge security controls.

- **Measure effectiveness of controls against known adversaries**

Assessing how the security stack performs against those most relevant attacks requires testing across the full attack. Active attacks executed safely and authentically in the environment are necessary to measure the true effectiveness of security controls. This includes evaluating how people, processes, and technologies work together against both adversary techniques and technical attacks.

- **Improve effectiveness through optimization and repeat assessment**

Optimizing controls based on the gaps and shortcomings that are revealed in the measurement stage is an ongoing process. Once controls are optimized, continuous testing will maintain a good baseline, enable measurement of improvements, and deliver quantitative data to demonstrate the value of security to the business.

- **Rationalize security investments with continuous security validation**

Security teams can capture data required to prove effectiveness of security to support rationalization of security investments. Additionally, the use of security validation can provide insight into the impact of a change or removal of a control within the security infrastructure and in the context of a company's risk tolerance. Once controls are optimized, security leaders can use validation data to continuously measure and demonstrate an improvement to the security program and investments. Equally important, companies can pinpoint where overlaps exist and find ways to cut costs without impacting risk.

- **Continuously monitor for environmental drift**

Changes naturally occur in the IT environment which may affect security effectiveness. To ensure cyber defenses are not weakened, it's critical to continually monitor, detect and alert on drift to accurately measure effectiveness. The completion of steps one through four gives security teams a baseline by which to conduct ongoing testing to ensure optimal effectiveness as these changes occur.

Misconfiguration and environmental drift are the silent killers of cyber security effectiveness.

Chris Key, Chief Product Officer, Mandiant

Validating the Effectiveness of Your Staff

It is equally important to validate your security team's technical capabilities, processes, and procedures to respond to attacks. Virtual environments and cyber ranges can be used to evaluate your staff and have them practice responding to real-world threats—without real-world consequences.

Security teams use virtualized environment that simulate typical IT infrastructure such as network segments, workstations, servers and applications.

These exercises are useful in the following ways:

- Identify areas for team improvement. Investigate real-world incidents to identify gaps in training, processes, procedures and communication plans.
- Investigate critical security incidents. Test your response and intelligence teams with the latest attack scenarios and attacker TTPs.
- Research and analyze identified threats. Learn to research attacker TTPs and identify IOCs from host- and network-based artifacts.

Cyber range exercises should cover various attack scenarios including ransomware, insider threats, data exfiltration, Active Directory attacks and lateral movement.

Validating the Design of Incident Response and Remediation Plans

It is important to have an incident response and remediation plan in place before it is needed. Tabletop exercises are a simple yet effective way to test the design of IR processes for multiple threat scenarios. By walking through a cyber incident and the associated IR processes, teams can identify gaps in investigative processes, available tools, communication processes, and stakeholders involved. Identifying and addressing these gaps before an incident occurs saves valuable time and prevents many mistakes. In addition to leveraging tabletop exercises for IR processes, it is recommended to leverage tabletop exercises for common remediation efforts such as large-scale password resets and restoration processes.

Mandiant experts recommend executive tabletop exercises be performed twice a year and technical tabletop exercises be performed quarterly.

Tabletop exercises help identify gaps in response and remediation plans and provide an opportunity to update plans based on the latest threats. Tabletop exercises should answer at a minimum:

- Who is part of the incident response and remediation team?
- Who will have the ultimate decision-making authority for the organization during an incident including business impacting decisions like—disconnecting from the Internet, conducting enterprise-wide password resets, payment or non-payment of ransom and public disclosure timing?
- Which partners will be brought in and when? Is there a service level agreement (SLA) in place?
- When should legal counsel be contacted?
- When should insurance provider be contacted?
- Who will handle crisis management and communications?
- How will the organization handle remediation efforts (e.g., Can a full password reset event be performed and by whom)?
- Is there a ransomware broker in place in the event ransom must be negotiated?
- Are there regulations around not paying ransom to identified groups?

Incident response, like all high stress, high risk activities, lowers an organization's performance to muscle memory. The adversaries rehearse every day. If you haven't rehearsed, they have.

Chris Calvert, Cyber Defense Entrepreneur



ACTIVATING CYBER DEFENSE

Most organizations do not have the resources nor the appetite to establish all the functions of Cyber Defense in-house. Instead, they leverage SaaS offerings, microservices, and partner with service providers to activate their cyber defenses.

Stakeholder Buy-in

It is important that organizations' Cyber Defense programs are armed with diverse tools and operations to effectively protect themselves against a wide variety of sophisticated attacks. Historically, acquiring funding to support these programs has been challenging. Achieving buy-in from key stakeholders for accelerating capabilities, such as a Board of Directors or executive leadership, is most effective when the needs for such programs are clearly tied to their impact on the business growth.

News stories focus on large attacks which confuse boards about where action is needed. It is the role of the Cyber Defense organization to help the board understand threats specific to the organization to gain buy-in for their strategic investment decisions.

Dawn Marie-Hutchinson, Chief Information Security Officer

Staffing Considerations

A well-known challenge in the cyber security industry is the shortage of skilled analysts. Automation of defenses can be implemented to relieve some of the workload and reduce the burnout that comes from sifting through mountains of data for many hours, day after day. Even with added automation providing more consistent analysis, the industry still relies heavily on analysts to collect and interpret data to piece together the big picture of an attack.

Security training for staff is one of the best investments an organization can make. Training increases employee satisfaction and matures their skill sets so they can provide a higher level of expertise back to the company. Training programs with development paths or certification programs offer better return on investment over piecemeal courses.

Organizations can also invest in automation to help repurpose traditionally lower tier staff roles by refocusing those resources away from overhead, “white noise” tasks. By having a program automate basic functions, analysts can stay focused on more critical work while also producing results at a faster pace. There is also opportunity to leverage automation for higher fidelity alerts, leaving the more difficult analysis to the most skilled at separating the signal from the noise.

Leveraging Accelerators

Another challenge to consider is the ever-persistent nature of business and day-to-day operations diverting resources away from being able to advance or mature an organization’s Cyber Defense program. Unfortunately, this can quickly lead to elements being neglected due to other priorities.

To address this, one approach is to apply accelerators or microservices to help bypass many traditional hurdles thus relying on external resources to address Cyber Defense components that an organization does not have the bandwidth to address on their own. By utilizing partner resources to analyze and provide objective solutions for identified issues, organizations get an unbiased perspective of improvement needs and on-demand, and often highly specialized, skills to fix the issues without having to hire new resources or maintain less-frequently called upon skill sets. Existing staff can also learn from experts and still complete their work.



Intelligence allows us to show management that the money spent protecting our business, our image, our reputation and the personal information of our customer is absolutely worth it. It has clearly shown us that we are a target, that we are being attacked daily, that we could never manage all of our cyber defenses in-house and that we had no idea before the availability of this intelligence how bad things really are.

Gary Winder, IT Network Engineer, Baptistcare

Engaging Managed Services

Another approach for accelerating Cyber Defense capabilities with a limited budget is to engage with a managed service provider. Managed services allow an organization to outsource a portion of Cyber Defense functions such as detection and response, hunting, or validation. The services provide confidence in 24x7 protection while benefiting from the providers intelligence gleaned from other customers and attacker visibility.

Managed services offer the additional benefit of intelligence gained from broad exposure to attacks. This exposure allows analysts within the managed service to gain experience responding to a wide range of incidents and activities. A managed service can observe campaigns as they unfold across their client-base and adjust response actions accordingly. This front-line experience is highly beneficial as it allows analysts to develop their skills and be able to respond quickly to events, reducing the impact on their clients' environments.

As a security team of one, it is impossible to keep up with the current volume of alerts. Partnering with a service provider to monitor threats is the only way to have confidence in our ability to detect compromise.

Andi Hill, System Administrator, B&M Roofing

Flexible Consumption Models

Organizations can supplement staff, outsource services with managed services or utilize microservices to achieve necessary CDO functions and capabilities. These services are typically purchased as six- or 12-month subscriptions. Many service providers also provide flexible spending options that allow organizations to make a single purchase of credits to be used over a twelve-month period. This flexible option is a good choice when organizations know they will have projects throughout the year but don't have a timeline for execution. This is also a common option utilized for training needs and for on-demand access to experts.

Example

The Mandiant Managed Defense organization received information about a zero-day vulnerability in a widely used product that was being exploited to deploy ransomware. The managed service provider initiated a threat hunting campaign to identify evidence of attacker activity across the entire customer base. Additional intelligence led the managed detection and response services team to begin scoping customer environments for hosts running the vulnerable software. Affected customers were quickly identified and advised to contain certain on-premises systems. Protections were put in place before the ransomware could be deployed. In this case, all customer of managed defense services and other SaaS offerings benefited from the adversary IOCs provided from intelligence gathered across the customer base.

Conclusion

In today's threat landscape, trying to gain an advantage in cyber security is not an easy task. In order to do so, defenders must work hard to be one step ahead of the attacker at every move. After all, attackers only need to be right once in order to make a defender the latest victim. Organizations can shift the odds of becoming a victim in their favor by activating the functions of Cyber Defense. This does not only refer to establishing the functions but operationalizing them to stand up against adversaries. Organizations need to focus on the areas that matter most and employing the accelerators described within this book to gain **The Defender's Advantage**.



Average and Median Ransom
Payments in Q1 2021*

Average:

\$220,298 USD

Median:

\$78,398 USD

APPENDIX A

Multifaceted Extortion

What is Multifaceted Extortion?

Multifaceted extortion combines traditional ransomware and other extortion tactics to coerce victims to comply with hefty demands. The nature of multifaceted extortion means that standard basic disaster recovery procedures used during a ransomware attack are no longer an adequate recovery strategy.

The first known ransomware was documented in 1989. The ransomware hid directories and encrypted file names on a victim's computer. Users had to pay \$189 to gain access to their files. Since then, attackers have matured their technology and tradecraft to demand sums up to \$50M. Today, ransomware spreads quickly through environments and encrypts entire drives, crippling business operations.

*<https://www.coveware.com/blog/ransomware-attack-vectors-shift-as-new-software-vulnerability-exploits-abound>



Financially motivated threat actors such as FIN11 employ ransomware-as-a-service to carry out their attacks. They outsource code development eliminating the need to maintain that expertise themselves. To maintain anonymity, attackers now demand payment in cryptocurrencies such as bitcoin, making it increasingly difficult to track and locate them.


Threat actors have realized they can demand higher ransoms by targeting larger organizations and applying pressure with additional coercion techniques.

Tactics that support multifaceted extortion include:



Impaired File Availability

Ransomware typically encrypts a target organization's sensitive files, making them unavailable to legitimate users. This can be combatted with best practices and disaster recovery planning.



Threats to Publish Data

Theft of sensitive data is followed by threats to publish the stolen data if the payment demands are not met. This form of extortion is more consequential because data breaches often carry more serious business consequences than service disruptions. According to the Mandiant M-Trends 2021 report¹, "A data breach can result in greater reputational damage, regulatory fines, class action lawsuits, and derailed digital transformation initiatives. These consequences were not typically seen with ransomware before 2019."



Name-and-Shame

Attackers will post parts of the stolen data on name-and-shame websites to prove they possess the stolen data. The attackers then engage with media organizations to inflict brand damage, further coercing victims into paying a ransom. Some attackers have even notified business partners of data theft, creating friction in third-party relationships and prompting breach disclosures.

Well-established Cyber Defenses help prevent the intrusions that precede ransomware deployment, hunt for active compromise, and respond to successful attacks.



APPENDIX B

Investigative Theory

The following can be used to steer an investigation and justify a course of action.

Ask Questions

Every investigation should be processed via questions. Certain questions are always relevant, while others will need to be developed based on the incident categorization or other details. Write down the questions within investigative notes. This helps to focus analytical work and realign the investigation after following leads.

Standard Questions

- What was the point of entry?
- Has malicious code executed?
- Has persistence been established?
- Were credentials compromised?

Investigation Specific Sample Questions

- What systems are accessible within this network or by a specific compromised host?
- What external IP addresses did the compromised asset communicate with?
- What is the parent process hierarchy for this malicious process?
- What scheduled tasks are normal within the environment, and which of these tasks are not part of the baseline?

Postulate on Attacker Intent

Based on training and experience, anticipate the attacker's intent. Use this anticipation to frame the attacker's actions. An attacker who wants to redirect customer payment data may inject malicious JavaScript into a web page, but never attempt to compromise subsequent machines. Conversely, an attacker wanting to deploy ransomware is likely to seek out administrative credentials across multiple machines to widely deploy their code.

As new evidence is collected, re-evaluate the assumed intent.

Continually ask yourself:

- Does the evidence support this theory?
- What evidence would disprove this theory? If it exists, where can it be found?
- What other competing theories does this evidence support?
- What is a non-malicious story that leads to this same behavior?

Identify Evidence

Use analysis and forensics to identify events of interest within the network and affected systems. The changes, behaviors, and observations should be paired together with an interpretation. In other words, the report should contain a description of the artifact along with a description of what it represents.

Example

Example: 22-3-2021 18:00:23Z - Artifact: The EXAMPLE malware altered HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Run adding an entry that references the malware.

Interpretation: This instantiates a backdoor by auto executing the listed program upon restart.



Organize Into a Story

As events of interest are identified and attributed to the attacker, capture the information within a unified timeline. Ultimately, the goal of the incident report is to reveal what the attacker accomplished or attempted. The report should read from beginning to end based on the attacker's timeline, not the analyst timeline. Do not start with the alert, start with the earliest observed evidence of attacker activity.

Leverage the attack lifecycle or other kill chain methodology, to organize the evidence into sequences that explain an attacker's overall intent or minor objectives.

Regularly ask:

- Does the story make sense?
- Why would an attacker do this?
- What sensors and logs are relevant?

Iterate

Repeat this process until there are diminishing returns on the output.





APPENDIX C

SOC Ransomware Procedure

If controls and other measures fail to prevent a ransomware attack, the SOC will work through the following high-level process to ensure an appropriate response:

Identify Ransomware Family/Type

To best combat a ransomware attack in an environment, the SOC understand the ransomware family and type of ransomware that they are dealing with. The SOC leverages all information and tools at their disposal. Such information may include:

- Host-Based Indicators (HBIs)
- Network-Based Indicators (NBIs)
- Threat Intelligence

This information can be derived from several sources from threat intelligence to security controls.

Isolate/Contain Affected Systems

Containment becomes the top priority as ransomware can spread quickly throughout an environment. To bolster the chances of a successful containment, the following actions should be taken:

1. Segmentation through the blocking of common ports and protocols that are used to spread ransomware is the critical first step. At a minimum, blocking should be applied where possible, to:
 - Server Message Block (SMB) – TCP/445, TCP/135
 - Remote Desktop Protocol (RDP) – TCP/3389
 - Windows Management Instrumentation (WMI) & Remote Power Shell – TCP/80, TCP/5985, TCP/5986
2. Restrict inbound connections on endpoints through the use of host-based firewall technologies (e.g., Windows Firewall via Group Policy).
3. Utilize network segmentation, where possible, to enforce logical isolation of endpoints. Examples of this may include applying ACLs to physical and virtual interfaces, firewalls configured with stateful ACLs, or applying routing configuration changes that null route traffic if specific pre-defined parameters are met.

Mitigate Lateral Movement/Propagation

In conjunction with containment and isolation, the SOC's next task is to mitigate common lateral movement techniques that ransomware attacks rely upon:

- Leverage UAC Token Filtering or `lor SID S-1-5-114 NT AUTHORITY\ Local Account` and setting the appropriate computer policies for user rights assignments through GPO.
- Quickly define a strategy to enforce password randomization on the built-in Local Administrator account. One such example is the deployment and use of Microsoft LAPS (Microsoft Local Administrator Password Solution).
- Restrict access to privileged accounts that may be used for propagation and lateral movement. Additionally, actions may be taken to restrict access to credentials or tokens in memory to minimize the exposure of credentials that may be used for movement.
- Disable Administrative and Hidden Shares to prevent the ransomware from binding to additional endpoints. Examples of these shares may include `ADMIN$`, `IPC$`, and `CS`

Protect & Secure Backups

While backups may play a crucial role in the remediation of ransomware attacks, they are certainly not immune to its reach. Typical modern threat actors and ransomware families are able to target a client's backups and backup environment by attempting to encrypt, delete or stop services associated with backups.

It is also important to ensure that key enterprise and network services, once thought unnecessary to backup given either time to rebuild or high levels of redundancy, are backed up as well. An example of this is Active Directory. Some organizations take the native replication/redundancy built into the solution as a reason to not have a mature or any real backup strategy. However, if those have been rendered unavailable by intentional reams, restoration from backup might be the only palatable recovery option.

Care must be taken to isolate the backup environment from access by endpoints and to secure known-good/clean backup copies, until the ransomware threat is remediated. A starting point for this is the 3-2-1 rule for backups. This rule requires three copies of the data, two different media types, and one offsite copy with strong consideration given to one of the copies, especially for critical assets, being an immutable source. It is also critical that backups are regularly tested and with different content/date sets.

Disable/Stop Maintenance Tasks & Services

Once the SOC has determined the ransomware family and worked with the client to assist in containment of the ransomware, the SOC will provide a list of tasks or services that should be stopped or disabled temporarily to prevent further interference from the ransomware.

Examples may include:

- Backup services
- Scheduled tasks
- Software deployment platforms (SCCM, Ghost)
- Administrative scripts or batch files

Note: The steps above are a guide and not an all-inclusive process.

Contributors

Kerry Matre

Alex Flores

Alexa Rzasa

Daniel Nutting

David Lindquist

Eric Scales

Evan Peña

Jay Christiansen

Jeff Compton

Jennifer Guzzetta

Jim Meyer

John DeLozier

John Doyle

Kelly Gordon

Kyle Baird

Lynn Harrington

Michael Reynolds

Nader Zaveri

Nicholas Slaughter

Nick Bartosch

Nick Bennett

Nick Pelletier

Omar Toor

Shayn Ronis

Trisha Alexander

Editor

Stephanie Wisdom

Design & Composition

Eclipse

The Defender's Advantage: A Guide to Activating Cyber Defense

© October 2021

Disclaimer

The information in this book is written as a general guide only. It should not be relied upon as a substitute for specific professional advice. Professional advice should always be sought before taking any action based on the information provided. Every effort has been made to ensure that the information in this guide is correct at the time of publication. The views expressed in this guide are those of the authors. The publishers and authors do not accept responsibility for any errors or omissions contained herein. It is your responsibility to verify any information contained in the guide before relying upon it.



We are facing off against adversaries in our own environments. This provides an advantage arising from the fact that we have control of the landscape that is under attack. Security organizations struggle to capitalize on this advantage. As security organizations, we must activate our cyber defenses, advancing capabilities from a prepared state to active duty. This activation is guided by Intelligence and orchestrated through the other Cyber Defense functions: Command and Control, Hunt, Detect, Respond and Validate. It is through this activation that we can take control and galvanize our defender's advantage.

About Mandiant

Effective security is based on the right combination of expertise, intelligence and technology. Since 2004, Mandiant has been a trusted security leader to organizations that can't afford to fail. Today Mandiant delivers decades of frontline insights at scale through easy-to-deploy and consume SaaS solutions for organizations of all sizes. Offerings span the range of consulting, automated defense, managed detection and response, threat intelligence and security validation for provable and transformative cyber defense.

www.mandiant.com