O'REILLY®

2nd Edition

Intelligence-Driven Incident Response

Outwitting the Adversary

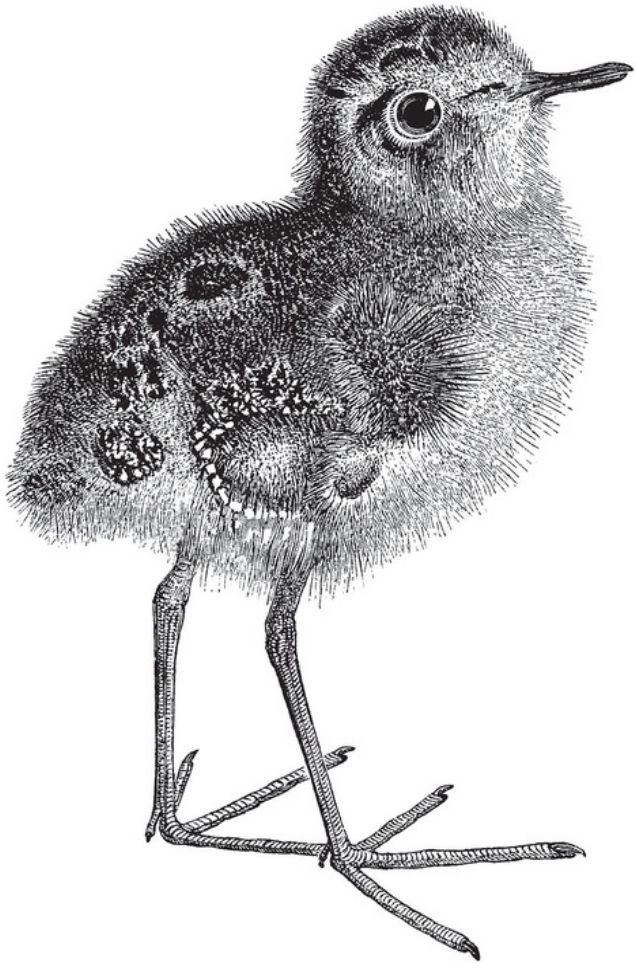

Early Release

RAW & UNEDITED

Rebekah Brown &
Scott J. Roberts

Intelligence-Driven Incident Response

SECOND EDITION

Outwitting the Adversary

Rebekah Brown and Scott J. Roberts

Intelligence-Driven Incident Response

by Scott J. Roberts and Rebekah Brown

Printed in the United States of America.

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

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

Acquisitions Editor: Jennifer Pollock

Development Editor: Angela Rufino

Production Editor: Elizabeth Faerm

Copyeditor: TO COME

Proofreader: TO COME

Indexer: TO COME

Interior Designer: David Futato

Cover Designer: Karen Montgomery

Illustrator: Kate Dullea

May 2023: Second Edition

Revision History for the Early Release

- 2022-02-25: First Release

See http://oreilly.com/catalog/errata.csp?isbn=9781098120689 for release details.

978-1-098-12062-7

[LSI]

Chapter 1. Introduction

A NOTE FOR EARLY RELEASE READERS

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

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

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

But I think the real tension lies in the relationship between what you might call the pursuer and his quarry, whether it's the writer or the spy.

—John le Carre

Once relegated to the secretive realms of national security and military operations, intelligence has become something that is fundamental to the daily functioning of many organizations around the world. At its core, intelligence seeks to give decision makers the information that they need to make the right choice in any given situation.

Previously, decision makers experienced significant uncertainty because they did not have enough information to make the right decisions. Today they are just as likely to feel like there is too much information, but just as much ambiguity and uncertainty. This is especially the case with network security, where there are fewer traditional indications that a significant action is actually about to take place. In order to make decisions about how to prepare for and respond to a network security incident, decision makers need analysts who understand intelligence fundamentals, the nuance of

network intrusions, and how to combine the two into an accurate assessment of a situation and what it means for their entire organization. In short, they need analysts who can conduct intelligence-driven incident response.

Before diving into the application of intelligence-driven incident response, it is important to understand the evolution of cyber security incidents and their responses, and why it is so relevant in this field. This chapter covers the basics of cyber threat intelligence, including its history, recent activity, and the way forward, and sets the stage for the concepts discussed in the rest of this book.

Intelligence as Part of Incident Response

As long as there has been conflict, there have been those who watched, analyzed, and reported observations about the enemy. Wars have been won and lost based on an ability to understand the way the enemy thinks and operates, to comprehend their motivations and identify their tactics, and to make decisions—large and small—based on this understanding. Regardless of the type of conflict, whether a war between nations or a stealthy intrusion against a sensitive network, intelligence guides both sides. The side that masters the art and science of intelligence, analyzing information about the intent, capability, and opportunities of adversaries, and is able to act on that information, will almost always be the side that wins.

History of Cyber Threat Intelligence

One of the best ways to understand the role of intelligence in incident response is by studying the history of the field. Each of the events listed below could (and often do!) fill entire books. From the iconic book *The Cuckoo's Egg* to recent revelations in decades old intrusions such as Moonlight Maze, the history of cyber threat intelligence is intriguing and engaging, and offers many lessons for those working in the field today.

The First Intrusion

In 1986, Cliff Stoll was a PhD student managing the computer lab at Lawrence Berkeley National Laboratory in California when he noticed a billing discrepancy in the amount of 75 cents, indicating that someone was using the laboratory's computer systems without paying for it. Our modern-day network security-focused brains see this and scream, "Unauthorized access!" but in 1986 few administrators would have jumped to that conclusion. Network intrusions were not something that made the news daily, with claims of millions or even billions of dollars stolen; most computers connected to the "internet" belonged to government and research institutes, not casual users, and it was easy to assume everyone using the system was friendly. The network defense staple tool tcpdump was a year from being started. Common network discovery tools such as Nmap would

not be created for another decade, and exploitation frameworks such as Metasploit would not appear for another 15 years. The discrepancy was more easily expected to be a software bug or bookkeeping error as it was that someone had simply not paid for their time.

Except that it wasn't. As Stoll would discover, he was not dealing with a computer glitch or a cheap mooch of a user. He was stalking a "wily hacker" who was using Berkeley's network as a jumping-off point to gain access to sensitive government computers, such as the White Sands Missile Range and the National Security Agency (NSA). Stoll monitored incoming network traffic with printers writing reams of packets onto paper to keep a record and began to profile the intruder responsible for the first documented case of cyber espionage. He learned the typical hours the attacker was active, monitored the commands he ran to move through the interconnected networks, and observed other patterns of activity. He discovered how the attacker was able to gain access to Berkeley's network in the first place by exploiting a vulnerability in the *movemail* function in GNU Emacs, a tactic that Stoll likened to a cuckoo bird leaving its egg in another bird's nest to hatch and inspiring the name of his book on the intrusion, *The Cuckoo's Egg*.

Understanding the attacker meant that it was possible to protect the network from further exploitation, identify where he may target next, and allowed a response, both on the micro level (identifying the individual carrying out the attacks) and on the macro level (realizing that nations were employing new tactics in their traditional intelligence-gathering arsenal and changing policies to respond to this change). Sharing this understanding ended up being key not just to protecting Lawrence Berkeley National Lab, but many other government organizations.

Destructive Attacks

In 1988 Cornell student Robert T. Morris hacked into a computer lab at the Massachusetts Institute of Technology (MIT) and released a computer program that was designed to replicate itself to as many computers as possible without being detected. It did this by exploiting a backdoor in the

Internet's email delivery system as well as a flaw in the *finger* program that identified network users. It may have just been a harmless experiment - or a prank, as some have described it - that few people ever knew about, however it did not work exactly as Morris intended and became part of cyber history (which is just like regular history, but cooler). The worm ended up crashing 6000 computers, which was about 10% of the internet of the time, with computers at Harvard, Princeton, Stanford, Johns Hopkins, NASA, and the Lawrence Livermore National Laboratory among the many victims. Investigators from the FBI were struggling to inspect the activity, unsure if it was an outage or an attack, when Morris called two friends and admitted that he was the one behind the worm. One of the friends called *The New York Times,* which led to the eventual identification and conviction of Morris for violations of the new Computer Fraud and Abuse Act of 1986 (CFAA).

While there was no nefarious intent behind this worm- it was really just another example of how programs don't always behave exactly the way that their creator intended them to - there were far reaching implications that can still be seen today. As the internet became more and more critical to operations, it became even more important to be able to quickly identify and remediate other intrusions or destructive attacks. The Computer Emergency Response Team (CERT) was established at Carnegie Mellon University as a professional, trained response team responsible for providing assessments and solutions for cyber attacks. It also highlights why it is sometimes important to be able to attribute an attack - in 1988 there had been significant "warming" of the Cold War, with Reagan and Gorbachev meeting in Moscow and Soviet Troops beginning to withdraw from Afghanistan. If this worm had inaccurately been blamed on Soviet activity - which was actually the case for the activity from the Cuckoo's Egg - it could have significantly changed the course of history.

Moonlight Maze

In the decade following 1986's Cuckoo's Egg and 1988's Morris Worm, the field of Incident Response improved, not just with the creation of CERT, but because of the professionalization of the field itself across the

government, military, and private sector. In addition, the emergence of proper network monitoring tools meant that defenders weren't reliant on printers scattered around a basement to identify malicious network activity. This increase in capabilities was fortuitous, as intrusions not only continued but grew in scope and sophistication. In 1998, the US government identified what is believed to still be the largest and longest running intrusion into government networks - codenamed Moonlight Maze.

In March of 1998, the US government noticed anomalous activity within several sensitive and restricted networks, including the Pentagon, NASA (yes, NASA has apparently always been pretty high on the priority list for intrusions), and the Department of Energy (DOE). Further analysis identified the same malicious activity at several universities and identified that the activity had been ongoing for at least two years. The sustained nature of this activity was unlike any of the previous intrusions (as far as was known at the time) which had seemed more targeted and short-lived. Unlike those instances, the adversaries had left strategic backdoors in different parts of the network so that they could return at will. Information was being gathered from numerous locations that often seemed unrelated.

We have been fortunate enough to work closely with people who directly supported the investigation and response to Moonlight Maze, both in the 1990s and today, as there are still many unknowns and many insights to be found from this intrusion. Talking with these individuals and reading through the numerous reports on the intrusion hammers home the point that intelligence work is critical to incident response, and that cyber threat intelligence in particular bridges the gap between what is happening on the strategic level with national interests and foreign adversaries, and how those adversarial goals and actions show up on a computer network. The US has had a full scope intelligence community actively looking and listening for signs of foreign interference, but the largest network attack went unnoticed until it was detected "by accident" because intelligence work had not yet been fully modified to account for actions taken against a network.

Moonlight Maze kicked cyber threat intelligence capabilities into the modern era. Computer networks were not something that might be impacted

from time to time, these networks were now being targeted directly for the information and access they held. Computer networks were part of intelligence collection, and intelligence needed to play a role in their defense.

Modern Cyber-Threat Intelligence

Over the decades, the threats have grown and morphed. Adversaries are not just foreign governments or curious students. Organized criminals, identity thieves, scammers, ideologically motivated activists, and others with various motivations all realized the impact that their activities could have when directed at digital targets instead of physical ones. These adversaries use an ever-expanding set of tools and tactics to attack their victims and actively attempt to evade detection. At the same time our reliance on our networks has increased, making incidents even more impactful. Understanding the attacker has gotten much more complicated, but no less important.

Understanding how to identify the attacker activity and how to use that information to protect networks is the fundamental concept behind a more recent addition to the incident responder's toolkit: cyber-threat intelligence. *Threat intelligence* is the analysis of adversaries—their capabilities, motivations, and goals, and cyber-*threat intelligence* (CTI) is the analysis of how adversaries use the cyber domain to accomplish their goals. See #fig0101 on how these levels of attacks play into one another.

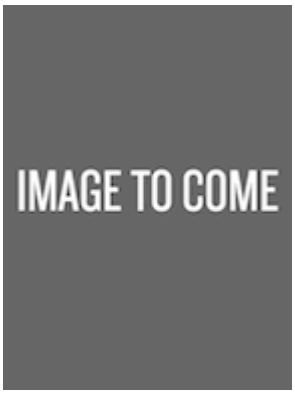

Figure 1-1. From intelligence to cyber-threat intelligence

Initially, intelligence analysts came into the picture after an intrusion like Moonlight Maze to understand what the overall intrusion told us about the adversary. What were their goals, their motivations, their capabilities, their organizational structure - all things that were important for a strategic understanding and long-term planning, but not things that would immediately provide value to those trying to defend their networks from the attacks of today, tomorrow, or sometimes even last week. Cyber-threat intelligence began to focus more on tactical and technical details that were more immediately actionable, learning along the way what types of information were most valuable in different situations. Cyber-threat intelligence analysts didn't just bring data, they also brought insights.

In information security, we traditionally focus on observable concepts; we like things that are testable and reproducible. Cyber threat intelligence, meanwhile, lives in the area between observations and interpretation. We may not know ***for sure*** that an adversary will attempt to access employee financial records, but we can conduct analysis on data about past intrusions and successful attackers outside of our network and make recommendations on the systems and types of data that may need additional protection. Not only do we need to be able to interpret information, but we also need to be able to convey it in a way that is meaningful to the intended audience to help them make decisions. Looking back on his historic analysis of the intrusions in the Cuckoo's Egg, Cliff Stoll identified the need for a story as one of his key takeaways from the entire experience. "I thought I could just show people the data and they would understand," he said "but I was wrong. You need to tell a story." (SANS CTI Summit 2017)

The Way Forward

New technologies give us more information about the actions that attackers take as well as additional ways to act on that information. However, we have found that with each new technology or concept implemented, the adversary adapted; worms and viruses with an alphabet soup of names changed faster than our appliances could identify them, and sophisticated, well-funded attackers were often more organized and motivated than many

network defenders. Ad-hoc and intuitive intelligence work would no longer suffice to keep defenders ahead of the threat. Analysis would need to evolve as well and become formal and structured. The scope would have to expand, and the goals would have to become more ambitious.

In, addition to detecting threats against an organization's often nebulous and ephemeral perimeter, analysts would need to look deeper within their networks for the attacks that got through the lines, down to individual user systems and servers themselves, as well as look outward into third-party services and to better understand the attackers who may be targeting them. The information would need to be analyzed and its implications understood, and then action would have to be taken to better prevent, detect, and eradicate threats. The actions taken to better understand adversaries would need to become a formal process and a critical part of information security operations: threat intelligence.

Incident Response as a Part of Intelligence

Intelligence is often defined as data that has been refined and analyzed to enable stakeholders to make better decisions. Intelligence, therefore, requires data. In intelligence-driven incident response, there are multiple ways to gather information that can be analyzed and used to support incident response. However, it is important to note that incident response will also generate cyber-threat intelligence. The traditional intelligence cycle—which we cover in depth in #basics_of_intelligence—involves direction, collection, processing, analysis, dissemination, and feedback. Intelligence-driven incident response involves all of these components and helps inform direction, collection, and analysis in other applications of threat intelligence as well, such as network defense, secure software development, and user awareness training. Intelligence-driven incident response doesn't end when the intrusion is understood and remediated; it generates information that will continue to feed the intelligence cycle.

Analysis of an intrusion, no matter if it was successful or failed, can provide a variety of information that can be used to better understand the overall threat to an environment. The root cause of the intrusion and the initial access vector can be analyzed to inform an organization of weaknesses in network defenses or of policies that attackers may be abusing. The malware that is identified on a system can help identify the tactics that attackers are using to evade traditional security measures such as antivirus or host-based intrusion-detection tools and the capabilities they have available to them. The way an attacker moves laterally through a network can be analyzed and used to create new ways to monitor for attacker activity in the network. The final actions that an attacker performed (such as stealing information or changing how systems function), whether they were successful or not, can help analysts understand the enemy's motivations and goals, which can be used to guide overall security efforts. There is essentially no part of an incident-response engagement that cannot be used to better understand the threats facing an organization and improve future defense and response.

For this reason, the various processes and cycles outlined in this book are aimed at ensuring that intelligence-driven incident response supports

overall intelligence operations. Although they provide specific guidance for utilizing cyber-threat intelligence in incident response, keep in mind that wider applications can be used as intelligence capabilities expand.

What Is Intelligence -Driven Incident Response?

Cyber-threat intelligence isn't a new concept, simply a new name for an old approach: applying a structured analytical process to understand an attack and the adversary behind it. The application of threat intelligence to network security is more recent, but the basics haven't changed. Cyber-threat intelligence involves applying intelligence processes and concepts—some of the oldest concepts that exist—and making them a part of the overall information security process. Threat intelligence has many applications, but one of the fundamental ways it can be utilized is as an integral part of the intrusion-detection and incident-response process. We call this *intelligence-driven incident response* and think it is something every security team can do, with or without a major capital investment. It's less about tools, although they certainly help sometimes, and more about a shift in the way we approach the incident-response process. Intelligence-driven incident response will help not only to identify, understand, and eradicate threats within a network, but also to strengthen the entire information security process to improve those responses in the future.

Why Intelligence -Driven Incident Response?

Over the past few decades, our world has become increasingly interconnected, both literally and figuratively, allowing attackers to carry out complex campaigns and intrusions against multiple organizations with the same effort that it used to take to target a single entity. We are long past the point where we can automatically assume that an intrusion is an isolated incident - in fact, while we used to be stunned to find overlaps and connections between intrusions, now we are suspicious when we *don't* see

overlap. When we better understand the adversary, we can more easily pick up on the patterns that show commonalities between intrusions. Intelligence-driven incident response ensures that we are gathering, analyzing, and sharing intelligence in a way that will help us identify and respond to these patterns more quickly.

Operation SMN

A good example of this is the analysis of the Axiom Group, which was identified and released as a part of a Coordinated Malware Eradication (CME) campaign in 2014 called Operation SMN.

WHAT'S IN A NAME?

The SMN in Operation SMN stands for some marketing name, a not-so-subtle but amusing jab indicating how widespread the belief is that marketing often takes over intelligence products. For better or worse, threat intelligence has been eagerly embraced by marketing forces all touting the best threat-intelligence products, feeds, and tools. The first time many people are exposed to threat intelligence is through marketing material, making it difficult for many to fully understand what threat intelligence actually is.

It is important that intelligence work is done with the end goal of better understanding and defending against adversaries. Sometimes marketing gets in the way of that, but ideally marketing can help with messaging and ensuring that the "story" behind threat intelligence reaches the right audience in the right way.

For more than six years, a group of attackers known as the Axiom Group stealthily targeted, infiltrated, and stole information from Fortune 500 companies, journalists, nongovernmental organizations, and a variety of other organizations. The group used sophisticated tools, and the attackers went to great lengths to maintain and expand access within the victims' networks. As malware was detected and the incident-response process

began within various victim organizations, coordinated research on one of the malware families used by this group identified that the issue was far more complex than originally thought. As more industry partners became involved and exchanged information, patterns began to emerge that showed not just malware behavior, but the behaviors of a threat actor group working with clear guidance. Strategic intelligence was identified, including regions and industries targeted.

This was an excellent example of the intelligence cycle at work in an incident-response scenario. Not only was information collected, processed, and analyzed, but it was disseminated in such a way as to generate new requirements and feedback, starting the process over again until the analysts had reached a solid conclusion and could act with decisiveness, eradicating 43,000 malware installations at the time that the report was published. The published report, also part of the dissemination phase, allowed incident responders to better understand the tactics and motivations of this actor group.

SolarWinds

In December of 2020, news broke of a massive intrusion at the Texas-based company SolarWinds, which makes software for monitoring and managing IT networks and is a very popular tool in many large networks, including cybersecurity companies, governments, and Fortune 500 companies. The activity was detected after the cybersecurity firm FireEye, a SolarWinds customer, identified that their networks had been compromised and a set of tools they developed for identifying intrusions had been accessed by an unknown entity. Their own investigation into the intrusion on their network led them to identify SolarWinds as the source of the intrusion.

Their analysis indicated that the SolarWinds' networks had been compromised in late 2019 and their software tampered with, so that a software update that was pushed to all its clients contained a backdoor that would allow the adversaries access to those networks as well. This was not the first software-based supply chain attack; however it was notable for its size and scale - estimates suggested that over 18,000 of SolarWinds'

customers were impacted. It was also notable for its response, which, 20 years after Moonlight Maze, showed how far cyber threat intelligence has come as a discipline. Once identified, FireEye published a blog post with details about the incident and ways for others to detect activity on their network. Additional teams jumped in to analyze activity on their networks and continued to share indicators and findings, both publicly and through established threat sharing groups, allowing a picture of the overall attack to quickly develop. The Department of Homeland Security published guidance on supply chain attacks, taking the response from a mentality of just reacting to one isolated incident to thinking about how this new information shapes the way the industry should think about preparing for intrusions in the future. While certainly not a perfect process, SolarWinds illustrates the direction that cyber threat intelligence can play in incident response and how it can help not only the organizations directly impacted but help to surface important lessons for others.

Both the Axiom Group attacks and the SolarWinds software supply chain intrusion were information-seeking, espionage-related attacks, but nation-state sponsored attackers aren't the only thing that incident responders have to worry about. Financially motivated criminal activity is also evolving, and those actors are also working hard to stay ahead of defenders and incident responders. One of the most significant tactic changes for financially motivated criminals in recent years is the move to ransomware. Ransomware attacks use tools to encrypt data on a network and then charge a ransom for the key to decrypt the data. The concept of ransomware has been around for decades; however its usage has increased drastically since 2012, along with its impact. Although ransomware attacks do not always involve strategic and coordinated attacks against multiple organizations, the groups executing ransomware attacks often target different victims using the same tactics, toolsets, and often even targeting information. Defenders working against these financially motivated attacks can also leverage intelligence-drive incident response to identify early indications that their networks have been breached by these actors before the actual encryption process has begun.

Conclusion

Despite the many advances in the computer security field, attackers continue to adapt - but they do not have to outpace defenders. Intelligence-driven incident response allows us to learn from attackers; to identify their motivations, processes, and behaviors; to identify their activities even as they seek to outwit our defenses and detection methods. The more we know about attackers, the better we can prevent, detect, and respond to their actions.

We have reached the point where a structured and repeatable process for implementing intelligence in the incident-response process is necessary, and this book aims to provide insight into that process. Throughout this book, we provide various models and methods that can be viewed as the building blocks of intelligence-driven incident response, as well as the background as to why these models are beneficial to and integrate with incident response. There is no one-size-fits-all approach. In many cases, the incident or the organization will dictate which specific combination of models and approaches fits best. Understanding the foundational principles of intelligence and incident response, as well as the specific methods for integrating them, will allow you to build a process for intelligence-driven incident response that will work for you and to develop that process to meet the needs of your organization.

Sources :

https://www.fbi.gov/news/stories/morris-worm-30-years-since-first-major-attack-on-internet-110218

https://www.sciencedirect.com/topics/computer-science/moonlight-maze

https://www.mandiant.com/resources/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor

https://threatpost.com/solarwinds-attackers-new-tactics-malware/176818/

Marcus Willett (2021) Lessons of the SolarWinds Hack, Survival, 63:2, 7-26, DOI: 10.1080/00396338.2021.1906001

http://www.novetta.com/wp-content/uploads/2014/11/Executive_Summary-Final_1.pdf

Chapter 2. Basics of Intelligence

A NOTE FOR EARLY RELEASE READERS

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

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

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

> *It consists of gathering facts. … It consists of forming hypotheses on the basis of these facts, of testing these hypotheses for traces of one's own ignorance or bias, of cleansing them if possible. The goal is to build better hypotheses than already exist and to establish them as relatively more true: it is to reveal a sharper picture of what happened and to make a closer approach to actuality than anyone has yet contrived.*
>
> —Sherman Kent

Intelligence analysis is one of the oldest and most consistent concepts in human history. Every morning people turn on the news or scroll through feeds on their phones, looking for information that will help them plan their day. What is the weather report? What implications does that have for their activities for that day? How is the traffic? Do they need to plan for extra time to get to where they need to go? External information is compared to

an internal set of experiences and priorities, and an assessment is made of the impact on the target subject—the individual in question.

This is the basic premise of intelligence: taking in external information from a variety of sources and analyzing it against existing requirements in order to provide an assessment that will affect decision making. This occurs at the individual level as well as at higher levels; this same process is implemented at the group, organization, and government level every single day. There is one big catch though - unlike many forms of day-to-day analysis, intelligence analysis involves trying to understand something about an adversary who very much wants to stay hidden from you. The weather report - although occasionally inaccurate - was not intentionally tricking you into leaving your umbrella at home so that you would get soaked in a downpour. Because of this, intelligence analysis almost always involves some aspect of secrecy. Even when it is not part of a classified government program, it involves an entity that does not want you to have the whole picture. Likewise, you don't want that entity to know what you know about it, otherwise it might change tactics and cause you to start over at the beginning. In fact, a significant debate in intelligence work is around the concept of intelligence-gain-loss (which we will cover in depth in chapter $) where analysts have to determine how much intelligence value will be lost by taking an action that would warn the adversary that their presence, tactics, or tools have been identified.

Although there are a growing number of intelligence training programs inside and outside of traditional government and military fields, many individuals currently conduct some form of intelligence analysis on their own without formal training, and many security teams work through similar processes as they conduct investigations without realizing that they are, in fact, engaged in intelligence analysis. While intuitive analysis can be beneficial to a security program, it is even more useful when basic structures, such as processes and models, are utilized to streamline intelligence work, account for biases, and make the analytic judgements defensible and repeatable.

When businesses and governments conduct intelligence operations, it is based on a formalized process and doctrine that have been captured over the years. In addition, there are formalized processes which have been specialized for intelligence operations in information security and incident response. This chapter walks through key concepts, some from intelligence —some from security, and a few that combine both. We'll start with abstract concepts that are primarily pulled from intelligence doctrine and move toward the more concrete concepts that can be applied directly to your incident-response investigations.

Intelligence and Research

Intelligence as a discipline follows the same basic principles of other types of applied research, however there are several significant differences; secrecy, timeliness, and lack of reproducibility. The first, as we mentioned briefly above, is the understanding that intelligence often deals with matters that the subject or the target is actively trying to keep hidden. Researchers cannot go to an archive or search online repositories and find *all* of the information that they need. They may be able to find a great deal of relevant information, but there will always be an unspoken understanding that key pieces of information are intentionally not included in public information.

The second difference is that timeliness is far more significant with intelligence analysis than in other forms of research. If the information is not analyzed and presented to decision makers ahead of when it is needed, then it is likely no longer relevant. The third difference is that reproducibility, or being able to replicate findings, is often not done because of the first two principles. Most analysts will not have access to the exact same information to provide external validation of an analytic judgment, and the timeliness required of intelligence products means that the peer-review process is a rarity in the field, and often only conducted in the aftermath of an intelligence failure.

Data Versus Intelligence

Before tackling anything else, it's important to clear up one of the most important distinctions of this discussion: the difference between information, data, and intelligence. All of these are significant terms in the security community and unfortunately are often used interchangeably, to the point that many practitioners have a difficult time articulating the difference between them.

"Joint Publication 2-0," the US military's primary joint intelligence doctrine, is one of the foundational intelligence documents used today. In its introduction, it states, "Information on its own may be of utility to the commander, but when related to other information about the operational environment and considered in the light of past experience, it gives rise to a new understanding of the information, which may be termed intelligence."

Data is a piece of information, a fact, or a statistic. Data is something that describes something that is. In our previous example about the weather report, the temperature is a piece of data. It is a fact, something that has been measured using a proven and repeatable process. Knowing the temperature is important, but in order to be useful for decision-making, it must be analyzed in the context of what else is going on that day. In information security, an IP address or a domain are pieces of data. Without any additional analysis to provide context, they are simply facts. When various data points are gathered and analyzed to provide insight around a particular requirement, it becomes intelligence.

Intelligence is derived from a process of collecting, processing, and analyzing data. Once it has been analyzed, it must be disseminated in order to be useful. Not only does intelligence need to be disseminated, but it needs to reach its intended audience in a timely manner. Intelligence that does not get to the right audience is wasted intelligence. Wilhelm Agrell, a Swedish writer and historian who studied peace and conflict, once famously said, "Intelligence analysis combines the dynamics of journalism with the problem solving of science."

The difference between data and true intelligence is analysis. Intelligence requires analysis that is based on a series of requirements and is aimed at answering questions relevant to decision makers. Without analysis, most of the data generated by the security industry remains simply data. That same data, however, once it has been properly analyzed in response to requirements, becomes intelligence, as it now contains the appropriate context needed to answer questions and support decision-making.

INDICATORS OF COMPROMISE

There was a time when many people considered indicators of compromise, or IOCs, to be synonymous with threat intelligence. IOCs, which we will reference a lot and cover in depth later in the book, are things to look for on a system or in network logs that may indicate that a compromise has taken place. This includes IP addresses and domains associated with command-and-control servers or malware downloads, hashes of malicious files, and other network- or host-based artifacts that may indicate an intrusion. As we will discuss throughout this book, however, there is far more to threat intelligence than IOCs, although IOCs still remain one of the most common types of technical intelligence around intrusions.

IOCs have gotten a bad reputation over the years, and while analysts once loved collecting as many as they could, IOCs may now be rejected out of hand as the pendulum swung too far in the other direction. Just because data is not intelligence does not mean that the data has no value - in fact without data there can be no intelligence! So rather than dismissing IOCs as useless artifacts of simpler times, value them for what they are - pieces that can help both detect threats on the network and be used in post-incident analysis and strategic research. We will discuss how to use IOCs for both of these use cases in later chapters.

Sources and Methods

Now that we have cleared up the distinction between information, data, and intelligence, the natural next question is, “Where should I get this data from so that I can analyze it and generate intelligence?”

Traditional intelligence sources are most often centered around the INTs, which describe where the data is collected from:

HUMINT

> Human-source intelligence is derived from humans, either through covert or clandestine methods or from overt collection such as from diplomats. Human-source intelligence is the oldest form of intelligence collection. There is serious debate about whether cyber-threat intelligence can be derived from HUMINT, however there is a growing body of evidence to suggest that HUMINT can be a crucial part of the story. One example is interviews or conversations with individuals who are involved with or have firsthand knowledge of intrusions, such as when researchers from Kaspersky Labs were able to connect with a systems administrator from a network compromised as part of Moonlight Maze. Not only did the sysadmin give inside information and perspective that only a person who experienced an intrusion can, he also provided access to a server from the intrusion that had been sitting under his desk - HUMINT can often lead to additional data collection. Another example that many describe as HUMINT is information gained from interactions with individuals via restricted or members-only online forums. This type of intelligence gathering could also be considered SIGINT, as it is derived from electronic communications.

SIGINT

> Signals intelligence includes intelligence derived from the interception of signals, including communications intelligence (COMINT), electronic intelligence (ELINT), and foreign instrumentation signals intelligence (FISINT). Most technical intelligence collection falls under SIGINT, because after all, computers function using electronic signals, so anything derived from a computer or other networking device could be considered SIGINT.

OSINT

Open source intelligence is gathered from publicly available sources, including news, social media, and commercial databases as well as a variety of other nonclassified sources. We discussed previously that intelligence analysis, including cyber threat intelligence, involves some aspect of secrecy. This does not mean, however, that ALL intelligence sources must be secret. Published reports on cyber-security threats is one type of OSINT that can be incredibly useful in intelligence analysis. When dealing with government-backed actors, publications that detail the organizational structure of that government's offensive cyber forces can provide a wealth of knowledge. OSINT can also help reveal technical details about things like IP addresses or domain names that are publicly accessible; for example, a WHOIS query detailing who registered a malicious domain.

IMINT

Imagery intelligence is collected from visual representations, including photography and radar. IMINT is not typically a source of cyber-threat intelligence, however there are always cases where visual representations can provide critical information, such as the ability to watch troop movements during large scale, government-backed denial of service attacks, as were previously seen with the Russia and Georgia conflict.[1]

MASINT

Measurement and signature intelligence is gathered from technical means, excluding signal and imagery. MASINT often includes signatures from nuclear, optical, radio frequency, acoustics, and seismic signatures. As MASINT specifically excludes signals intelligence, it is also not a typical source of cyber-threat intelligence.

GEOINT

Geospatial intelligence is collected from geospatial data, including satellite and reconnaissance imagery, maps, GPS data, and other sources of data related to locations. Some organizations consider IMINT to be a part of GEOINT, and some believe it is a separate discipline. Similar to IMINT, GEOINT is not a typical source of cyber-threat intelligence, but it can provide contextual information on threats to help you understand how attackers may use the cyber domain to achieve their goals.

At this point, many would bring up a variety of other INTs that have popped up over the years, including cyber intelligence (CYBINT), technical intelligence (TECHINT), financial intelligence (FININT), and the most recent one we found - CyberHumint - but most of these new terms are already covered by other intelligence-collection methods. For example, cyber intelligence is primarily derived from ELINT and SIGINT. It is not important to get into an argument about the number of INTs in existence; the important thing is to understand the source of the data. At the end of the day, if it helps to refer to a specific collection type as its own INT, then go ahead; just be prepared to deal with the eventual terminology conflicts that tend to pop up in this field.

In addition to the traditional intelligence-collection disciplines listed here, some collection methods are often utilized specifically in cyber threat intelligence. It is useful to have a solid understanding of where this specific threat data comes from:

Incidents and investigations

This data is collected from the investigation of data breaches and incident-response activities. This is often one of the richest data sets used in cyber-threat intelligence because investigators are able to identify multiple aspects of the threat, including the tools and techniques that are used, and often can identify the intent and motivation behind the intrusion.

Honey pots and honey nets

These devices are set up to emulate machines or entire networks and gather information about interactions with these devices. There are many types of honey pots: low interaction, high interaction, internal honeypots, and honeypots on the public internet. Honeypot information can be useful as long as you know the type of honeypots it comes from, what they were monitoring for, and the nature of the interactions. Traffic gathered from a honeypot that captures exploit attempts or attempts to install malware on a system are far more useful in analysis than scanning or web-scrapping traffic.

Forums and websites

A variety of companies claim to have deep web or dark web collections. In many cases, these companies are referring to forums and chatrooms with restricted access that are not easily accessible from the internet. In these forums and sites, individuals often exchange information that is valuable after it's analyzed. There are so many of these types of sites that it is nearly impossible for any one company to have complete coverage of *the dark net*, so be aware that the collection is often limited in scope and will differ from that of other companies that claim to have similar data.

Even these techniques are new iterations of common techniques of the past. What's old is new as technology evolves, and intelligence is no different. The philosopher George Santayana's missive about forgetting the past is as true as ever.

MILITARY JARGON

One common point of contention in information security is the use of military terminology. Although intelligence has existed for centuries, it was codified in doctrine by military entities in documents such as the US Army's "Joint Publication 2-0: Joint Intelligence," and the UK's "Joint Doctrine Publication 2-00 - Understanding and Intelligence Support to Joint Operations." The majority of nonmilitary intelligence applications still pull heavily from the general principles captured in these documents, which results in a high volume of military terms in modern intelligence analysis. This means that related fields, such as cyber-threat intelligence, often pull heavily from military doctrine. However, just as with marketing, military jargon is useful in some situations and not useful in others. If the use of military terminology gets in the way of conveying your message, it may be a good time to use different terms.

Models

Models are a critical tool in the analyst's toolkit. Without models, many analysts would not be able to keep up with the quantity of data that exists in the world or the requirements that they provide meaning around this data. They would struggle to synthesize the information that they collected from all the various sources we just discussed and make meaning and insight that are relevant to decision makers.

One of the best working definitions of models for intelligence purposes comes from the book "Quantitative Intelligence Analysis" - "Models refer to the cognitive, conceptual, mathematical, physical, or otherwise logical representation of systems, entities, phenomena, or processes." The two most common types of models used in cyber threat intelligence are mental models and conceptual models. Mental models are cognitive models that hold an analyst's perception of reality. Even when analysts have gone through similar courses of formalized training, their mental models may be

very different because of the many diverse ways that individuals perceive and process information.

Conceptual models on the other hand, are representations of explicit knowledge, and are commonly the result of intelligence synthesis based on mental models that have been codified. Some of the most useful models have been derived from an analyst capturing and codifying a way that they intuitively approached or thought about problems. One example of this is the Diamond Model for Intrusion Analysis, which we cover in depth in Chapter 3, The Basics of Incident Response. Sergio Caltegrione, one of the creators of the model, described the generation of a new model as "a long progression of understanding the work we were repeatedly doing, until we understood it well enough to abstract it. Once abstracted, it became infinitely more useful because we could then ask new questions [of the model]." Sergio also described models as formulas - you have to understand its parts and their relationships before you can fully utilize them.

The goal of codifying a conceptual model is to be able to structure information so that it can be analyzed and acted on. Many models that are frequently used in intelligence analysis, including the diamond model, are covered further in Chapters 3, the Basics of Incident Response, and 7, Analyze.

Using Models for Collaboration

One of the key benefits of using explicit, conceptual models (rather than exclusively using mental models) is that it enables collaboration. Intelligence collaboration has been likened to "thinking in public", because it requires analysts to verbalize or otherwise articulate mental processes of analysis and synthesis that can be difficult to describe. Models are critical to analytic collaboration and will result in higher quality intelligence, and therefore it is worth the time investment to make sure that team members understand the different conceptual models that are commonly used and also understand the process for codifying the mental models that they commonly use in analysis.

Process Models

Models can be divided into two broad categories: models that represent our thinking, and models that represent the subject of analysis. The first type of model can be thought of as being used to give structure to processes - such as how we think, or the process of generating intelligence. This section covers two models that are used to effectively generate and act on intelligence. The first is the OODA loop, which can be used in making quick, time-sensitive decisions. The second is the intelligence cycle, or intelligence process, which can be used to generate more-formal intelligence products that will be used in a variety of ways, from informing policy to setting future intelligence requirements.

USING MODELS EFFECTIVELY

George E.P. Box said, “All models are wrong; some models are useful.” Every model is an abstraction that’s useful for understanding a problem. On the other hand, by its very nature, every model is reductionist and throws out important details. It’s not important to fit all data into a particular model, but it’s always valuable to use models as a way to understand and improve your thought processes.

OODA

One of the most referenced military concepts in security is *OODA*, an acronym for observe, orient, decide, act. The OODA loop, shown in Figure 2-1, was developed by fighter pilot, military researcher, and strategist John Boyd in the 1960s. He believed that a fighter pilot who was at a disadvantage against an adversary with more-advanced equipment or capabilities could be victorious by using OODA to respond more quickly to stimuli and effectively attack the opponent’s mind through decisive actions.

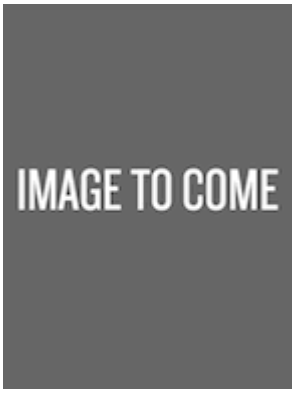

Figure 2-1. The OODA loop

Here's an introduction to each of the four stages.

Observe

The Observe phase centers around the collection of information. In this phase, an individual collects information from the outside world—anything and everything that could be useful. If the individual is planning to catch a baseball, this phase is about observing the baseball to determine its velocity and trajectory. If the individual is trying to catch a network attacker, the observation includes gathering logs, monitoring systems, and collecting any outside information that could help identify the attacker.

Orient

The Orient phase puts the information collected during the Observe phase into context with already known information. This takes into account past experience, preconceived notions, expectations, and models. For the baseball example, orientation uses what the observer knows about how a ball moves, taking into account its velocity and trajectory, to predict where it will go and how much force the impact will generate when it is caught. In the example of a network attacker, orientation takes the telemetry pulled from the logs and combines it with knowledge about the network, relevant attack groups, and previously identified artifacts such as specific IP addresses or process names. The orientation phase is heavily reliant on the mental models we previously discussed. Without a way to quickly and

accurately “sort” the data you have observed, orientation becomes a very difficult task.

Decide

At this point, information has been collected (observed) and contextualized (oriented), and now it is time to determine a course of action. The decide phase is not about executing an action. It is about debating various courses of action until the final course of action is determined.

In the baseball example, this phase includes determining where to run and how fast, how the fielder should move and position her hand, and anything else needed to attempt to catch the ball. In the case of dealing with a network attacker, it means deciding whether to wait and continue to observe the attacker’s actions, whether to start an incident-response action, or whether to ignore the activity. In either case, the defender *decides* on the next steps to achieve their goal.

Act

After all that, the Act phase is relatively straightforward: the individual follows through with the chosen course of action. This doesn’t mean it’s 100% guaranteed to be successful. That determination is made in the observation phase of the next OODA loop, as the cycle begins again back at the observation phase.

OODA is a generalization of the basic decision-making process that everyone goes through thousands of times a day. This explains how individuals make decisions, but also how teams and organizations do so. It explains the process a network defender or incident responder goes through when gathering information and figuring out how to use it.

The OODA loop is used by not only one side. While we, as defenders, go through the process of observing, orienting, deciding, and acting, in many cases the attacker is as well. The attacker is observing the network and the defender’s actions in that network, and deciding how to respond to changes in the environment and attempts to kick them out. As with many things, the

side that can observe and adapt faster tends to win. Figure 2-2 shows the OODA loop for both an attacker and defender.

One thing to note about the OODA loop is that it can often trip analysts up as they try to think through adversarial responses to their actions - “first I will do x, then they will do y, and then I will…” - when in reality it can be difficult to know how the adversary will respond. There will always be some level of uncertainty when dealing with adversaries who are human and may act unpredictably, which can make the risk-adverse wary of making any moves. When in doubt, reach back to your requirements, or even further back to the goals and mission of your team or program. Are you tasked with defending the network? With safeguarding user data or public safety? Make sure those requirements and missions are included in the orientation phase, which can help to avoid decision-paralysis when it comes time to decide what the best course of action is.

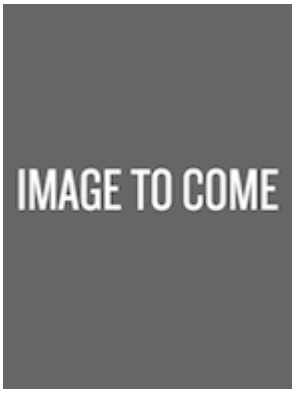

Figure 2-2. Competing OODA loop of the defender and the attacker

>

MULTIPLE DEFENDER OODA LOOPS

Beyond Attacker-Defender OODA loops, it's also useful to think about Defender-Defender OODA loops—that is, how the decisions we make as defenders can impact other defenders as well. Many decisions that defensive teams can make may essentially set up race conditions for other defenders. For example, if a defender executes an incident response and then publicly shares information about the attack, then the first defender has started the clock on all other defenders to ingest that intelligence and use it. If an attacker can move through the OODA loop faster, find the public information about their activities, and change their tactics before the second defender can use the information, then they've turned inside (outmaneuvered and achieved a more ideal position) the second defender and can avoid serious consequences.

For this reason, it's important to consider how your actions and sharing impact other organizations, both adversaries and allies. In all cases, computer network defense is about slowing down the OODA loops of the adversary and speeding up the OODA loops of defenders.

This generalized decision model provides a template for understanding the decisions of both defenders and attackers. We'll discuss the cycle more moving forward, but in the end, this model focuses on understanding the decision-making processes of all parties involved.

Intelligence Cycle

The intelligence cycle, pictured in Figure 2-3, is the formal process for generating and evaluating intelligence. The cycle begins where the last intelligence process ended and continues to build off itself. The intelligence cycle doesn't need to be followed to the letter. In fact, processes explored later in this book will build upon it. You do have to be careful not to omit critical steps, however. If you start skipping entire steps, you run the risk of ending up with more data and questions instead of intelligence.

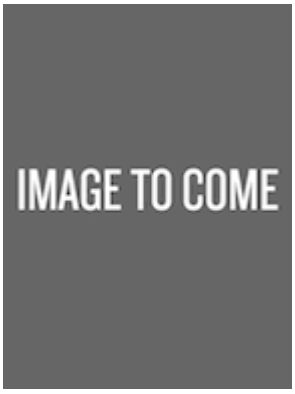

The intelligence cycle

To properly utilize the intelligence cycle, you need to know what is involved in the steps, which we will dive into next.

Direction

The first step in the intelligence cycle is direction - also known as requirements. *Direction* is the process of establishing the question that the intelligence is meant to answer. This question can be delivered from an outside source, developed by the intelligence team itself, or developed by stakeholders and the intelligence team. (This is sometimes called the *RFI process*, which we'll discuss in Chapter 4, Find. The ideal outcome of this process is a clear, concise question whose answer the stakeholders will find usable.

Much intelligence work is based around requirements. In the intelligence community, any work done or reports circulated needs to be tied directly to an intelligence requirement, either standing (long term) or priority (urgent and time sensitive) requirements.

Collection

The next step is collection of the data necessary to answer the question. This is a wide-ranging exercise that should focus on gathering as much data as possible from many sources. Redundant information adds value here, because corroboration is often important.

This leads to a key idea of developing an effective intelligence program: building a collection capability. It's difficult to know exactly what data might eventually prove useful, so building a broad capability to collect a wide variety of information is important. This includes tactical information such as infrastructure, malware, and exploits, as well as operational strategic information such as attacker goals, social media monitoring, news monitoring, and high-level document exploitation (identifying reports, such as those that vendors release about groups, and gathering information from them). Be sure to document the sources and take care: news stories often republish or reference the same original material, making it difficult to know what's corroboration and what's just a rehash of the same material. If it is impossible to determine the source of a particular data set, you may want to avoid using it as a collection source.

Collection is a process, not a one-time action. Using information from the first round of collection (such as gathering IP addresses) leads to a second round (such as using reverse DNS to find domains related to those IP addresses), which leads to a third round (using WHOIS to gather information about those domains). This exploitation becomes exponential as it builds upon itself. The focus at this point is not understanding how the data relates but simply developing as much information as possible. Combining it comes later. Also, don't forget to consider internal sources, such as an incident-management system. It's common for organizations to discover actors or attacks they're already intimately familiar with.

NAME DECONFLICTION

Naming presents a significant challenge in intelligence collection. While in the old days this focused on aliases and cover terms, today the field struggles with the fractured nature of intelligence collection and naming conventions. Every company, every intelligence sharing group, and every intelligence agency has its own names for various threat groups. The intrusion group APT1 is a great example: most commonly referred to as Comment Crew, this group was also known as ShadyRat, WebC2, and GIF89a by industry groups. Mandiant called them APT1. CrowdStrike called them Comment Panda. Ongoing intelligence determined their actual identity as Peoples Liberation Army Military Unit 61398. Collection against all of these names matters, as overlooking reporting that uses a particular name could lead to missing critical data.

Processing

Data is not always immediately usable in its raw format or in the format in which it was collected. In addition, data from different sources may come in different formats, and it is necessary to get it into the same format so it can be analyzed together. The *processing* necessary to make data usable is often an overlooked task, but without it, generating intelligence would be nearly impossible. In the traditional intelligence cycle, processing is part of collection. However, when dealing with the types of data and organizations involved in incident response, it may be useful to consider processing separately. Here are some of the most common ways to process data related to cyber threats.

Normalization

> Processing includes normalizing collected data into uniform formats for analysis. The collection process will generate nearly every conceivable kind of data result. Intelligence data comes in a variety of formats, from JSON to XML to CSV to plain text from email. Vendors share

information on websites in blog posts or tables, but also in PDF-based reports or even YouTube videos. At the same time, organizations tend to store data in different formats. Some organizations use a purpose-built threat-intelligence platform, while other organizations build customized solutions from wikis or internal applications.

Indexing

Large volumes of data need to be made searchable. Whether dealing with observables such as network addresses and mutexes or operational data such as forum posts and social media, analysts need to be able to search quickly and efficiently.

Translation

In some cases, regional analysts may provide human translation of source documents, but this is generally not feasible for most organizations dealing with information from all over the world. Machine translation, while imperfect, usually provides sufficient value so that analysts can find items of interest. If necessary, they can then be escalated to specialists for a more accurate translation.

Enrichment

Providing additional metadata for a piece of information is important. For example, domain addresses need to be resolved to IP addresses, and WHOIS registration data fetched. Google Analytics tracking codes should be cross-referenced to find other sites using the same code. This enrichment process should be done automatically so that the relevant data is immediately available to analysts.

Filtering

Not all data provides equal value, and analysts can be overwhelmed when presented with endless streams of irrelevant data. Algorithms can filter out information known to be useless (though it may still be searchable) and bubble up the most useful and relevant data.

Prioritization

The data that has been collected may need to be ranked so that analysts can allocate resources to the most important items. Analyst time is valuable and should be focused correctly for maximum benefit to the intelligence product.

Visualization

Data visualization has advanced significantly. While many analysts fear vendor dashboards because of the clutter they typically contain, designing a visualization based on what analysts need (rather than what marketing and executives think looks good) can assist in reducing cognitive load.

Taking the time to process data effectively enables and improves future intelligence efforts.

Analysis

Analysis, as much an art as it is a science, seeks to answer the questions that were identified in the Direction phase. In intelligence analysis, data that has been collected is characterized and considered against other available data, and an assessment is made as to its meanings and implications. Predictions are often made as to future implications. There are various methods for conducting analysis, but the most common is to use analytic models to evaluate and structure the information, identify connections, and make assessments about their implications. In addition to preexisting models, which we cover later in this chapter, it is also common for analysts to develop their own models that work with their particular data sets or way of interpreting information.

The goal of the Analysis phase is to answer the questions identified in the Direction phase of the intelligence cycle. The type of answer will be determined by the nature of the question. In some cases, the analysis may generate a new intelligence product in the form of a report or could be as

simple as a yes/no answer, most often backed up with a confidence value. It is important to understand what the output will be before beginning the analysis.

Analysis is not a perfect science and must often be conducted with incomplete information and uncertainty. It is important that analysts identify and clearly state any information gaps in their analysis so that decision makers can be aware of potential blind spots in the analysis. Information gaps can also drive the collection process to identify new sources in order to reduce those gaps. If the gaps are significant enough that an analyst does not think it is possible to complete the analysis with the current information, then it may be necessary to go back to the Collection phase and gather additional data. It is much better to delay the final analysis than to provide an assessment that the analyst knows is flawed.

It is important to note that *all intelligence analysis is generated by a human*. If it is automated, it is actually processing instead, which is a critical step in the intelligence cycle but is not by itself analysis.

Dissemination

At this point, the process has generated real intelligence: a contextualized answer to the question posed in the Direction phase. A report with an answer is useless until it's shared with the relevant stakeholders: those who can use this intelligence. In plenty of documented intelligence failures, analysis was spot-on but dissemination failed. *Intelligence must be shared with relevant stakeholders in the form they find the most useful*. This makes dissemination dependent on the audience. If the product is aimed at executives, it's important to consider length and phrasing. If it's aimed at implementation in technical systems (such as IDS or firewalls), this could require vendor-specific programmatic formats. In any case, intelligence must be usable by the relevant stakeholders.

Feedback

Often forgotten, the Feedback phase is key to continuing intelligence efforts. *The Feedback phase asks whether the intelligence that was*

generated answers the direction successfully. This results in one of two outcomes:

Success

If the intelligence process answered the question, the cycle may be over. In many cases, though, a successful intelligence process leads to a request for more intelligence based on either new questions or the actions taken based on the answer given.

Failure

In some cases, the intelligence process failed. In this case, the Feedback phase should focus heavily on identifying the aspect of the original direction that was not properly answered. The following Direction phase should take special care to address the reasons for that failure. This usually comes down to a poorly structured Direction phase that didn't narrow the goal enough, or an incomplete Collection phase that was unable to gather enough data to answer the question, or improper analysis that did not extract correct (or at least useful) answers from the data available.

Using the Intelligence Cycle

Let's consider how the intelligence cycle can be used to start learning about a new adversary.

One of the most common questions a chief information security officer, often abbreviated as CISO, asks (hopefully before she gets asked it herself) is, "What do we know about this threat group I heard about?" A CISO will want a basic understanding of a group's capabilities and intention, as well as an assessment of relevance to a given organization. So what does the intelligence process look like in this situation? Here is an example of what is involved in each step of the intelligence cycle to meet the CISO's needs:

Direction

This came from a key stakeholder: the CISO. “What do we know about *X* threat group?” The real answer sought is a target package, which we’ll explore in detail later.

Collection

Start with the original source, most likely a news article or report. That document will usually provide at least some context for beginning the collection. If the source material includes information about a specific intrusion or attack, it can be helpful to understand more about the entity that had been targeted and identify any potential motivations or goals of the intrusion. If indicators (IPs, URLs, etc.) exist, explore those as deeply as possible by pivoting and enriching. The source may itself point to additional reporting with IOCs, tactics, techniques, and procedures (TTPs), or other analyses.

Processing

This is very workflow/organization dependent. Getting all the collected information into a place where it can be used most effectively may be as simple as putting all the information into a single text document, or it may require importing it all into an analysis framework. Additional enrichment can be done with technical details related to the group. In addition, translation of reports or other documents may be necessary.

Analysis :

Using the collected information, the analyst will start by attempting to answer key questions:

- What are these attackers interested in?
- What tactics and tools do they typically use?
- How can defenders detect those tools or tactics?
- Who are these attackers? (Although this is always a question, it is not always one worth taking the time to answer.)

Dissemination

For a product like this that has a specific requester, the CISO, a simple email may suffice. Although in some cases limiting a response to this makes sense, a real product for proactive distribution to others will almost always create greater value.

Feedback

The key question: is the CISO pleased with the results? Does it lead to other questions? These pieces of feedback help close the loop and may begin a new series of collections.

The intelligence cycle is a generalized model that can be used to answer questions large and small. However, it is important to note that following the steps will not automatically result in good intelligence.

Qualities of Good Intelligence

The quality of intelligence relies primarily on two things: collection sources and analysis. Many times in cyber-threat intelligence we end up working with data that we did not collect ourselves, and therefore it is critical that we understand as much as possible about the information. When generating intelligence ourselves, we also need to ensure that we understand collection sources and are addressing biases in our analysis. Here are some things that should be considered to ensure that quality intelligence is produced:

Collection method

> It is important to understand whether the information is collected primarily from incidents or investigations, or whether it is being collected from an automated collection system such as a honeypot or a network sensor. Although knowing the exact details of the collection is not imperative—some providers prefer to keep their sources confidential—it is possible to have a basic understanding of where the data comes from without compromising collection resources. The more details you have about the way information was collected, the better

your analysis of this information will be. For example, it is good to know that data comes from a honeypot; it is better to know that it comes from a honeypot configured to identify brute-force attempts against remote web administration tools.

Date of collection

The majority of cyber-threat data that is collected is perishable. The lifespan of that data varies from minutes to potentially months or even years, but there is always a period of time when this information is relevant. Understanding when data was collected can help defenders understand how it can be acted upon. It is difficult to properly analyze or utilize any data when you do not know when it was collected.

Context

The collection method and date can both provide some level of context around the data, but the more context that is available, the easier it will be to analyze. Context can include additional details, such as specific activities related to the information and relationships between pieces of information.

Addressing biases in analysis

All analysts have biases, and identifying and countering those biases so that they do not influence analysis is a key component of quality intelligence. Some biases that analysts should seek to avoid include *confirmation bias*, which seeks to identify information that will support a previously formulated conclusion, and *anchoring bias*, which leads analysts to focus too heavily on a single piece of information while disregarding other, potentially more valuable information. We cover biases in depth in Chapter 7, Analyze.

Levels of Intelligence

The intelligence models we have examined thus far focus on a logical flow of information through a sort of analysis pipeline. But just as with incident analysis, this approach is not the only way to model the information. We can think about intelligence at different levels of abstraction, ranging from the highly specific (tactical) to the logistical (operational) to the very general (strategic). As we examine these levels of intelligence, keep in mind that this model represents a continuous spectrum with gray areas between them, not discrete buckets.

Tactical Intelligence

Tactical intelligence is low-level, highly perishable information that supports security operations and incident response. The customers for tactical intelligence include security operations center (SOC) analysts and computer incident response team (CIRT) investigators. In the military, this level of intelligence supports small-unit actions. In cyber-threat intelligence (CTI), this usually includes IOCs and observables as well as highly granular TTPs describing precisely how an adversary deploys a particular capability. Tactical intelligence enables defenders to respond directly to threats using methods such as hunting in the network for signs of adversary activity, prioritizing critical patching based on reports of active exploitation, or issuing information to employees of active phishing campaigns targeting their organization.

An example of tactical intelligence is IOCs related to an exploitation of a newly discovered vulnerability. These tactical-level IOCs include IP addresses conducting scans searching for the vulnerability, domains hosting malware that will be downloaded to the host if exploitation is successful, and various host-based artifacts that are generated during exploitation and installation of malware. This tactical intelligence would enable security operations teams to effectively hunt for malicious activity in the network while taking concrete steps to mitigate any future exploitation.

Operational Intelligence

In the military, *operational intelligence* is a step up from tactical. This information supports logistics and analyzes effects of terrain and weather on larger operations - in other words, it involves far more context than just the tactical mission at hand. In CTI, this usually includes information on campaigns and higher-order TTPs, as well as anticipated responses from adversaries and how those actions may impact other operations. It may also include information on specific actor attribution as well as capabilities and intent. Customers for operational intelligence include senior-level Digit Forensics and Incident Response (DFIR) analysts and other CTI teams.

This is one of the harder levels of intelligence for many analysts to work with, because it sometimes feels too general to be tactical but too specific to be strategic. Operational intelligence needs to be acted on somewhat urgently, however there are potentially far reaching implications of any action taken, so both the “observe” and “orient” phases of the OODA loop have far more components and complexities. In addition, the necessity of making a decision and acting on it is felt differently by analysts who are working on the data and decision makers, over at the executive level, who may want to spend a considerable amount of time weighing the best course of action.

Following the preceding example about tactical-level indicators of active exploitation of a vulnerability, operational-level intelligence would include information on how widespread the exploitation is, whether it is targeted or opportunistic, who else is being targeted, the purpose of the malware that is being installed, and any details on the actors who are carrying out the attacks. Understanding these details can support the generation of follow-up intelligence, including what other actions that may be seen, and should include information on the severity of the threat to help plan a response.

Strategic Intelligence

In the military or government, *strategic intelligence* deals with national and policy-level information and is often the accumulation of years of analytic work aimed at providing a holistic picture of a situation. In CTI, we think of this as supporting C-level executives and boards of directors in making

serious decisions about risk assessments, resource allocation, and organizational strategy. This information includes threat trends and actor motivations, along with additional information relevant to the organization. In the preceding example, strategic intelligence would include information on the motivations of the attackers, especially if the activity indicates a new or previously unidentified threat, and any information that indicates new tactics or attacker targeting that may require higher-level responses, such as new policies or an architecture change.

Confidence Levels

As mentioned previously, intelligence typically has different confidence levels associated with it. These confidence levels reflect the analysts' trust that the information is correct and accurate. For some types of data, this confidence may be on a numeric scale (for example, 0 to 100) and calculated using traditional statistical methods, while in other cases the confidence assessment is provided on a qualitative, subjective basis by analysts directly. It is important to identify confidence in two important areas: confidence in the source of the information, and confidence in an analyst's conclusions.

One common way of describing source reliability is the Admiralty Code or NATO System found in FM 2-22.3. This consists of two scales. The first evaluates the reliability of a source based on previous information, ranging from A (Reliable) to E (Unreliable). The second scale evaluates the degree of confidence in the information content itself, ranging from 1 (Confirmed) to 5 (Improbable). These two scores are combined for a particular piece of information based on the source and specific content, so that information known to be true from a source with a history of valid information might be evaluated as B1, but information that is improbable from a source with a history of invalid information would be evaluated as E5.

Sherman Kent, often referred to as the father of intelligence analysis, wrote an essay in 1964 called "Words of Estimative Probability," which describes various qualitative ways to describe confidence in an analyst's judgment. In

that essay, Kent shares one of the charts that he and his team use to assign and describe confidence (shown in Figure 2-4) but also writes that other terms may be used in their place as long as the meaning is understood and the terms are used consistently.

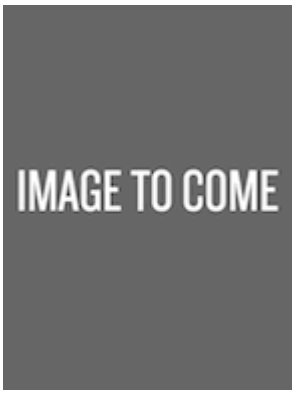

Figure 2-4. Sherman Kent's chart on estimative probability

Conclusion

Intelligence has a long and fascinating history and is a field that has adapted time and time again to changes in technology and to the growing complexities of the world. At its foundation are principles that are part of human nature - curiosity, assessing situations for danger, making connections. These foundations have been built upon with structured processes and models that make these often instinctive mental processes more academically rigorous and defensible. Biases, which humans cannot completely escape, can influence analysis and need to be identified and countered in analysis to avoid inaccurate assessments.

Intelligence is also a critical component of incident response, and many processes can be used to integrate intelligence principles into incident response investigations. It is important to understand the sources of intelligence that you will be relying on; there is a big difference in the way that you treat intelligence that came from a previous IR investigation in your network and the way you treat information that comes from a honeypot. Both types of information are valuable; they just have different applications. The next chapter dives into the specifics of incident response

and the models that help analysts implement intelligence-driven incident response.

---

1 https://www.history.com/news/russia-georgia-war-military-nato