Threat Intelligence in Practice

A Practical Guide to Threat Intelligence from Successful Organizations

Allan Liska

Beijing • Boston • Farnham • Sebastopol • Tokyo O'REILLY®

Threat Intelligence in Practice
by Allan Liska

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*.

Editor: Courtney Allen
Production Editor: Colleen Cole
Copyeditor: Dwight Ramsey
Interior Designer: David Futato
Cover Designer: Karen Montgomery
Illustrator: Rebecca Demarest

January 2018: First Edition

Revision History for the First Edition
2017-12-11: First Release

978-1-491-98208-2

[LSI]

As always, for Kris and Bruce

Table of Contents

Preface

There aren't many topics in cyber security that generate more arguments than threat intelligence. Security professionals have a wide range of views on the topic that range from severe eye rolls to a critical part of a well-run security program. What I present in this book are my thoughts about what threat intelligence is and how organizations can use threat intelligence to better protect themselves against all manner of threats.

These thoughts are gathered from my years spent as an intelligence analyst and from the thousands of organizations I have talked to about their threat intelligence programs. Not everyone will agree with everything I have written, and that is a good thing because hopefully these disagreements will start a conversation.

The goal of this book is to act as a primer for organizations who are considering building or rebuilding a threat intelligence program. This book is not designed to be a step-by-step guide, instead it is meant to be a spark. There should be enough information contained between these covers to get a team thinking about how to improve the security of an organization through the effective use of threat intelligence.

If you have any thoughts or questions about the tools I have laid out here, I would love to hear from you. Reach out to me any time. You can find me on Twitter as *@uuallan* or send me an email to *allan@allan.org*.

Acknowledgments

No book is ever solely the work of the author, there are a lot of people involved in the process. In that spirit there are quite a few people I need to thank. From O'Reilly I would like to thank superstar editor Courtney Allen and our word ninja, Virginia Wilson. I also would like to thank the great O'Reilly editors Colleen Cole, Nan Barber, and Dwight Ramsey.

In addition to the team at O'Reilly I would like to thank the smart technical reviewers whose feedback proved invaluable: Tim Gallo, Melissa Kelley, Amanda Berlin, and Tony Godfrey. I have so much respect for all four of you and hope I was able to successfully incorporate your suggestions.

Finally, I cannot express my thanks enough to Robert Morton and Eric Kobrin at Akamai and Jay Nancarrow for taking the time to share your thoughts on threat intelligence not only with me, but with everyone reading this book.

CHAPTER 1
Defining Threat Intelligence

Threat intelligence is gaining a more prominent role in running a modern security team. Of course, this prominence means that every security professional and vendor also wants the world to adopt their vision of threat intelligence. This leaves many organizations with two questions: what is threat intelligence, and can it can really help improve security?

The short answer to the second question is: it can and does, when implemented correctly. But, as with any complex system, there is no "Easy Button" for threat intelligence. The goal of this book is to provide an introduction to some of the basic themes of threat intelligence. This book is not designed to be comprehensive; instead, it is designed to start a conversation about building a successful threat intelligence program. This book provides guidelines and exposes pitfalls for any organization that is ready to build a Threat Intelligence Unit for the first time, or is looking to improve their existing intelligence team.

This chapter starts by defining threat intelligence. As silly as this may sound, without a common definition of the term, it is hard to build an effective program. The rest of the book revolves around the definition and the basic tenets of threat intelligence defined in this chapter.

NOTE

Military Terms

Threat intelligence in information security draws heavily upon years of intelligence experience from the military. Not just because the military has established intelligence frameworks, but because many information security threat intelligence professionals got their start in the military. To that end there are military intelligence terms used throughout this book, in part because these terms are commonly used by threat intelligence teams.

What Is Threat Intelligence?

There are a number of characteristics that define threat intelligence, but there is a general industry consensus around the definition proposed by Gartner:[1]

> Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard.

This definition covers all aspects of threat intelligence from collection, to processing, to the decision-making process. It also covers the many different types of information that should be collected as part of the intelligence process.

In short, threat intelligence is any information that can be correlated with additional information in a manner that allows an organization to improve their security in a tangible manner. For example, if a third-party vendor tells an organization that Anonymous has been launching attacks against other organizations in the same vertical, and provides a list of tools that Anonymous is using to launch these attacks, that is threat intelligence.

In order for data from outside sources to be considered threat intelligence, it must be:

- Relevant: It must impact the organization in some way.

1 McMillan, Rob, "Definition: Threat Intelligence" (*https://www.gartner.com/doc/2487216/definition-threat-intelligence*), Technology Research, May 16, 2013, accessed January 03, 2017.

- Actionable: Concrete steps can be taken by security teams to protect the organization.
- Contextual: There should be enough evidence included to enable an intelligence analyst to effectively rank the threat.

New threats spring up all the time, but not all of those threats are relevant to every organization. For example, a new vulnerability against the webmail application SquirrelMail is not relevant to an organization that is not running SquirrelMail, no matter how severe the vulnerability is.

Simply knowing that Anonymous has launched new attacks is not sufficient because there is no actionable information. That basic knowledge does not provide an organization with enough information to take steps to ensure it is protected against those attacks.

Finally, in order for information to be considered threat intelligence there must be context surrounding it. The SquirrelMail vulnerability mentioned above is a perfect example. While a vulnerability in a popular Internet-facing application could be a cause for concern, if an organization is properly managing and scanning its network assets, the security team should be able to quickly determine whether SquirrelMail is installed anywhere on the network. If SquirrelMail is not installed, then there is most likely no threat to the organization from the vulnerability. It is that additional context—knowledge of an organization's network assets—that can turn information into threat intelligence.

Of course, context does not always have to originate from within an organization. In fact, context can be entirely external. For example, if a third party provides an organization with the email address *wowsmith123456@posteo.net* as an indicator associated with the infamous NotPetya ransomware attacks, that provides some level of context. If an organization matches against that address in their logs or in collected NetFlow data, they will know there is a good chance they have a NotPetya attack. The first part of the intelligence, that *wowsmith123456@posteo.net* is associated with the NotPetya ransomware attack, is not something that would be picked just reviewing internal data, unless the organization happened to fall victim to the NotPetya attack. Instead, the context around the email address would come from a third-party source that is collecting data from thousands of different networks.

Context can be deeper than just a single connection. Knowing that the email address is associated with the NotPetya attack is a baseline of useful information. But tying that email address to other indicators associated with the attack, such as IP addresses and domains used for command and control purposes or file hashes associated with the malware is even better. Providing information about attack atmospherics—such as what organizations are being targeted or the fact that the email address has been disabled, so if an organization is infected they should not try to pay the ransom—is the ideal level of context.

What Threat Intelligence Isn't

Now that there is a consensus definition of threat intelligence, it is important to take a step back and explain what threat intelligence isn't. Threat intelligence is not a list of IP addresses or domain names with no context. Threat intelligence is also not a proprietary platform that exists solely inside a security vendor's tool. Finally, threat intelligence is not data that rests solely in a portal or in a report that is isolated from the rest of an organization's network.

Each of the examples listed in the previous paragraph can help create threat intelligence, but they are not threat intelligence in and of themselves.

Examples of "not" threat intelligence that are seen most often tend to be around the context of the indicators. A third party might share that IP address 101.200.81.187 is malicious. That isn't threat intelligence because there is no context surrounding it. Knowing that a third party classifies an IP address, domain, file hash, or some other indicator as malicious, without understanding how they came to that conclusion, is not threat intelligence. In fact, these types of context-free data feeds can make the security team's job more difficult, as they don't have any context around which to judge a threat.

Problems with context can even happen with internal intelligence, which is why documentation between teams is so important. For example, a firewall administrator may receive a report that several hosts on the network are attempting to call out to a single IP address several times a minute. Based on that report, the firewall administrator puts a rule into the firewall blocking traffic to that IP address and goes on to the next task.

With no documentation in place, when the firewall administrator goes home for the evening and the night administrator comes into work, there is no information about why the rule was created. In addition, there is no information about whether or not the original traffic was really a threat, or, if it was, what the associated malware was and whether it has been as successfully cleaned. Actions, even those seemingly benign as a firewall rule change, that pass between different security teams can become threat intelligence for the organization. But they require context to understand what the threat is and determine whether there are additional actions that need to be taken.

NOTE

Data Cleansing

Data cleansing is the process of cleaning up data to remove obvious mistakes or incorrect entries and it is an important part of good threat intelligence. Intelligence delivered without at least some attempt at data cleansing is not threat intelligence, in fact it usually makes more work for intelligence teams and can make an organization less secure.

This is often seen in data feeds where RFC 1918 (the private IP space) addresses are accidentally slipped into the list or users are encouraged to block well-known IP addresses, such as 8.8.8.8 (one of Google's DNS servers) with no explanation.

Even the best third-party providers occasionally make mistakes, but threat lists with no sanity check are not threat intelligence.

Data Feeds versus Threat Intelligence

Many organizations get their start with threat intelligence through data feeds. A data feed is a list of indicators provided by a third party that can be correlated against internal security systems to find matches that can be acted upon.

The appeal of data feeds is understandable—there is a lot of malicious activity happening at any given time on the Internet. No single organization can see all of that activity. By aggregating indicators from security incidents across the Internet and combining them into a single data feed, organizations can better protect themselves

against coming attacks or use the indicators to find attacks they may have missed.

As most organizations quickly find out, data feeds rarely live up to their promise. Too often, data feeds contain outdated data and don't provide context or give security teams enough information to make the data actionable. Even the purported purpose of data feed correlation, to provide relevance, doesn't always work as expected.

Many organizations who start down the path of a threat intelligence program with data feeds are surprised to discover that these feeds often create more work than expected. Correlating data feeds against logs from other network devices in a Security Information and Event Manager (SIEM) seems like a natural fit—after all, it will help security teams identify security incidents that may have been missed otherwise. Unfortunately, what many security teams quickly find out is that, rather than making their lives easier, data feeds generate more work. By identifying even more security incidents and, often, many more false positives, data feeds can quickly become a burden to security teams.

That is the primary difference between a data feed and threat intelligence, even threat intelligence delivered in feed format. Threat intelligence helps security teams improve the security posture of the organization and makes the security team more effective and responsive to potential security incidents.

That doesn't mean that data feeds don't serve a purpose. In fact, the data feed concept is an important one for improving an organization's security posture. One of the biggest challenges that organizations of all sizes face when it comes to network security is that for 30 years, the security industry, as a whole, has solved the latest security problem by "adding a box" to the mix. First, it was firewalls, then intrusion detection systems (IDS), proxies, endpoint protection, web application firewalls (WAF), data loss prevention (DLP), and so on, to the point that many organizations have 20 or more security systems in their network, none of which talk to each other.

This leads many security teams to suffer from "console fatigue." They constantly have to jump from one security vendor's console to another's, and correlation is often done manually, as in "Oh wait, I think I saw that same indicator in an alert from my endpoint vendor. Let me go see if I can find it."

Data feeds, delivered in an automated fashion, from one security system to another can help improve the security of an organization by having those systems communicate. That automation between systems, even cloud systems, allows the security team to have a better view of what is happening within their organization and allows them to make more effective and faster decisions when it comes to prioritizing incidents.

Even external data feeds, when they are processed correctly, can help an organization improve security. Some data feeds can be fed directly into a firewall, mail server, or even fed directly into a DNS server as a Response Policy Zone (RPZ). These tend to be specialty feeds that are well-vetted and serve a specific purpose. Most threat data feeds from reliable sources can contain valuable data that, when combined with other information, can eventually become threat intelligence.

Threat Intelligence from the Inside Out

Inevitably when an organization is serious about building out a threat intelligence program they start by looking externally. Whether the start involves looking for threat feeds, as described above, or reaching out to vendors who sell threat intelligence, there is a strong draw to rely on vendors to deliver threat intelligence.

The problem is that even the best threat intelligence providers can't deliver effective threat intelligence unless the organization knows what their needs are. Remember, actual threat intelligence has to be relevant, provide context, and be actionable. In order for these requirements to be met there must be something against which to correlate the external collection.

To that end, the strongest threat intelligence programs start internally and work their way out. Often, this is the hardest part of building a threat intelligence program: getting the internal data from the network into the hands of the people who need to be able to analyze it.

Make no mistake, there is a lot of valuable security data inside an organization of any size that is owned by different teams and this data is very siloed. For example, the vulnerability management team gets threat intelligence from their vendors regarding new vulnerabilities and exploits that target those vulnerabilities. However, the vul-

nerability team is usually separate from the security team, so the security team doesn't always learn about these new threats in a timely fashion. But it goes deeper than just getting data from one team to another—a successful threat intelligence program has to encompass the entire organization, and it needs to start from the top.

Defining a Mission

When starting a threat intelligence program, many organizations don't bother answering the most fundamental question, "Why does our organization need threat intelligence?" In security engineering parlance this is known as the "What problem are we trying to solve?" question. Before doing anything else, this question must be answered in order for a threat intelligence program to be successful.

On the surface this may seem like an easy question to answer: "To protect our organization," "To understand the threats facing our organization," or "Because the CEO said to do it," are all very common answers. And while those are important components of a threat intelligence program, they are not complete answers. Those are all surface answers (with the exception of the CEO answer), but they don't really define the unique intelligence needs of an organization.

Instead, the goal of any threat intelligence program should be to protect the most valuable asset of an organization, the asset that if it wound up on a "paste site" would potentially irreparably damage the reputation or value of an organization. That asset could be a customer database, it could be the designs for new automobiles, it could be the proprietary formulas for beauty products. Whatever that asset is, it should be the core mission of a threat intelligence team to protect it; and all actions the team takes should derive from that mission. This is why it is so important to have senior management and the board of directors involved in creating a threat intelligence program, it helps to ensure the goals of the threat analysts align perfectly with the goals of the organization at large.

NOTE

What Is a Paste Site?

A "paste site" is a website that is used to share plain text documents, usually anonymously, such as code. Some attackers use paste sites to post data they have collected from their attacks, especially if that data will be damaging or embarrassing to the victim organization. The most popular paste sites are Pastebin, Pastie, and Ghostbin, but there are dozens of others that are used by attackers.

It is not just attackers that use paste sites, there are a number of legitimate uses for them as well. Programmers will often paste large snippets of code to paste sites so that others can review the code. Many legitimate users think of paste sites as a safe place to store information, but the truth is that most of these sites are indexed and searchable both from the paste site's search function and through larger search engines. Anything published on a paste site will most likely be viewable by anyone on the Internet.

That doesn't mean that threat intelligence collection and analysis should only involve the core mission. What it does mean is that all intelligence requirements should stem from the core mission and support the core mission in some way. In other words, a threat intelligence program should start with a core mission and then ask the questions that will help support that mission.

One of the purposes of threat intelligence is to provide answers to questions that the leaders of an organization have. These questions are more rightfully referred to as intelligence requirements. There are two military definitions of intelligence requirements. The first is: *Any subject, general or specific, upon which there is a need for the collection of information, or the production of intelligence.* The second definition is: *A requirement for intelligence to fill a gap in the command's knowledge or understanding of the operational environment or threat forces.*[2]

The first definition focuses on longer-term strategic intelligence, while the second definition is more immediate and revolves around

2 Headquarters, Department of the Army. "Joint Intelligence (JP 2-0)" (*http://www.dtic.mil/doctrine/new_pubs/jp2_0.pdf*), 2013.

tactical intelligence. Ultimately, in the realm of threat intelligence, the purpose of both types of intelligence requirements is to answer questions that leaders in the organization have about threats to the organization.

Understanding the Threat

Now that that the core mission has been defined, the threat intelligence team has to understand what the threats are to those assets which are vital to the mission. This may seem counterintuitive; after all, most people think that one of the purposes of threat intelligence is to inform organizations of threats. But that's not exactly correct: good threat intelligence from a third party will inform an organization about specific threats, for example new vulnerabilities in a database or new techniques used to gain access to an organization, but it can't determine what an organization sees as its threats. However, good third-party providers of threat intelligence will somewhat tailor intelligence to the specific needs of their customer (there are some caveats to this that will be discussed in Chapter 2).

To illustrate this idea more clearly let's take the example of the customer database mentioned above. There are a number of potential threats against against a customer database, a few of these threats include:

1. Vulnerabilities in the database software itself
2. The risk of attackers accessing the data and selling it
3. The risk of an employee accessing the database and taking that information to a competitor.

By first understanding and outlining the threats to the most valuable assets to an organization a threat intelligence team can start to build requirements that need to be satisfied both internally and externally.

Externally, the threat intelligence team can ask their providers for information about new vulnerabilities in the database software. They can also ask what are the most popular underground forums or carding sites where data similar to the organization's is being sold, and if the provider will monitor those forums for mentions of the organization's name. Even beyond that, the organization can ask which attack groups are targeting this type of data and what their tactics, techniques, and protocols (TTPs) are. Knowing the methods

the attackers use might help an organization distinguish between an external attack and an insider threat.

These are just some examples of external-facing questions threat intelligence teams can ask. However, not all of the questions are external facing, there are also a lot of questions that the threat intelligence team needs to ask internally (more on this next).

Collecting the Data

Getting answers to the internal questions can sometimes be more difficult than getting answers to the external questions. After all, an organization pays security vendors and threat intelligence providers for that type of information, so they are incentivized to respond quickly and accurately. On the other hand, there are often years of rivalries between departments and groups within an organization. This can make data sharing difficult. But knowing the mission means that intelligence teams can focus on collecting the data needed to carry out that mission. That means having the full support of the organization is critical. Collecting the necessary data may require access to systems to which security teams may not have had previous access.

Being able to answer those internal questions is just as important as answering the external questions. In fact, it is necessary to have a thorough understanding of the internal functions of an organization before the data collected from external providers can become true threat intelligence.

Log collection should not be the first step in internal data collection. The first instinct many threat intelligence teams have when they think about data collection is log collection. At an operational level this makes sense—intelligence analysts want more data and want systems that can consume as much data as possible. Logs are an excellent source of data (log collection will be discussed in more detail in Chapter 2).

Instead, threat intelligence teams should start by thinking strategically within the organization. This builds on the knowledge collected during the process of understanding the threats and expanding outward to learn more about different business units within the organization and their processes. It is not just processes that need to be understood; intelligence teams must understand data flow throughout the organization. For example, knowing that programs

A, B, and C feed into database X and that cloud service Z pulls from that database is also an important part of the assessment process. This threat assessment process should always have the goal of understanding the potential threats associated with these relevant organizational processes.

Breadth is an important part of any threat intelligence program. A threat intelligence team should always be looking to broaden their sources of information. Threat intelligence teams should strive to acquire data from as many sources and in as many forms as possible, which is why reaching out to other groups in the organization is so important.

Collaboration is Key

Over the years, information security teams and, by extension, threat intelligence teams have gotten a reputation for being the team of "No!" Because security teams are steeped in the world of malware and exploits, they are quick to see the vulnerabilities in activities in which other parts of the organization engage. They are also quick to try to stop those activities.

While the intentions here are good, they can often lead to an organization being less secure. If other groups think that security only serves as a blocker, they will stop reaching out to security teams before deploying a new capability.

When determining a core mission, the process of internal collection has to be collaborative and should be used as an opportunity to build a working relationship between organizations that involves respect and trust. So, rather than saying, "You want to start an external blog using a WordPress installation with 20 unpatched plug-ins? No! No! No!" It is more productive to say, "There are some potential security issues with WordPress. We have been recommending [platform] to other parts of the organization that want to set up external-facing blogs." Or, "If WordPress is the only option that meets your needs, then the following security steps need to be followed, with which we are happy to assist."

The process of collecting and monitoring necessary data becomes a lot easier if the intelligence team is seen as a help rather than a hindrance.

Understanding the processes and workflow of the different teams in the organization helps the threat intelligence team develop a broad series of intelligence requirements. Some of those requirements will be shared with threat intelligence providers. However, many of those requirements will remain internal, and the intelligence team will be able to source the necessary information to satisfy those requirements entirely within data collected by the organization. Sometimes that source will be mined log data, for example. Whether it is from logs or other sources, this will all become part of the threat intelligence cycle discussed in the next chapter.

Case Study: Facebook's Collaboration on a Large Scale

Most organizations think about collaboration within the organization or with a few select partners. However, when your company is worth almost $500 billion and you have 1.2 billion users you have to think about collaboration on a larger scale.

This is the dilemma that the security team at Facebook faces. Because Facebook plays such a vital role in sharing information around the world, in order to successfully secure the organization and their users the security team must work closely with a range of organizations and ingest information from a large number of sources:

> We have made concerted efforts to collaborate with peers both inside the technology sector and in other areas, including governments, journalists and news organizations, and together we will develop the work described here to meet new challenges and make additional advances that protect authentic communication online and support strong, informed, and civically engaged communities.[3]

Each of the sources that Facebook works with undoubtedly supplies information in a different format and at different levels of reliability. The security and threat intelligence teams must distill this information in a timely fashion and standardized format to make it actionable to Facebook itself.

3 Weedon, Jen, William Nuland, and Alex Stamos, "Information Operations and Facebook" (*https://fbnewsroomus.files.wordpress.com/2017/04/facebook-and-information-operations-v1.pdf*), Facebook, April 27, 2017, accessed 19 June 2017.

Summary

Threat intelligence is a complex topic that can be daunting to organizations that are just starting to build out a threat intelligence program. It doesn't have to be though. Starting with the definition outlined in this chapter, it is possible to build out a threat intelligence program that uses relevant, actionable indicators to provide context to threat or a potential threat.

The ultimate purpose of threat intelligence and the threat intelligence team is to protect the core assets of an organization. In order to do that effectively, the threat intelligence team needs to understand the core mission of the organization and what the organization considers to be its most valuable data. All intelligence requirements should stem from that.

While third-party threat intelligence can be invaluable, it doesn't help an organization improve its security unless it can be correlated against data collected within an organization. Some of that internally collected data will be log data, but internal intelligence is more than just log dat—it also involves understanding the processes of other groups within the organization.

CHAPTER 2

The Threat Intelligence Cycle

As established in Chapter 1, threat intelligence is not a data feed. Instead, threat intelligence is a system. Good threat intelligence teams have a process in place that gives them the ability to continuously adjust to new threats and quickly incorporate new data sources into their intelligence process. Almost all threat intelligence organizations use the intelligence cycle model, with some variation in the terms and numbers of phases.

The Intelligence Cycle

The most commonly used threat intelligence model is the intelligence cycle, shown in Figure 2-1, or a variant on this model.

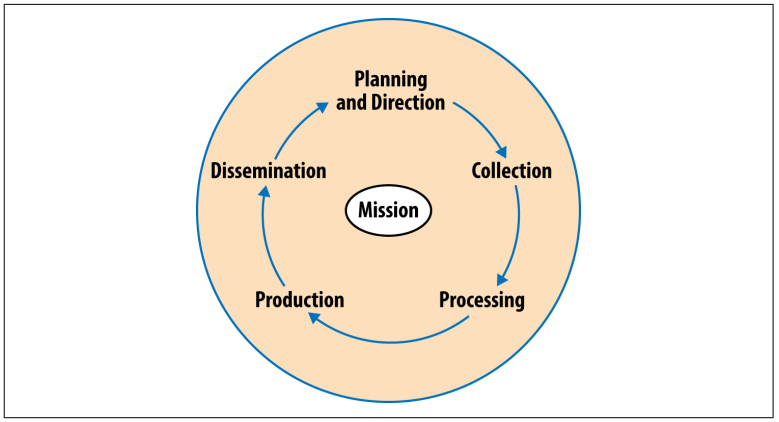

Figure 2-1. The Intelligence Cycle

This is the model that is used by military intelligence, and it consists of five parts, some of which have already been discussed:

- Planning and Direction
- Collection
- Processing
- Production
- Dissemination

At the core of the intelligence cycle is the mission. The five components of the intelligence cycle revolve around helping the organization succeed in its mission. Note that no one part of the intelligence cycle is more important than the other parts. In order for a threat intelligence program to be effective, all components of the threat intelligence cycle have to work equally well.

Intelligence Requirements

The flow of the intelligence cycle allows the threat intelligence team to sift through the incredible amounts of data that is collected by the organization and produce actionable intelligence that makes the security team more effective and improves the security of the organization. The process is truly circular in nature. So, while it may seem obvious to start the intelligence cycle by talking about requirements, as discussed in Chapter 1, requirements generally stem from collected data and analysis of that data. A requirement is a subject about which the threat intelligence team has to collect information or produce reporting. Requirements are often, though not always, requested by someone outside of the threat intelligence organization.

Because the process has to start somewhere, let's start with the requirements phase. Producing good requirements involves asking good questions and properly prioritizing the requirements that are determined based on the responses to those questions. Military intelligence uses the term Priority Intelligence Requirements (PIR) to refer to those requirements that are most critical to the organization, or most time sensitive.

Whether being used with military intelligence nomenclature or developing a different method, every organization has to determine how it is going to prioritize intelligence requirements. The simple fact is that no matter how many people threat intelligence team has

or how large the budget is, there are simply more intelligence requirements than there is time to resolve them all. Therefore, it is important to ensure that the most critical intelligence requirements are resolved first.

High priority intelligence requirements should primarily be those that are most closely tied to the core mission of the organization, however that is not always the case. There may be time-sensitive requirements that need to receive a higher priority simply because there is a smaller window in which to get answers. For example, an organization that is sponsoring FIFA's World Cup may be concerned about both physical and cyber threats surrounding the event, as any such incidents could damage the reputation of their brand. Those intelligence requirements would receive a higher priority because of the fixed time period for their relevance as well as the high profile of the World Cup. Lower priority intelligence requirements tend to be more technical, and ongoing. An example of a lower priority intelligence requirement might include an organization monitoring underground forums for mentions of network blocks assigned to that organization. This is an ongoing and repetitive task that is not time sensitive.

In addition to being prioritized, intelligence requirements must be specific in order to be effective. Crafting effective intelligence requirements is almost an art form and is one of the reasons that good threat intelligence analysts are in such high demand. According the Army's Intelligence Officer's Handbook a good intelligence requirement has three components:[1]

- It asks a single question
- Focuses on a specific fact, event, or activity
- Provides intelligence required to support a single decision.

In other words, broad questions such as "What are all the threats against our organization and what tools are they using?" or "Who is talking about our organization on the underground forums, what are they saying, and are they a threat?" are not effective intelligence requirements. Instead, those questions can be broken down and refined into more close-ended requirements that can be readily sat-

1 U.S. Army, Intelligence Officer's Handbook, FM 34-8-2 (Washington, D.C.: Government Printing Office, 1998), Appendix D.

isfied. Examples of better intelligence requirements based on the broader questions are "What attackers are currently targeting organizations in our sector?" "What are the current TTPs associated with the Syrian Electronic Army?"

The idea of focused requirements may seem counterintuitive, especially considering that most threat intelligence teams are going to be understaffed and underfunded. Focused intelligence requirements, even if there are significantly more of them, actually enable a threat intelligence organization to be more effective. Narrower requirements allow the analysts to get a specific answer without having to guess what the original intent was behind the question. So, while there are more requirements to respond to, the threat intelligence team is able to respond to them faster and in a more complete fashion.

Collection

Collection is the phase of the threat intelligence cycle with which security teams and threat analysts are probably the most familiar. Collection is an inherent part of almost any security program—gathering log data from as many sources within an organization is seen as a critical component to success. Collection is the process of gathering data to fulfill the requirements.

Most organizations rely heavily on their Security Incident and Event Manager (SIEM) to act as the collection point for both security and threat intelligence purposes. There is nothing wrong with that, and SIEMs are very powerful tools. However, there are also limits to SIEM collection—starting with the fact that it is log centric, which can lead to limited thinking when it comes to collection. After all, if an organization's primary collection source is seen primarily as a log aggregator, the collection phase is going to focus on logs.

True threat intelligence collection requires a breadth of sources. Log data is an important source, but it should only be one type of data that is included in the data collection process. It is important for a threat intelligence organization to think about other ways to collect data that may not fit into a traditional log format. For example, there is very valuable intelligence to be gained by mining NetFlow data, but NetFlow data doesn't always fit nicely into a SIEM architecture. Collecting DNS resolution data for passive DNS analysis can also be

very valuable, but again doesn't always fit into a SIEM-based structure.

The good news is that there are a growing number of tools that allow for collection of unusual data sources in their raw form and still allow those sources to be correlated against log data to look for potential threats.

The scope of the data to be collected is dependent on the requirements that are laid out in the planning and requirements phase. That may seem obvious, but it is not always as easy as it sounds. One example of a problem that often arises is that organizations are concerned about the security, both physical and cyber, of remote offices. A reasonable set of requirements to answer this priority might be, "Based on recent examples, what are some of the physical security threats to our offices in the Philippines?" or "What attack groups are active the Philippines at this time?"

To effectively respond to those requirements may involve collecting news stories from local or international papers. The collection point needs to be able to ingest that information along with basic information like the addresses of any offices in the Philippines. Similarly, an organization's threat intelligence provider may share a list of attack groups currently active in the Philippines along with their associated TTPs, but that list probably won't be in log format.

A threat intelligence team should constantly be expanding the idea of what a threat to the organization looks like. It should be trying to find new sources of data to ingest, which requires a platform that can easily adapt as these sources continue to grow.

Collection and Cloud Providers

One of the biggest challenges that faces security and, by extension, threat intelligence teams today is the expanded footprint of the organization. The "network" is no longer defined as "everything behind the firewall." Many critical business functions now reside in data centers around the world and under control of third-party providers.

The reliance on cloud providers can be a challenge when it comes to the collection phase of the threat intelligence cycle. Many providers don't offer the ability to collect log data or provide any insight into activity happening on the remote servers. Often, even provid-

ers that do provide access to log data don't do so in a manner that can be easily integrated with local log sources.

Part of the planning and direction phase of the threat intelligence cycle should include cataloging all of the cloud services used by the organization, determining the security risks associated with the data housed by each provider, and understanding the scope of log data with each provider and the organization's ability to collect that data.

The inability to collect logs may not matter for some providers. For example, not being able to collect log data from a corporate SurveyMonkey account is probably not going to have an impact on organizational security. However, providers with the most sensitive data, such as SalesForce, will require a log collection strategy.

That breadth of intelligence shouldn't apply just to intelligence generated from internal sources. Any third-party threat intelligence providers that an organization uses should have that same goal. Many threat intelligence providers are stuck in the mind-set that indicators are exclusively:

- IP addresses
- Domains
- File hashes

Again, these are useful indicators and will be important for any organization that is just starting the process of building a threat intelligence program. However, there is much more data that a third-party threat intelligence provider can deliver. Other indicators that a third-party provider can deliver include vulnerability information, account numbers associated with their customers (either reward cards or company credit cards), email addresses and subject lines associated with spam or phishing campaigns, proprietary code leaked to the internet, and more. It is important to partner with a third-party provider that can grow with an organization and help an organization expand its capabilities.

Processing

Even though they are distinct components of the threat intelligence cycle, analysis and processing often get lumped together because, in most organizations, these two tasks are carried out in the same plat-

form. Whether an organization is using a SIEM for processing log and network data or aggregating threat intelligence into a Threat Intelligence Platform (TIP), for the most part the threat analyst team is relying on an underlying platform to handle much of the initial processing of collected data.

There is nothing wrong with that approach—in fact, in most cases, relying on anything but a platform that automates processing and correlation doesn't make sense. There are simply too many data sources for a threat intelligence team to be able to analyze manually. An organization of any size is going to have millions of log events and hundreds of thousands of indicators to process each day. It is important to rely on correlation and automation to take the first pass at identifying potential threats.

SIEMs and TIPs

As mentioned earlier, today's threat intelligence teams primarily rely on two platforms for analysis and production. The first is the SIEM, typically those are platforms like ArcSight, LogRhythm, QRadar, or Splunk. The second platform, not as widely used today but gaining market share is the TIP. Some of the most commonly deployed TIPs are Anomali, ThreatConnect, and ThreatQ. Both SIEMs and TIPs have their pluses and minuses, but they both can serve a vital role in the threat intelligence life cycle.

Security teams beginning the process of building out a threat intelligence program generally start with a SIEM. SIEMs are very powerful because they take a morass of data and provide a structure that allows for easy correlation. By dumping log data and data feeds into the same platform with all the data conforming to a singular framework, threat intelligence analysts can easily write correlation rules. In fact, most SIEM platforms offer a large number of out-of-the-box correlation rules that can be used for security, compliance, auditing, and other purposes. So, while a SIEM does require a great deal of "care and feeding" to operate in an efficient manner, a threat intelligence team can get up and running on a SIEM platform in a relatively quick time frame.

However, the SIEM's greatest strength can also be a drawback: the rigid framework of the data that is populated into the SIEM can limit the data types that can be used in the automated analysis. Again, for organizations just starting to build out a threat intelli-

gence program, that is probably OK. But as the threat intelligence program matures, the organization may need to move beyond the SIEM.

TIPs help to fill the gap that is left by some SIEMs, at least when it comes to threat intelligence ingestion. Because TIPs aggregate multiple types of sources of threat intelligence, they tend to be more flexible in their data structures than SIEMs. This allows them to ingest threat intelligence in different forms, often structured and unstructured. Despite the less structured nature of the data that is inputted into the platform, TIPs can produce results in a structured manner that can then be delivered to SIEMs or other platforms to provide context and actionable results.

The strength of the TIP platform is its flexibility. However, TIPs are not designed to ingest large amounts of log data—their focus is strictly on threat intelligence. This means that threat analysts are still required to manage at least two platforms, the SIEM and the TIP.

A Better Solution

A better solution, for more advanced threat intelligence programs, is to bypass the idea of structured data entirely and instead to drop all log and threat intelligence data into a single unstructured platform. Using tools like the Elastic Stack, which allow organizations to include all sorts of data sets, in any format, and analyze those data sets, no matter how disparate they seem, can enable threat intelligence analysts to make connections that are not able to be made by a SIEM.

In order to challenge systems like Elastic Stack, a few SIEM vendors have adopted this type of backend infrastructure, creating a greater flexibility for their customers. Surprisingly, this approach can also allow for greater scalability than the traditional, more structured, SIEM solutions. That being said, working with unstructured data usually involves more work by the threat intelligence team or an organization's Security Operations Center (SOC), as often queries will need to be built from scratch and correlation rules will need to be developed by the threat analyst team. An Elastic Stack solution, or other similar solutions, can take significantly longer to get up and running than a traditional SIEM and may require more adjustments over time.

That is not to say that this type of flexibility is not worth pursuing. Some of the largest threat intelligence programs in the world use this type of solution for threat analysis. It is simply important to be aware of the challenges and time commitment that is often required to make these types of solutions work.

Don't Forget About Vulnerability Analysis

Indicators are what is "sexy" in the world of threat intelligence. But IP addresses, domains, and file hashes were not always the focus of threat intelligence. The earliest exposure many cybersecurity professionals had to threat intelligence involved vulnerability intelligence. Having intelligence around new vulnerabilities helps compliance teams prioritize patching and makes Security Operations Center (SOC) personnel aware of potential threats for which to monitor.

Vulnerability intelligence is still very important, but it has taken on a new dimension over the last few years. Vulnerability intelligence used to rely primarily on reporting around well-known vulnerabilities, such as those presented in the National Vulnerability Database (NVD). This reporting was fed into a vulnerability scanner, such as Tenable or Qualys, to match against internal network scans. This type of automated analysis not only allows for prioritization or patching but is important for compliance reporting.

However, as the underground market for vulnerabilities has expanded and attack groups have grown more sophisticated, that model has proven insufficient. Organizations are relying more heavily on reporting of new vulnerabilities before they are well-published. So-called "zero-day" exploits play a more prominent role in modern vulnerability intelligence, which doesn't always lend itself to structured reporting.

Vulnerability intelligence is still a critical component of any threat intelligence program but, as with other forms of threat intelligence, it is evolving.

Analysis Strategy

From a threat intelligence program perspective there are two goals to be achieved from the analysis of collected data. The first goal is to uncover potential security issues and alert the relevant teams about them. The second goal is to create responses to the ongoing intelligence requirements.

Handling immediate security alerts generated by analysis of collected data is generally handled by the Security Operations Center (SOC) analysts (though, in many organizations there is no difference between SOC analysts and threat intelligence analysts). This type of reporting involves correlating third-party threat intelligence against logs generated by security systems looking for matches against common Indicator of Compromise (IOC) types, such as IP addresses, domain names, or file hashes.

This type of reporting is a way to show immediate value from third-party threat intelligence providers. Most SOCs are overwhelmed by the number of alerts being generated by the various security platforms in the organization. Assuming a third-party threat intelligence provider has done a good job of filtering the indictors they are sending to their customers, and they can provide context around those indicators, third-party threat intelligence can help SOC analysts better prioritize alerts and respond to immediate threats more efficiently.

Even the best SOC analysts can only respond to about eight security incidents per day.[2] A SOC of any size will get hundreds of incidents per day, which means that a lot of security incidents will be ignored. Correlating contextualized threat intelligence from a third-party provider against internal logs helps the SOC identify and prioritize critical incidents based on the criticality of the indicators provided by the provider.

2 Goldfarb, Joshua, "Spear Alerting: Improving Efficiency of Security Operations and Incident Response" (*http://www.securityweek.com/spear-alerting-improving-efficiency-security-operations-and-incident-response*), SecurityWeek, 15 Dec 2014, accessed 31 Jan 2017.

Automated Correlation

The importance of automated correlation cannot be overstated in a modern SOC, and it is a critical part of a threat intelligence program. Tying together different aspects of an attack by linking different sources of data through timestamps or common indicators allows threat intelligence teams to extract intelligence in a much faster manner than before these tools were available.

But automated correlation should not be the only focus of the threat intelligence program. This type of response is almost entirely reactive. The other side of the threat intelligence program, creating and responding to intelligence requirements, allows the threat intelligence program to be more proactive.

Clearly defined intelligence priorities, help build a security strategy and can help build both short-term and long-term goals. Depending on how it is structured, the same collection system can be used by both the SOC team (short-term) and threat intelligence analysts (long-term).

For example, a SOC analyst may notice unusual traffic to an external IP address. The traffic doesn't seem malicious on the surface, but it is odd enough that it is worth creating an intelligence requirement for deeper analysis. The threat intelligence team can review collected data to see if that pattern has been repeated elsewhere, or at an earlier time, in the network. The threat intelligence team can also determine if there have been indicators of malicious behavior associated with the unusual traffic pattern. Beyond what is available in the collected systems, threat analysts can reach out to other organizations to see if anyone else has experienced that pattern and if those organizations had uncovered any malicious activity associated with it.

Again, this is where actionable and contextual intelligence from third-party providers is important. If that unusual traffic pattern is associated with malicious activity, a threat intelligence provider should codify it in such a way that the two activities are tied together. That way, when the next SOC analyst notices the pattern, all of the relevant information will be automatically connected, and the analyst can take immediate action.

Security Orchestration

As security organizations have become overwhelmed with alerts, more SOCs are incorporating automation into their process. Security orchestration companies, such as Swimlane, Phantom, and FireEye, offer solutions that automate the process of an alert being generated to a firewall rule being created, or some other common security action. They can even open and close tickets after the changes have been made.

Security orchestration platforms relieve some of the burden on SOC teams by automating the common repetitive tasks that the analysts perform hundreds of times each day.

Production

Up to this point in the intelligence cycle the discussion has revolved around raw data. Even third-party threat intelligence ingested into platforms is just data at this point. The production phase of the threat intelligence is where the raw data becomes threat intelligence. Production is the process of turning the raw collected data into threat intelligence.

Production can take many forms, each with a specific purpose and audience in mind. Traditionally, the production phase involved the creation of reports that could be delivered to customers as part of the dissemination and feedback phase. Report production is still an important part of the threat intelligence cycle and a critical function of the threat intelligence team.

Reports provide a peer-reviewed, in-depth look at a threat to the organization that is generally in response to an intelligence requirement. But production is not limited to just reporting. Processing the raw collected data can result in the production of internal threat lists that can be distributed to other teams within the organization to create firewall rules, domains to block on a proxy, or signatures that can be incorporated into incident response platforms.

NOTE

Bias and Groupthink

The review process is an important part of intelligence production. All humans are subject to biases—the best threat intelligence analysts try to be aware of their biases and to keep those biases from creeping into their reporting. However, even the best analyst can sometimes unwittingly act based on biases.

Having another analyst review reporting before it is released can help to alleviate the risks associated with biases in reporting. Ideally, the reviewer should be someone in a separate group, to avoid any biases associated with groupthink. For large threat intelligence teams broken into smaller groups, having an analyst from one group review reporting from another group should be sufficient. Smaller threat intelligence teams may need to designate reviewers in another group, such as the SOC, in order to ensure that biases and groupthink are kept to a minimum in threat intelligence reporting.

Production is not just creating reports. It is about getting the data in a format that can be easily consumed by the customers of the threat intelligence team in a timely manner and tracking the delivery of these finished pieces of intelligence.

Ticketing Systems

In order for the different stages of the intelligence cycle to function as a whole there must be something that connects them all. There has to be a way to get from the requirements phase through processing, production, and dissemination in a cohesive manner.

The most common way to do this is through a ticketing system, like JIRA, Service Desk, Remedy, or ServiceNow. Ticketing systems generally provide the flexibility for organizations to use them to create and respond to intelligence requirements while also allowing the SOC to use them to report the results of an investigation or put in an IT management request.

Beyond their basic purposes, ticketing systems can also be used by the threat intelligence team as a way to track more in-depth reporting, especially in response to PIRs. Because ticketing systems connect to both SIEMs and TIPs, as well as other security systems,

supporting data from those platforms can be passed directly to the ticketing system. Which means that the response to the PIR can contain the write-up as well as any necessary raw data.

Ticketing systems have a few other advantages that make them attractive in the production process. The first is that they can be used as part of the review process. Generally, when a response to a PIR is written up, it should be reviewed by at least one other analyst before dissemination, and often there is a multi-step review process. All of this can be tracked within the ticketing system.

Secondly, a ticketing system is great for tracking time-sensitive information. When requirements are created, a deadline for response should be included in the ticketing system. The system will then automatically alert the relevant parties as the deadline approaches and even produce dashboards that allow the manager of the threat intelligence team to track the progress of the intelligence requirements.

Finally, going back to the requirements creation process, ticketing systems allow the threat intelligence team to create custom priority levels for tickets. Prioritization is an important part of the intelligence requirements process. Using a ticket system to create custom prioritization levels and then tracking the process of those requirements will ensure that the threat intelligence team is devoting resources to the right priorities and doing so in a timely fashion.

Dissemination

Dissemination is the process of distributing finished threat intelligence to the original customer as well as anyone else in the organization who requires access. Whether it is a finished report or a threat list, it is important to ensure that it is delivered to the original requestor in a timely fashion and any feedback from the customer is processed. This is usually done through a ticketing system, as described above.

Keep in mind that intelligence reporting doesn't necessarily have to be written up inside the ticketing system—the ticketing system can just be used for tracking purposes. In fact, in many cases, it makes more sense to use standard office tools, such as Microsoft Word, to deliver reporting based on requirements.

The reason it makes sense to use widely available office tools is that, if a threat intelligence team is successful, many of its requirements will originate from outside the team. Most organizations are not going to require that executives or members of the board of directors use a ticketing system to generate intelligence requirements or read the resulting reporting. Forcing customers outside of the threat intelligence or SOC teams to use a ticketing system to create intelligence requirements may result in fewer requirements being created and create the perception that the threat intelligence team is less valuable.

NOTE

Customers versus Coworkers

This book uses the term customers to refer to other people in the organization that submit requirements to the threat intelligence team.

Threat intelligence organizations often adopt the nomenclature of referring to the leadership that creates the intelligence requirements as customers. In this case, a customer is not someone external to the larger organization, instead it is a frame of reference for the threat analysts.

The term customer generally implies that there is an obligation to provide a certain level of service, and can serve as a reminder of that expected service.

Threat intelligence teams often set up an email address that people can use to submit intelligence requests. That email address can automatically open a ticket and be used to track intelligence requests and PIRs, even if the ticketing system is invisible to the customer.

Case Study: Facebook and Dissemination

Dissemination within an organization can be challenging enough, but what if the dissemination needs to occur either outside the organization or to an external customer of the of the organization? Facebook runs into this problem often. In addition to protecting Facebook itself, Alex Stamos, the Chief Security Officer of Facebook, and his team must protect the 1.2 billion Facebook users.

The Facebook Security team has a big challenge:

> Facebook has long focused on helping people protect their accounts from compromise. Our security team closely monitors a range of threats to our platform, including bad actors with differing skillsets and missions, in order to defend people on Facebook (and our company) against targeted data collection and account takeover.[3]

Their team is collecting intelligence from a wide variety of sources, as well as doing their own investigative work and finding potential threats to their customers. But how do they warn those customers that their accounts may be compromised? More importantly, how do they warn them before the compromise happens and account for the variety of platforms that customers use to access Facebook?

It really breaks down to five different types of notifications:

1. Even before accounts are potentially compromised Facebook starts by giving their customers the ability to harden the authentication process, including enabling two-factor authentication.
2. If a Facebook user has been targeted, the Facebook Security team will reach out to that user to let them know and offer suggestions for how best to protect their account.
3. The Facebook Security also reaches out to users who have not been targeted, but may be targeted based on ongoing activity by specific malicious actors that the Facebook team may be tracking.
4. If a threat is serious enough, the Facebook team will reach out to the user directly to warn that user of the potential danger.
5. The Facebook Security team will also work with organizations, especially governments, directly—when appropriate—to educate larger groups about the threat they are facing.

Aside from tracking, using a ticketing system has another advantage—it helps to ensure that the responses to these requests are fed back into the data collection process. The responses, when appropriate, should be stored with other collected data and used to help respond to future intelligence requests, in other words, they become part of the intelligence cycle.

3 Weedon, Jen, William Nuland, and Alex Stamos, "Information Operations and Facebook", Facebook, 27 Apr 2017, accessed 19 June 2017.

Marking Sensitive Responses

Most people with a government background are used to the idea classification markings in documents and know the importance of properly marking documents. The same is not always true in the private sector, but there can be sensitivity around responses to intelligence requirements that should be carefully monitored.

For example, if a company is considering an acquisition of a competitor and has tasked the threat intelligence team with several PIRs around the target company the responses should not be widely disseminated.

Since organizations outside of governments generally don't have a classification system, many rely on the Traffic Light Protocol (TLP). The TLP was created in the early 2000s by the National Infrastructure Security Coordination Centre (NISCC) in Great Britain. It is fairly intuitive in that it assigns one of four classification levels to a document: White, Green, Amber, or Red. Documents with a TLP of White can be widely released, while documents marked TLP Red can be released only to the person or persons who submitted the initial intelligence request.

A system like this is easy to track and provides a level of protection to intelligence requests without being overly complicated to implement.

When a response to an intelligence requirement has been sent to the original requestor, it is important to reach out to requestor to ensure that the request was satisfied and there are no follow-up requirements. This is the part of the intelligence cycle that is most often ignored by threat intelligence teams. Organizations are overworked and already have too many intelligence requirements to answer in a timely fashion. It is easier to just assume that if there are any questions or follow-up the customer will reach out. But that is not always the case—the customer may be confused by a response but not want to reach out, or the customer just gets busy and doesn't have time to review.

Reaching out to ensure that the requirement was properly responded to often generates additional requirements and can help improve processes within the threat intelligence team. Especially

when a customer is confused about language or terms used in the response. Responding to that feedback can help the threat intelligence team become more precise or get better at explaining complex technical terms.

Tracking intelligence requirements through the entire cycle, from creation to dissemination and feedback, is important. It is also important to create dashboards documenting the overall success of the threat intelligence team. The idea of creating dashboards for intelligence requirements is anathema to many threat intelligence professionals, but it is a quick and easily distillable way to illustrate the value of the threat intelligence team. Like it or not, security and threat intelligence teams cannot be seen as separate from the business, instead they must be seen as a core of the business and must be able to demonstrate value in terms that those on the business side of the house can understand. Dashboards are one way to do that.

Of course, the other way for the threat intelligence team to show value is to follow up with their customers during the dissemination and feedback phase of the intelligence cycle. Ensuring that reports are read and understood, getting feedback, and channeling that feedback back into the intelligence cycle also demonstrate the value of the threat intelligence program.

Summary

There are a number of of different threat intelligence models that an organization can follow, but most use the threat intelligence cycle. This cycle involves five phases that all build on each other in a continuous loop: Planning and Direction, Collection, Processing, Production, and Dissemination. All these should revolve around the overall mission of the organization.

Gathering requirements and prioritizing those requirements is an important task for the threat intelligence team. Most people will not know how to properly create intelligence requirements. It is up the threat intelligence team to get to the heart of the questions being asked and create useful intelligence requirements based on those questions.

Collection should be from as many sources, internal and external, as possible and in as many forms as possible. Organizations just starting a threat intelligence program often rely on SIEMs as the collec-

tion point, but it is important to make sure that the SIEM will meet all the collection needs of the organization. At least initially, the collection platform may also serve as the analysis and production platform.

The threat intelligence team doesn't just take in requirements; it also has to produce responses during the dissemination phase. Tracking intelligence requirements through the entire intelligence cycle usually requires a ticketing system of some sort. The ticketing system can track the creation, analysis, and production of the intelligence requirement, but the dissemination is often done using standard office systems, such as Microsoft Word. Even when documents are delivered in a standard format they should be tracked, and it is important to follow up with customers to see if there is any feedback or new intelligence requirements that originated from the reporting.

CHAPTER 3

Applied Threat Intelligence

Chapters 1 and 2 focused on the scope and creation of threat intelligence, but there is also a practical side of threat intelligence that needs to be explored. Many organizations that know they need threat intelligence also know that they are not going to be able to build a full threat intelligence team or build out a threat intelligence process.

In cases like this, the organization may rely almost exclusively on third-party threat intelligence providers to assist with a practical application of threat intelligence in their environment. This is a common scenario for mid-size organizations or those that simply don't know where to start when it comes to threat intelligence. Using the right third-party providers can help an organization experience the benefits of a threat intelligence program without having to invest the, often, millions of dollars involved in building one.

In addition to understanding third-party threat intelligence providers, this chapter will also discuss the Diamond Model of threat intelligence and delve into the roles that strategic, tactical, and operational intelligence play in delivering effective intelligence across an organization.

Relevant Threat Intelligence at All Levels

There are a number of ways that an organization can access third-party threat intelligence. The most common way is through a subscription to a third-party threat intelligence provider. Such

providers offer a wide range of services designed to help their customers gain access to up-to-the-minute threat intelligence. The best threat intelligence providers go beyond just delivering reports and data feeds, they also offer integration and will work with customers to help them build out intelligence services. In fact, threat intelligence providers that offer solutions without knowing specifics around an organization's threats, risks, and maturity level should be treated with some skepticism.

Aside from commercial threat intelligence providers there are also industry—specific Information Sharing and Analysis Centers (ISACs). For the longest time ISACs were considered the domain of the government or financial services industry (FS-ISAC). That is no longer the case: the presence of ISACs has increased greatly over the last few years—there are now ISACs for retail, health care, aviation, and other sectors. ISACs can provide a lot of valuable industry-specific threat intelligence, but they only work as well as the effort an organization is willing to put into them.

Choosing a Threat Intelligence Provider

For organizations that have adopted an intelligence-led security program and are ready to move from open source intelligence to a third-party provider, the challenge becomes how to find the right provider. Every organization has individual needs and priorities, and there are a number of threat intelligence providers that could be a good match. It starts with an honest assessment of the maturity of the organization. Is the organization ready to ingest automated data? Does the security staff have the capacity to handle additional alerts and work to fine-tune incoming intelligence? Is the Incident Response team mature enough to tackle researching adversaries and producing detailed reporting on attacks? Is the organization fully committed to becoming intelligence-led?

If the answer to these questions is yes, or close to yes, the first step is to talk to existing security vendors, especially those that the security or intelligence teams already work with closely. Most security vendors have a threat intelligence offering.

There is a downside to relying on an existing vendor: their offering is constrained by their view of malicious activity, so much of what they offer in threat intelligence may already be part of their other security offerings. It might make sense to reach out to other vendors that specialize in threat intelligence and have a different per-

spective on malicious activity, providing more diverse protection for the organization.

Some third-party threat intelligence providers to consider are:

- Digital Shadows
- Farsight Security
- Intel 471
- iSight Partners (part of FireEye)
- Recorded Future
- SenseCy

Each of these providers have their strengths and weaknesses, but they all have a strong presence in the market and offer services that can cater to different levels of intelligence maturity within an organization.

The same caveats to threat intelligence that were discussed in the first two chapters apply here: a list of indicators without context or relevance and that are not actionable is not threat intelligence. That being said, using third-party threat intelligence providers can be an effective and budget-friendly way to start building out a threat intelligence program.

If a data feed of indicators is not considered threat intelligence, what can third-party providers deliver that constitutes threat intelligence? An organization should look for a threat intelligence provider that delivers three different types of intelligence:

- Strategic Intelligence
- Tactical Intelligence
- Operational Intelligence

But simply providing all three levels of threat intelligence is not sufficient. The intelligence must be delivered in a way that allows an organization to tie all three levels together using something like the Diamond Model for threat intelligence[1] (covered in the next section). A threat intelligence provider that can deliver intelligence in

1 Caltagirone, Sergio, Andrew Pendergast, and Christopher Betz, "Diamond Model of Intrusion Analysis" (*http://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf*), Center for Cyber Threat Intelligence and Threat Research, Hanover, MD, Technical Report ADA586960, 05 July 2013.

such a way enables an organization to pivot from one part of an attack to another with a clear understanding of the bigger picture. A third-party provider can give some level of context to a threat and deliver intelligence in an actionable form. However, the organization must make it relevant by applying the intelligence in a meaningful way.

Case Study: Facebook and Fake Accounts

Facebook has a problem with fake accounts, both from a revenue and an information operations perspective. The Facebook Security team is constantly trying to find new ways to stay ahead of the attack groups creating fake accounts and disabling those accounts before they can cause damage.

The Facebook Security team has to do this without violating the privacy of legitimate Facebook account holders. This means constantly adjusting their strategy and tweaking algorithms to find these accounts:

> We've made recent improvements to recognize these inauthentic accounts more easily by identifying patterns of activity—without assessing account contents themselves. For example, our systems may detect repeated posting of the same content, or aberrations in the volume of content creation.[2]

In other words, identifying fake accounts requires context around the normal account creation process. Having this type of intelligence helps the security team better understand when anomalous behavior is just an unusual account being created versus an account intended to be used for malicious purposes.

Of course, identifying fraudulent accounts is not enough. In order for this intelligence generated by these algorithm tweaks to be useful it must result in action being taken. In this case, that action is the disabling of the suspect accounts. According to an April 2017 security update, these changes allowed Facebook to disable 30,000 accounts.[3]

---

2 Weedon, Jen, William Nuland, and Alex Stamos, "Information Operations and Facebook", Facebook, 27 Apr 2017, accessed 19 June 2017.

3 Shaik, Shabnam, "Improvements in Protecting the Integrity of Activity on Facebook" (*http://bit.ly/protect-activity-fb*), Facebook, 12 Apr 2017, accessed 26 June 2017.

The Diamond Model

The Diamond Model is a way of cataloging all aspects of an intrusion in a manner that allows analysts to easily pivot from one point to another in a single attack or over time. Figure 3-1 provides a high-level illustration of the Diamond Model.

The Diamond Model breaks an attack down into its four major components:

- Adversary: The person or group behind the attack
- Capabilities: The methods of attack, from simple to complex, the adversary has at their disposal to carry out their mission
- Infrastructure: The structures the adversary has in place that can be used to initiate attacks, provide command and control infrastructure, and to exfiltrate stolen data
- Victim: A victim can be something as simple as a targeted host, an organization, or a series of organizations that develop into a pattern.

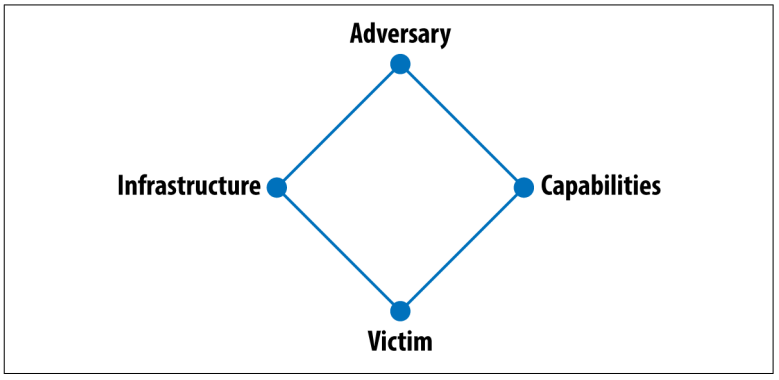

Figure 3-1. The Diamond Model

This model tracks attackers along with their capabilities, the infrastructure they are using, and whom they have targeted or successfully breached. It is an intuitive way of cataloging and understanding the full scope of an attack because it allows for the inclusion of context around the attack. This means that an intelligence analyst can jump from an indicator, such as a domain name, to the group behind that domain and their preferred attack methods.

The Diamond Model also accounts for changes in behavior over time. Attack groups evolve. As they mature they may change tactics,

they may add new zero-day exploits to their arsenal, and their victims will change over time. Even the nature of the group itself may change, members may come and go or the group may change its name.

The Diamond Model is useful because it allows threat analysts to put a testable structure in place when identifying new threats and understanding the nature of those threats in context. It allows analysts to see if current behavior conforms to previous actions before making statements of attribution.

Attribution

Often teams that are setting up a threat intelligence program for the first time will decry the need for attribution. There is a common feeling that it doesn't matter if the attack is coming from a group in China or the Syrian Electronic Army, all that matters is stopping it. With security teams already overwhelmed, who has time to worry about attribution? That type of stuff is for large companies and governments, right?

In an intelligence-led security program attribution doesn't have to be about satellite images and getting the names of specific attackers as well as their favorite vodka. Instead, attribution is about being able to tie information together and tie it to a specific actor/group so the organization can better protect itself against that group. Alternatively, if a successful breach has occurred, attribution helps the incident-response team know which artifacts they should be hunting for.

In the end it doesn't matter if an attack group is named after a bear, a kitten, or it just is assigned the label APT+[Number]. What matters is that there is a rigorous process in place to tie all of the components of an attack together in a way that will be useful to the rest of the organization.

More importantly, the Diamond Model is a way to expand the capabilities of a threat intelligence organization. It helps an organization move from crawling, where this type of complex reporting is primarily derived from outside vendors, to walking, where this type of reporting is built and disseminated in-house.

Strategic Intelligence

Whether an organization uses the Diamond Model to track attack data or another model, such as the Cyber Kill Chain,[4] the goal is still to extract the three types of intelligence. Strategic intelligence is intelligence that helps an organization set policy and understand the "big picture" around threats to the organization. Strategic intelligence tends to be report-based, longer form, intelligence that senior management uses to guide the decision-making process. In the Diamond Model, strategic intelligence sits at the top and bottom of the diamond.

Strategic intelligence allows senior management to direct budget and resources toward systems that will be effective against attacks. For example, say an organization determines that there has been a 30% year-over-year increase in email–attachment based phishing attacks. In addition, the organization receives confirmation from third party-threat intelligence providers that attackers are investing more resources into these types of attacks. Senior management now has strategic intelligence that allows them to allocate budget toward fighting those types of attacks.

Strategic intelligence tries to answer the questions:

- Who may be attacking the organization?
- Why are they attacking?

The best answers to these questions usually involve data from third-party intelligence providers combined with data collected internally.

Strategic intelligence requires that a threat intelligence team has a strong understanding of the threat landscape and can communicate that understanding in an easily digestible way to senior management. There are thousands of adversaries operating at any given time, but not all of them are going to be of concern to an organization. Being able to communicate which groups are a threat and why they are a threat is strategic intelligence. A group does not have to necessarily be specifically targeting an organization in order to be considered a threat. For example, the Dridex team that is behind the Locky ransomware campaign may not be targeting an organization

4 Lockheed Martin, "The Cyber Kill Chain" (*http://www.lockheedmartin.com/us/what-we-do/aerospace-defense/cyber/cyber-kill-chain.html*).

specifically, but they pose a threat to organizations nonetheless. Knowing what the emerging threats are, even those distributed indiscriminately, is an important part of guiding budget and policy and a critical part of strategic intelligence.

It is also important to understand why an organization is being targeted. This connects back to what was discussed in Chapter 1: understanding an organization's most valuable assets. If an organization is targeted by an attacker, what will they be trying to steal?

Of course, attacks aren't always about what an attacker can steal. Many organizations have to worry about so-called "hacktivist" attacks, which are attacks launched against a company for activities in which the company has engaged. Banks are often the target of hacktivist campaigns as are organizations that sponsor controversial events (even if the organization doesn't realize it is controversial at the time). Again, this is why it is important to understand events inside and outside of the company. If the threat intelligence team is not aware that an organization is sponsoring an event that is being targeted by hacktivists, they cannot brief management on the potential risk.

Strategic intelligence should be non-technical and focus on the business impact of threats as well providing recommendations, again at a high level, for protecting against those threats. Strategic intelligence is, by its nature, forward thinking and proactive. It helps an organization put plans in place to protect against upcoming threats.

Tactical Intelligence

Unlike strategic intelligence, tactical intelligence is focused on the capabilities of the attackers. Tactical intelligence sits on the left side of the Diamond Model. It is technical in nature and it is tied to how groups operate.

The purpose of tactical intelligence is to document the TTPs of an adversary in a meaningful way. This documentation can be report based, but it works best when it is integrated into an incident response platform (e.g., Resilient Systems or Resolution1), a ticketing system, or even a TIP. The goal is to make it easy for security teams, especially during a hunt mission, to understand the tools that specific attack groups are using.

For example, knowing that an attacker prefers to use PowerShell to jump from machine to machine once they have breached a network means that incident response teams know to search for PowerShell related artifacts as they are trying to determine the extent of a breach.

On the one hand, tactical intelligence is about tools—knowing the preferred exploit kit of an attacker and knowing what type of loader they prefer to use is important—but tactical intelligence is more than just tools. Tactical intelligence also involves understanding when they work (e.g., do they tend to engage in activity during business hours Monday through Friday GMT +3) and anything that might be useful in identifying an attacker.

Tactical intelligence can also bleed into strategic intelligence at times. For example, if the threat intelligence team learns there is a new zero-day Adobe Flash vulnerability that is being actively exploited in the wild, they will prioritize patching Adobe Flash within the organization over other vulnerability patching. As with other types of threat intelligence, this process can be automated, to an extent. By feeding vulnerability threat intelligence into a Governance Risk and Compliance (GRC) system or directly into a vulnerability scanner, such as Qualys or Tenable, prioritization of patching can be automated. However, it usually requires the threat intelligence team working with the patch management team to prioritize patching vulnerabilities that pose the greatest risk to the organization.

Even the largest threat intelligence teams don't have the resources to gather tactical data on the thousands of threat actors that are active at any given time. Most organizations rely on third-party threat intelligence providers to deliver that type of information. These providers have access to a great deal of resources and teams that do nothing but gather data around different threat actors. They also have the advantage of seeing attacks against all different sectors from many different attack groups.

Tactical intelligence is one area in which threat intelligence providers can provide a depth of coverage that a single organization would have trouble compiling itself.

Operational Intelligence

Operational intelligence is the type of threat intelligence that is most commonly used by the information security community. Opera-

tional intelligence sits on the right side of the Diamond Model and focuses on infrastructure and indicators that can be tied to attackers.

Generally, operational intelligence is fed into SIEMs and TIPs. However, it is also fed into endpoints, proxies, firewalls, and any security tool that can ingest external indicators and make them immediately actionable. Operational intelligence includes practical indicators of compromise (IOCs) such as:

- IP addresses
- Domain names
- File hashes
- Registry entries
- Filenames

Operational intelligence is usually delivered as a feed that is designed to be programmatically ingested into third-party platforms. While this type of intelligence can offer immediately actionable value, if it is not delivered in a way that ties it to tactical and strategic intelligence it can create more work for the SOC and Intel teams that rely on it. As has been emphasized throughout the book: indicators without context and relevance is not threat intelligence.

The ideal application of operational threat intelligence within security platforms is to deliver it in a way that allows SOC analysts to pivot from the IOC to context around why the IOC is bad and what type of attack methods are associated with it.

In a threat intelligence–led security program, a SOC analyst shouldn't just receive an alert from the SIEM that a "bad" IP address showed up in the logs. Instead, the alert should provide information about the type of activity the IP address is associated with (e.g., command and control host, attack infrastructure, etc.) along with other indicators that are tied to that IP address (domain names, file hashes) and the tools the attackers use during a breach as well as to which group all of this information is potentially tied.

But there is a danger in operational intelligence in that it is often fleeting. Unlike tactical and strategic intelligence, which tend to be stable for longer periods, operational intelligence can disappear quickly. For example, an IP address tied to a server that is used for redirection in an attack may get patched and not be "bad" any longer, or a website that was used to deliver ransomware may have been compromised and is now fixed. When relying on third parties

for operational threat intelligence, it is important to understand how those providers age indicators and lower severity ratings over time, because the truth is, some providers do not age or remove indicators in a timely fashion.

The more information a threat intelligence provider offers about their scoring system for indicators and how scores fluctuate over time, the easier it is for threat intelligence teams to determine how much or how little weight to assign to those indicators. Different organizations have different risk profiles. If a threat provider offers threat scores on a scale of 1–10 some organizations will only want to be alerted on anything that has a score of 7 or above; others will want all the data, but may not alert on everything.

This is another area where threat intelligence teams play a crucial role. By understanding the available security resources and the risk tolerance of the organization, the threat intelligence team can work with third-party providers to transform the operational intelligence delivery into something manageable. The truth is that the majority of IPv4 IP addresses have found themselves on someone's watch list at some point, so culling the incoming data makes the threat intelligence team more effective.

Summary

Many organizations do not have the ability to collect and distribute large amounts of threat intelligence internally, so they rely on third-party providers to deliver that intelligence for them. Using a third-party provider can be very effective and it is a great way for an organization just starting a threat intelligence team to build out quickly.

The other advantage of using a third-party provider is their intelligence can be programmatically applied to security systems within the network. However, using a third-party provider does not change the fact that intelligence needs to be delivered at multiple levels: Strategic, Tactical, and Operational. Each level of threat intelligence should allow the threat intelligence team to tie pieces of an attack together to understand the big picture.

Tying the pieces together using something like the Diamond Model allows an organization to deliver the type of threat intelligence each group needs in a way that is useful. Whether it is a SOC analyst,

incident response team, or senior management there are different views of threat intelligence that these teams need, and it is up to the threat intelligence team to provide what they need in the format they need it.

Getting threat intelligence that covers all three phases and then applying that threat intelligence in a way that is actionable, relevant, and provides context will help an organization become intelligence-led.

CHAPTER 4

Case Study: Akamai Technologies

Getting a detailed case study of how well-run organizations use threat intelligence is challenging, because many organizations do not want to give away their "secret sauce," especially if it might open them up to attack. This chapter will provide an overview of the ways in which Akamai Technologies defines and uses threat intelligence to protect not only their organization and employees, but their customers as well. This chapter will help the reader understand how Akamai defines threat intelligence, how they have structured their team, sources of their data collection, and some of their frustrations.

Akamai was selected for this case study because they are well-known and respected in the industry, and they have one of the most sophisticated threat intelligence teams out there. The author of this book has no affiliation with Akamai.

Akamai may not be a household name to most Internet users, but the organizations that rely on Akamai absolutely are. Founded in 1998, Akamai is the world's largest and most trusted cloud delivery platform. Akamai helps some of the world's largest websites distribute traffic to ensure that no single web server is overloaded, that content is served from the closest location possible, and Akamai can even help to distribute regionally specific content.

Akamai also offers Distributed Denial of Service (DDoS) prevention services. Using its massive infrastructure, Akamai can help to prevent even the largest DDoS attack. Whether the attack is application-based or protocol-based, Akamai has a DDoS prevention service.

In addition to Content Delivery Network (CDN) and DDoS services, Akamai offers a number of other security and cloud-based services for their clients.

All this added up to $2.3 billion in revenue in 2016 and is supported by more than 6,600 employees in 60 offices around the world.

Threat Intelligence at Akamai

Akamai differs from most companies in that their sprawling infrastructure generates a lot of threat data that can be distilled into intelligence. Akamai has deployed more than 233,000 servers in 130+ countries that are seeing traffic from 1,600 networks around the world. Very few organizations have a better global view of Internet traffic than Akamai.

Studying this data falls under the purview of Eric Kobrin, Akamai's Director of Security Intelligence, who leads Akamai's Security Intelligence Response Team (SIRT). The SIRT is responsible for distilling the petabytes of data into actionable intelligence that Akamai can use to protect the organization and its customers. The SIRT is managed by Lisa Beegle and Mike Kun, who report to Eric.

The SIRT is one of many teams responsible for intelligence at Akamai. Eric's responsibility is for internal intelligence, but those lines can often be blurred. His team is responsible for ensuring that the customers' services are safe and that the customers themselves are safe. Other teams are responsible for directly driving product direction based on threat research.

Defining Intelligence at Akamai

It is always a good idea to start with a baseline definition of intelligence, since the terms means different things to different organizations. To Eric, threat intelligence is "…non-public information about the Tactics, Techniques, and Procedures (TTPs) of attackers as well as insight into what attacks are coming and who the likely targets of those attacks are."

This is a great definition because it moves beyond the traditional indicators of compromise and focuses on information that is actionable and provides context. Eric provides the following example of something that he considers threat intelligence: A security

researcher finds a vulnerability in software that Akamai uses and reaches out to alert the team at Akamai. That vulnerability may be in software that Akamai built or something they use (e.g. a particular library on which Akamai is very reliant). This type of responsible disclosure means that Akamai can get a patch released or in place before any potential damage occurs to their organization or their customers.

Of course, threat intelligence is not just reported by third parties or collected from network traffic, SIRT also monitors underground forums for information related to upcoming attacks. They look for attackers targeting specific organizations, learn the attacker's timeline, monitor for the attacks, and warn Akamai's customers of the coming attack, as well as monitor for potential collateral damage. As Eric says, given the extent of their reach they are able to use underground data to say, "Here's the attack, it is coming at this time, and we can set up to watch the attack."

Threat Intelligence Sources

As mentioned, Akamai produces a great deal of internal and external threat intelligence as part of their day-to-day operations. But Akamai does not see every part of the Internet, so they also rely on third parties to supplement their threat intelligence. In fact, there are four sources that Akamai uses for threat intelligence:

- Active observations
- Passive observations
- Information sharing: both formal and informal
- Malware analysis

Quite a bit of Akamai's intelligence is informal, and informal information sharing tends to have a lot more impact than formal sharing. When your friends tell you something, they know; and they are not going to waste your time.

Informal sharing doesn't happen automatically, it takes years of cultivating relationships and requires hiring the best analysts, who will have strong relationships with analysts in other organizations. It also means being willing to share information as well, whether that is through informal channels or through more formal channels, such as presenting at security conferences and providing training to other organizations.

Eric encourages his team to present at conferences around the world, and Akamai supports its employees by providing education reimbursement. Building a stronger team and encouraging them to constantly challenge themselves and improve their skills means that Akamai's respect within the industry will continue to grow.

Formal sharing agreements, whether through an ISAC or a third-party threat intelligence provider can also be very valuable. The problem with these sources is that they can be subject to too many false positives. In Eric's view, context is important in these cases. The better the context the third-party provider can present around their intelligence, the easier it is for the analyst team to work with. Another important aspect of third-party threat intelligence relationships is reliability over time. A provider who consistently provides good intelligence with context around the threats is more valuable than a provider who presents a lot of intelligence with no context and too many false positives.

The Akamai Team

Eric's security intelligence team at Akamai has several overlapping skill sets, working together to fuse data to produce a threat intelligence stream that the entire company can use. The different roles are:

- Reverse engineering specialists: The more obfuscated and hidden the malware is, the better. If you give them a problem, they will work on it until they are finished. Sometimes he wonders if they go home.
- Dark web research specialists: Underground experts who spend time talking to bad guys
- Industry research specialists: Well-connected analysts who spend time talking to industry experts
- Threat analysis specialists: They spend time analyzing Akamai data and looking for patterns—reviewing traffic, logs, and other data that helps them understand the threats hiding in the dataflow.
- Writing specialists: One of the most important roles, they are responsible for communicating the findings of the other teams in a clear, concise, and accessible manner.

There are a couple hundred people working in security across Akamai at any given time. The number of people dedicated to intelligence is not as important as the number of resources that Akamai can bring to bear to resolve a problem. But security is everyone's concern, and if there is a security problem, his team will reach out across the company to try to find someone who can help resolve it.

Education is an important role for Eric's group. They train both Akamai employees and other organizations in security architecture and secure development. Eric's group offers several internal education programs, both self-taught and in person. They will also fly out to a customer site to educate their staff. As discussed, Eric's team speaks at industry events as well as Akamai's own EDGE conference, and Eric personally trains new hires on how Akamai approaches security as a company, which includes a call to action akin to "if you see something, say something."

Lack of Standardization Challenges

One frustration that Eric has is the lack of standardization around information sharing across the industry. Transparency and willingness to share information between organizations is still a challenge. He understands why it is done, but one of the things that is harmful is when intelligence is shared under onerous agreements. This happens when an external group or organization provides direct actionable intelligence but dictates with which internal groups it can be shared.

Every organization has their own story around information sharing, but there has to be less differentiation between information sharing agreements.

As Eric says, "Akamai wants the Internet to be a better place, and the work we do in this space is to make the internet safer." He wants people to find more efficient ways to share information with each other with less restriction on how the information can be shared in the defense of the internet. While it is understandable that organizations want to keep sensitive intelligence from being used by sales or marketing teams, overly restrictive sharing agreements makes it hard to get intelligence to the people who need it.

A better standardized sharing and redistribution framework is necessary going forward. Eric expects that there will be a new standard developed in the near future. There almost has to be.

Final Word

Setting up a threat intelligence program from scratch or revamping an existing program is challenging. There are a lot of challenges and pitfalls that can hinder the ability of a good threat intelligence team to be as effective as possible. That shouldn't stop an organization from trying.

The goal of this book was to provide a framework that allows organizations to get started, as well as some practical advice to assist during the launch of a threat intelligence team. The next step is to actually get started—make the leap from being reactionary to threats to getting ahead of them. There are a lot of great resources, outside of this book, to assist in the process. SANS has published several excellent white papers on the topic of threat intelligence. A number of good threat intelligence articles also regularly appear on Dark Reading and SC Media.

Beyond reading material, organizations should not be afraid to reach out to industry-specific groups to find out what other organizations in the same vertical are doing to build threat intelligence programs. If there is no industry-specific group, organizations can talk to their security vendors about what other organizations in their vertical are doing about threat intelligence.

The point is, the best way to improve threat intelligence posture is to start doing something. Even if that something turns out to be the wrong direction, it will make a good lesson learned and the team can move forward. Again, it will take a lot of work, but the payoff in terms of better security is worth the effort.

About the Author

Allan Liska, security architect at Recorded Future, has more than 15 years of experience in the world of cyber security. Mr. Liska has worked both as a security practitioner and an ethical hacker, so he is familiar with both sides of the security aisle and, through his work at Symantec and iSIGHT Partners, has helped countless organizations improve their security posture using more effective intelligence.

In addition to security experience, Mr. Liska also authored the books *The Practice of Network Security*, and *Building an Intelligence-Led Security Program*, and he coauthored the book *DNS Security* and contributed the security-focused chapters to *The Apache Administrators Handbook*.