Best Practices for Applying Threat Intelligence

Recorded Future

brookcourt
Collaborative Solutions

1. Introduction

What Threat Intelligence Is (and Isn't)

Threat intelligence is one of the most talked-about areas of information security today. Recent research conducted by SC Media revealed that 46 percent of security professionals expect threat intelligence to be a very important part of their strategy in 2017. At the same time, vendors, service providers, consultants, and integrators are desperately searching for ways to use threat intelligence, and offering businesses help in applying intelligence on current and emerging cyber threats to protect valuable data and systems.

But when it comes time to choose between these services and products, it can be hard to know where to start. Establishing what types of threat intelligence will prove truly beneficial to your organization will be critical to getting the greatest return on investment in this emerging field.

46%

of security professionals expect
threat intelligence
to be a very important part
of their strategy in 2017

Key Takeaways From This White Paper:

- Understand the important distinction between threat data and intelligence
- Gain insight into the value of different intelligence sources and how to work with them
- Learn about the importance of context to threat intelligence
- Get best practices and case studies for implementing threat intelligence as part of your information security strategy

In the early stages of building a threat intelligence capability, it's vital to develop an understanding of the services, providers, tools, and platforms that are currently available. Unfortunately, as interest in this area of security has increased, the term "threat intelligence" has been adopted and applied in many places where perhaps it doesn't belong. In particular, the terms "data", "information", and "intelligence" are often used interchangeably, which can make drawing distinctions between competing products extremely difficult.

Concerningly, many services provide raw threat *data*, but label it threat *intelligence*.

So what exactly *is* the difference? And how can threat data be processed into information, and (ultimately) intelligence? Simply, the key factors to consider are volume and usability.

Data is typically available in huge volumes, and while it has the potential to become information, it first requires selective extraction, organization, and sometimes analysis and formatting for presentation. For example, a point-of-sale (POS) terminal in a busy retail store collects massive amounts of raw data each day, but that data doesn't become information until it's processed. Details of individual connection requests are an excellent example of data because they're simple statements of fact that aren't open to discussion, and can be processed easily by machines.

Introduction

Information is produced when a series of data points are combined to answer a simple question. In our shopping example, the data collected by our POS terminal may indicate which items each customer buys, when they buy them, and at what price. If these data points are plotted onto a chart to determine regional buying habits and spend, the resulting information is far more valuable than the sum of its parts. Note that although this is a far more useful output than the raw data, it still doesn't directly inform a specific action.

Intelligence takes this process a stage further by interrogating data and information to tell a story (a forecast, for example) that can be used to inform decision making. Crucially, intelligence never answers a simple question, but rather it paints a picture that can be used to help people answer much more complicated questions. To revisit our retail analogy, buying trends information could be used in combination with relevant behavioral psychology research to help make it easier for shoppers to find the items they want. This intelligence doesn't directly answer the question of how to make people buy more, but it does aid in a business decision-making process.

Unsurprisingly, as we progress along the path from data to information to intelligence, the quantity of outputs drops off dramatically while the value of those outputs rises exponentially. The image below, taken from the U.S. Department of Defense's "Joint Publication 2-0: Joint Intelligence" report, does a good job of demonstrating this process.

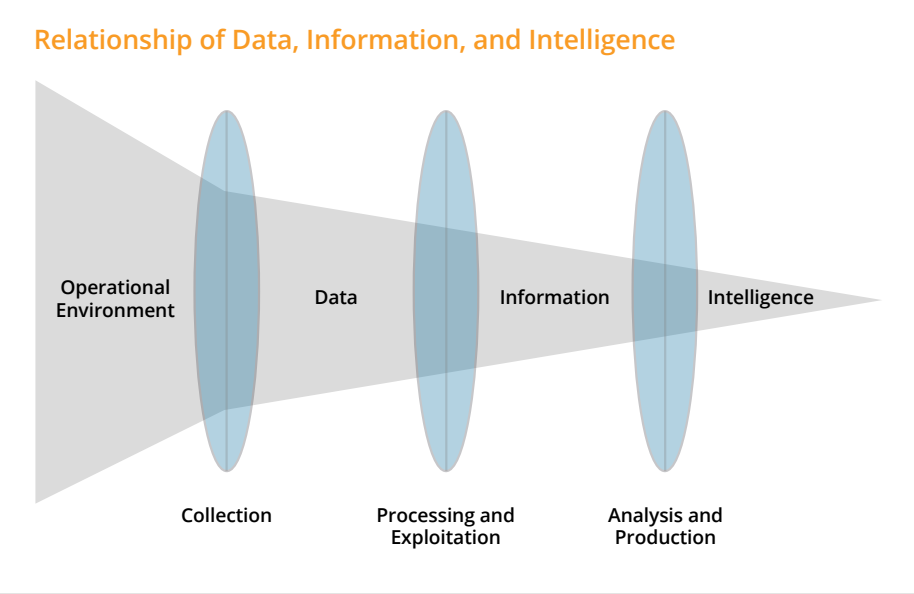

Drowning in Data

In many cases, accessing data from threat feeds is seen as the "on" switch for a threat intelligence capability. Because these tools are often open source and deal with technical indicators, they're frequently touted as a good starting point for developing a strategy.

But what do you do with this volume of data? Ingesting it into your SIEM (security information and event management) solution would seem to be the most obvious answer, but the risk with this approach is that you end up with increasingly huge quantities of data that can't be processed into intelligence. When this happens, so-called "alert fatigue" is almost always the result.

Introduction

Avoiding the Hype

Threat intelligence is a relatively new area of information security, and those with services and technology to offer are keen to ensure organizations understand the benefits they'll see from this type of capability. But as with any emerging technology, the hype occasionally overtakes reality.

The same SC Media survey mentioned earlier provides some insight here, with 43 percent of respondents expecting threat intelligence to offer "early warning of new threats and tactics." While a mature and effective threat intelligence capability should aim to eventually deliver this, there are many other benefits of threat intelligence that can be more realistically achieved in the short term. Businesses need to be pragmatic when it comes to defining how threat intelligence can be applied as part of their whole information security strategy, and consider which phases will align to their most pressing risks, as well as available resources and existing technology.

The Adaptive Security Architecture: 12 Critical Capabilities of Security

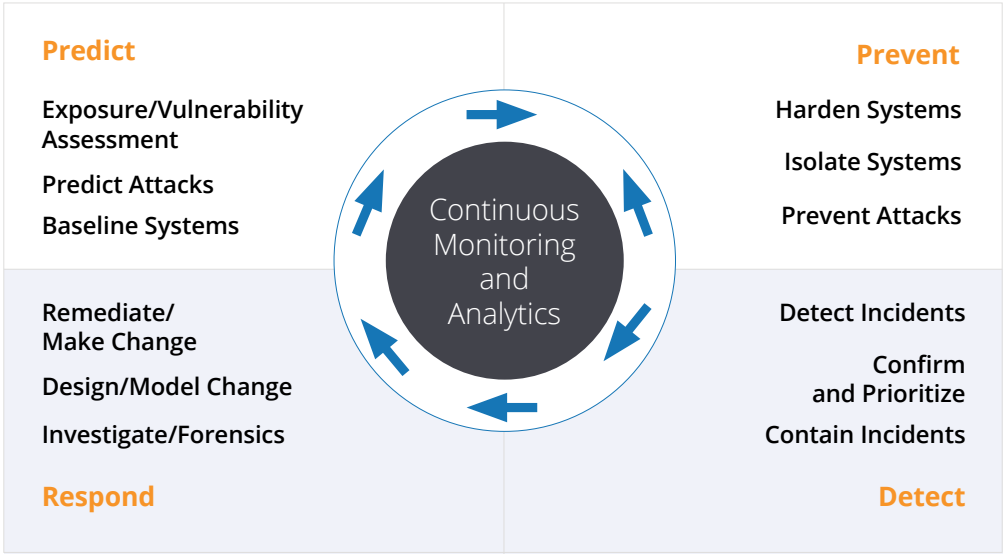

We can see this ethos reflected in Gartner's proposed Adaptive Security Architecture. The diagram above details the operational functions of information security, and their roles in reducing the risk from incoming attacks. But, of course, these are only effective when augmented with the outputs from continuous monitoring and analytics. Historically, you might consider this monitoring and analysis to be from internal log data, but modern threat intelligence has an important role to play in bringing context to the entire information security lifecycle.

2. Sources of Threat Data (and Why They Aren't Intelligence)

As there's exponentially more data than ever before, there are also many more opportunities to derive intelligence from it. But, with so many sources and so much data, this is hard to do manually. Not only do you have to collect this data from the right sources, which can take a great deal of effort to identify, you must also have the resources and expertise to analyze it.

A lot of the time, the term "threat intelligence" is used to describe the sources of all this data, but in reality they're simply origins of data that must be processed before they can be considered intelligence.

To illustrate this point, consider a large stack of reconnaissance photos. Once they have been reviewed by an expert, analyzed, and used to tell a story, they can be considered intelligence. But until then? It's just a large stack of reconnaissance photos. This list, which is far from exhaustive, broadly defines the available sources of threat data. Note, once again, that at this stage we call it threat *data*, not threat *intelligence*.

- **Technical** (e.g., threat lists, spam, malware, malicious infrastructure)
  - This type of data is available in huge quantities, often for free. Due to its binary nature, integrating it with existing security technologies is easy, although a great deal of further analysis will be needed to derive real context. These sources present a high chance of false positives, and results are frequently outdated.
- **Media** (e.g., news, information security sites, vendor research, blogs, vulnerability disclosures)
  - These sources often provide useful indicators of new and emerging threats, but it will prove hard to connect with relevant technical indicators in order to measure genuine risk.
- **Social Media**
  - Undoubtedly, there's masses of potentially useful data on social media channels, but it's hard to determine false positives and misinformation. Typically, you'll find many references to the same threats and tactics, which can place a heavy burden on human analysts.
- **Forums**
  - Because these channels are specifically designed to host relevant discussions, they are a potentially valuable source of threat information. With that said, you'll still need to spend time on collection and analysis to identify what is truly valuable.
- **Dark Web** (multiple tiers of underground communities)
  - Often the source of very specific tactical and technical threat information, but incredibly hard to access, particularly for the higher tier criminal communities. Additionally, as many of these communities are non-English speaking, language is often a challenge.

Sources of Threat Data (and Why They Aren't Intelligence)

Of course, if your goal is to develop a complete picture of your threat landscape, the only route forward is to combine references from various sources of intelligence. But as we've already mentioned, many of the sources above routinely present language barriers, which can prove to be a significant hindrance to effective analysis.

Thankfully, advances in machine learning and natural language processing (NLP) mean that with the right technology, references to threats can be rendered language neutral, and therefore analyzed by humans or machines regardless of the original language used. Perhaps it's even more amazing that we've now reached a point where intelligence solutions that incorporate artificial intelligence (AI) components have successfully learned the *language of threats*, and are able to accurately identify "malicious" terms.

Clearly, this combination of machine learning, NLP, and AI poses a huge opportunity for organizations looking to incorporate threat intelligence. The reduction of analyst workload and removal of language barriers in particular are hugely beneficial, and when combined with the ability to consider multiple data and information sources concurrently to produce genuine threat *intelligence*, it becomes far easier to build a comprehensible map of the threat landscape.

3. The Threat Intelligence Balancing Act: Time vs. Context

As with every aspect of security, the challenge is to find balance. When considering how to access and apply threat intelligence, your two key concerns will be **time** and **context**.

Time to discovery of a security incident was examined in some detail in the recently published 2017 Verizon Data Breach Report, which highlighted that rapid discovery of an attack will substantially reduce the risk of an eventual data breach.

In the report, they use the example of outbound traffic back to a command and control (C2) server. In this particular scenario, threat intelligence is used to provide evidence that the C2 infrastructure in question is malicious, which must be confirmed before action can be taken to block it.

If you currently have no access to threat intelligence, you'll be entirely reliant on your security solutions to identify this threat and protect your organization. By contrast, if you rely on a list of threat data, you may be able to make a fast decision if the IP address or domain appears, but you'll have no context. Under these circumstances, you can't know if you're looking at a false positive, or old, inaccurate data.

It's at this point in the process that time becomes a factor. To find the context you need, you'll need to explore available sources to find relevant and timely references that confirm the infrastructure as malicious. And throughout this process, you'll be grappling with multiple tools, numerous sources, and encountering varying terminology in different languages. Unsurprisingly, this process can prove highly time consuming, and the longer it takes the wider your window of risk grows.

"Alert fatigue" is also a factor here, as security operations staff are forced to deal with a constantly increasing mountain of alerts. A survey from the Cloud Security Alliance highlighted that 40 percent of analysts don't have the intelligence necessary to investigate alerts, and more than a third regularly ignore alerts due to the number of false positives.

4. Best Practices for Utilizing Threat Intelligence

It's human nature to assume that more of something is always better than less. But as we've already seen, the enormous mass of data points provided by the typical threat feed can often lead to nothing more than "alert fatigue."

It's a shame, that so many organizations view pure threat feeds as their best opportunity to "get started" with threat intelligence. Of course, the reality is that if producing valuable, contextualized threat intelligence is your goal, there's a need for massive quantities of incoming data from varied sources somewhere in the process. But, and this is vital, **there's no need for human analysts to ever see it.**

In our opinion, rather than simply collating massive streams of threat data, a successful implementation of threat intelligence should provide analysts with *only* the intelligence they need to make proactive and reactive security decisions.

When defining a strategy to implement threat intelligence, it's far better **not** to start out by investigating which technologies or vendors are available. As we have already highlighted, there are scores of sources of threat data, as well as many intelligence and threat feed providers.

Instead, begin by considering three key questions:

1. Do you understand your greatest risk?
2. Which areas of your information security strategy have you already invested in and do you plan to invest in further?
3. How will human resources and capability impact how you implement any strategy?

In the introduction to this white paper we discussed the importance of threat intelligence as the centerpiece of an adaptive information security architecture. By answering the questions above, it becomes much easier to define which operational areas threat intelligence can usefully influence, and in turn how your overall security profile can be enhanced. It's also valuable to consider how your implementation of threat intelligence has the potential to result in greater efficiency, and more effective use of available resources.

The diagram on the next page shows a suggested pathway for implementing threat intelligence, although naturally your own journey will be highly dependent on how you answered the questions above. The examples given here are by no means exhaustive, and there are a multitude of ways to make threat intelligence work in the context of your organization's individual security strategy and architecture. With that said, the following case studies and suggested applications for threat intelligence should give you a strong indication of where you can reap benefits.

Best Practices for Utilizing Threat Intelligence

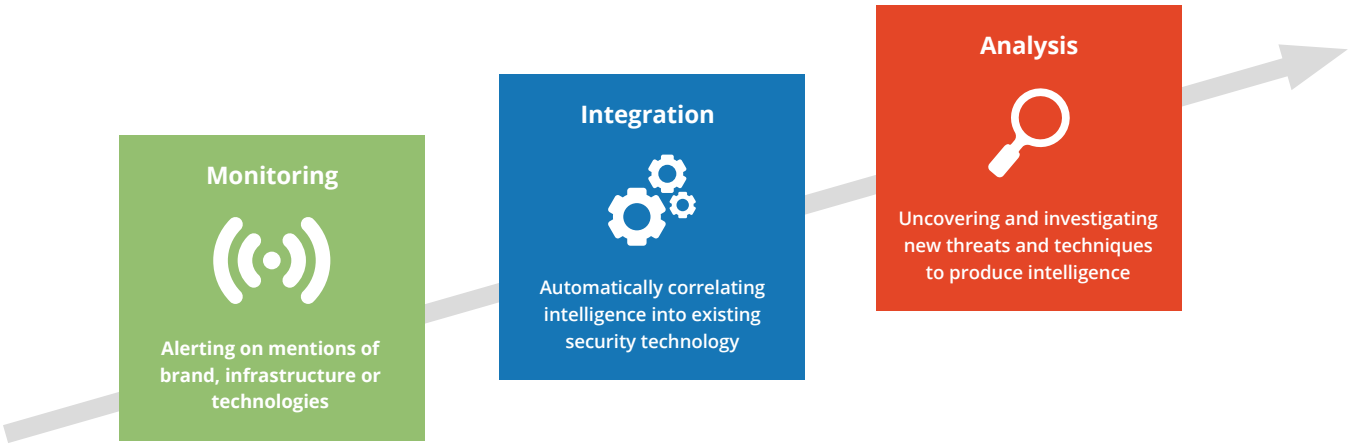

Monitoring

As you start to build your threat intelligence capability, it's likely you won't have the relevant expertise or time resource necessary to support proactive analysis of threat intelligence. Nevertheless, you can still gain significant advantages, and start to collect information from different sources by monitoring them for content that's relevant to your business, and responding as part of your information security strategy.

The types of intelligence you can uncover with this approach include:

› Leaked corporate credentials, data, and code.

› Visibility of new vulnerabilities.

› Threat trends that highlight potential new risks.

Case Study: Leaked Credentials

As recent research in Verizon's 2017 Data Breach Investigations Report highlights, "81 percent of hacking-related breaches leveraged either stolen and/or weak passwords."

Most every day "script kiddies," hacktivists, and cyber criminals upload massive caches of usernames and passwords to paste sites and the dark web, or make them available for sale on underground marketplaces. These dumps typically include corporate email addresses and passwords found when third-party websites are exploited through SQL injection or other weaknesses.

In these cases, internal security measures often prove ineffective, as around 60% of users admit to reusing passwords to access third-party websites and other IT resources. And with many organizations still not using multi-factor authentication (MFA), these exposures pose a truly significant risk.

Best Practices for Utilizing Threat Intelligence – Monitoring

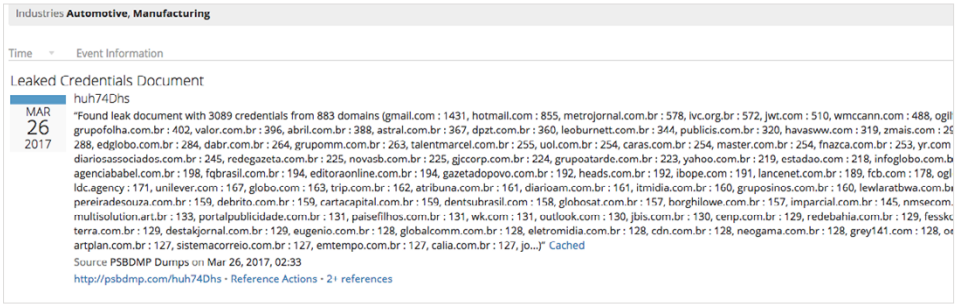

Leaked credential alert in Recorded Future.

The risks of corporate email and password reuse are clearly illustrated in this case from March 2017.

Recorded Future has been configured to alert when a specific set of corporate domains are seen in connection with leaked credentials. We can see from the upload to a paste site that it seems as though a number of Brazil-based websites have been hacked. The database of the website's users has been dumped including email addresses, hashed passwords, and names.

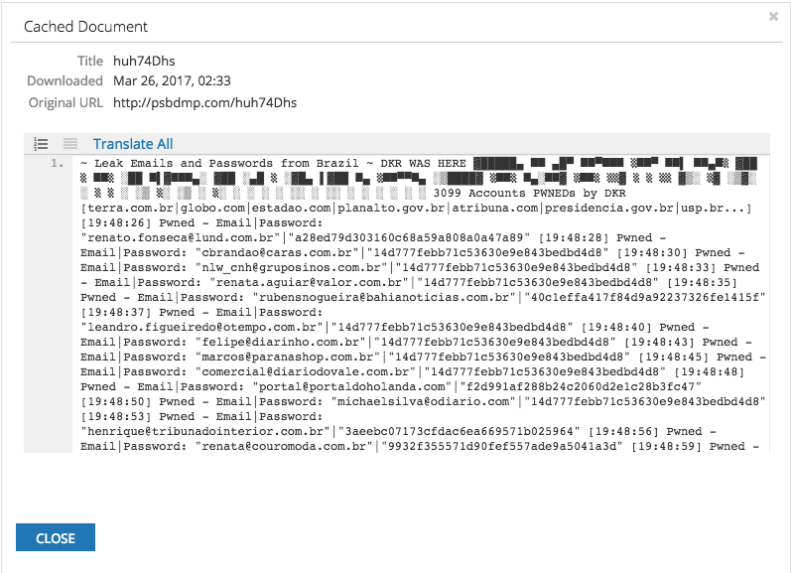

Cached paste site content in Recorded Future.

In this list of over 3000 credentials, some clearly belong to a large number of wide-ranging corporate users, and these leaked credentials open a door for attackers looking to target that organization with spear phishing or other forms of social engineering.

Monitoring external sources for this type of intelligence will dramatically increase your visibility, not just in terms of uncovering leaked credentials, but also potential breaches of corporate data and code.

Best Practices for Utilizing Threat Intelligence – Integration

Integration

If you've invested significantly in security operations and supporting technologies, there are a number of ways you can utilize external threat intelligence to help combat "alert fatigue" within your organization. This intelligence can provide a great deal of context to the indicators you're seeing from internal sources, which in turn can bring significant advantages including:

- Much faster identification of which alerts matter.
- Enriched intelligence on uncovered indicators.
- More context from more sources than technical threat feeds.

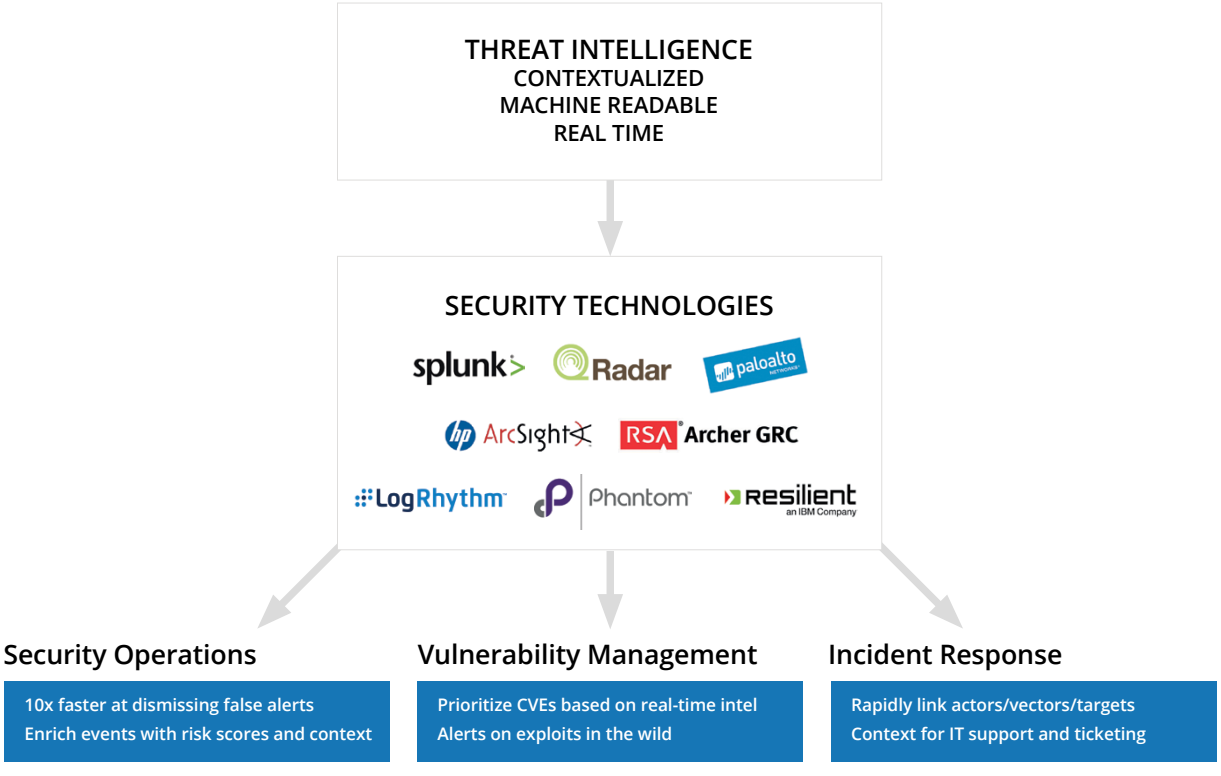

Integration of threat intelligence with security technology

Best Practices for Utilizing Threat Intelligence – Integration

Lab Test: Measuring Productivity Gains From Threat Intelligence

Naturally, you want threat intelligence to add tangible and quantifiable value to your organization's security. As a provider of threat intelligence, we strive to provide measurable benefits to our customers, who in turn have reported back some highly impressive results.

For example, one customer went on record to say that Recorded Future helped reduced the amount of malicious traffic entering their network by 63 percent.

Inspired by anecdotal feedback from our customers, we commissioned a lab test to be conducted by Codis Technologies, an information security consulting firm specializing in incident detection, incident recognition, and process automation. The test measured the quantifiable value (in terms of productivity and security) that a security operations center (SOC) analyst can gain from integrating Recorded Future with a SIEM solution.

The results were conclusive. In a controlled environment, one SOC analyst experienced a 10 times gain in productivity after Recorded Future's real-time threat intelligence was integrated with a SIEM solution.

Best Practices for Utilizing Threat Intelligence – Analysis

Analysis

As your threat intelligence function begins to mature, there's no doubt you'll seek ways to proactively identify emerging threats, and more closely examine the trends that pose risks to your industry, competitors, vendors, and supply chain. And with this level of threat intelligence capability, you'll be gathering the kind of insight that will not only uncover new threats and risks, but also show strategic value. For instance:

› Intelligence that informs your entire security strategy.

› Uncovering new threats, methods, and exploits.

› Analyzing threat trends related to your industry.

Throughout the development of your threat intelligence capability, this should be your aim. After all, while it's incredibly valuable to be able to identify malicious traffic (for example) and respond instantaneously, the true value of threat intelligence can only be realized if you take a more strategic view.

Earlier in this white paper, we asked you to consider three questions:

1. Do you understand your greatest risk?

2. Which areas of your information security strategy have you already invested in and do you plan to invest in further?

3. How will human resources and capability impact how you implement any strategy?

Now, with the benefit of highly contextualized threat intelligence, you can continue to consider these questions as you expand the scope of your security operations. Naturally, over time, your priorities may well change, particularly if your organization or industry is growing or changing. But with a powerful threat intelligence capability, you'll be uniquely positioned to adapt your information security strategy and operations in line with the needs of your organization.

Vulnerability Visibility

A lot of the time, organizations take a volumetric approach to security, particularly when it comes to addressing vulnerabilities. And of course, without threat intelligence to inform your strategy, it only makes sense to prioritize vulnerabilities based on the number of systems susceptible.

But with a strong threat intelligence program, which provides analysis of vulnerabilities from a wide breadth of available sources, you'll be able to take a much more strategic and risk-based approach. Instead of painting by numbers, you can query a range of sources and be alerted to the specific indicators that increase the risk of a CVE being exploited.

Best Practices for Utilizing Threat Intelligence – Analysis

But since vulnerability disclosure and reporting is sporadic (to say the least), what are the right sources to gather this intelligence from?

Data from Recorded Future reveals that 75 percent of disclosed vulnerabilities since the beginning of 2016 appear on web and social media sites an average of seven days **before** official reporting channels. And as references to disclosed vulnerabilities increase, so too does the likelihood of exploitation.

The nature of sources also becomes a factor here. An uptick in references on criminal forums or dark web communities will also contribute to an increased risk score, as threat actors begin to discuss and share methods for exploit. The risk will increase yet again when indicators show the vulnerability becoming part of an exploit kit.

A recent real world example of this process involves a vulnerability in a tool called "Active Management Technology," which is used in some intelligence products. Our Insikt Group analyzed the available intelligence:

"Recorded Future had seen a sharp increase in content regarding the recent exploit CVE-2017-5689 (SA-00075) nicknamed "Silent Bob is Silent," which includes proof of concept (POC) code. The flaw was originally discovered by researchers at Embedi in mid-February of 2017, and they recently released the details in their published white paper."

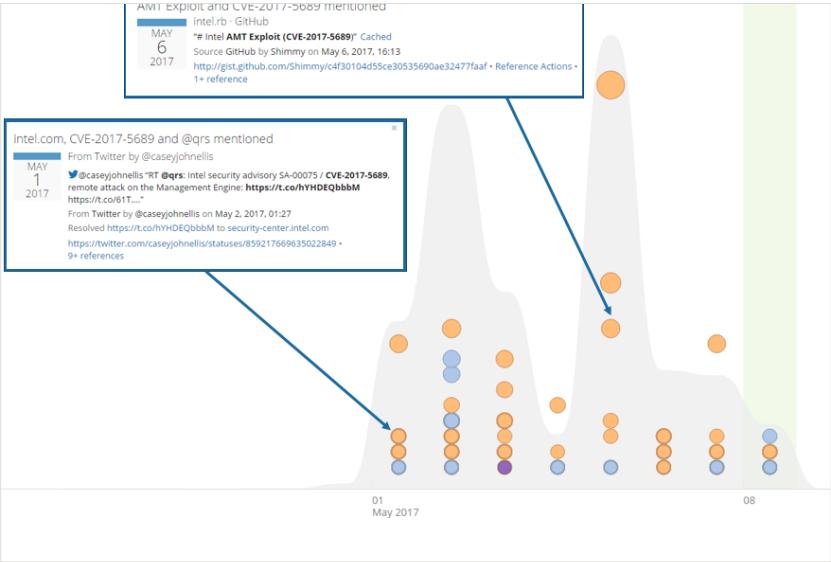

Timeline of CVE-2017-5689 in Recorded Future.

Best Practices for Utilizing Threat Intelligence – Analysis

In the image above, you'll note that as time goes on, references to the exploit quickly increase in volume. During the time period displayed, the CVE was added to a scanner, which enabled threat actors to easily identify vulnerable systems. Following this, the proof-of-concept (POC) code was added to GitHub, which quickly led to the exploit being actively discussed in criminal forums and dark web communities:

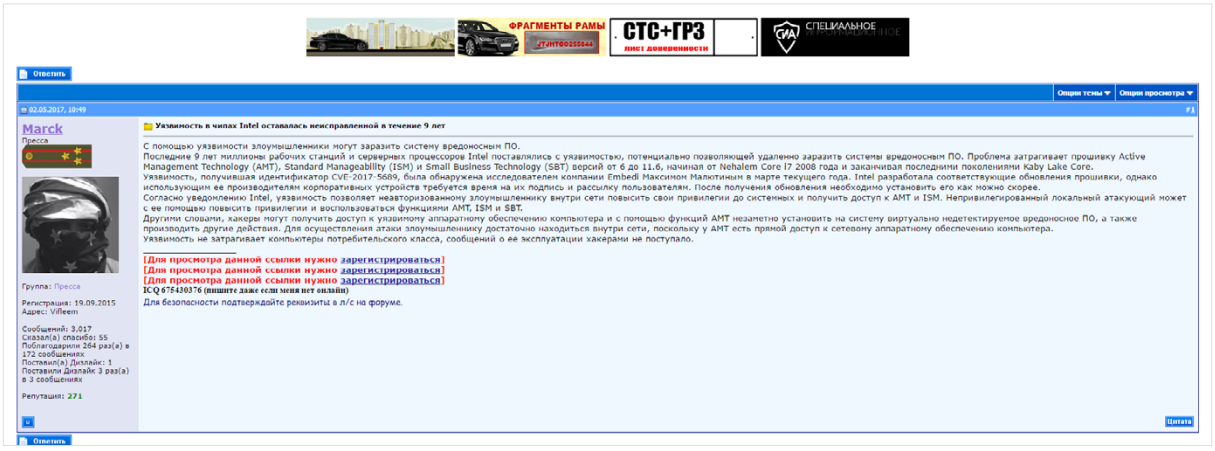

Reference to the POC exploit for CVE-2017-5689 on a criminal forum.

Clearly, having this type of intelligence makes the task of prioritizing vulnerabilities far simpler and more powerful. After all, no matter how few of your assets it might impact, if an exploit is being actively discussed on dark web forums, and is being weaponized at an alarming rate (as it was in this case), it should immediately jump to the top of your remediation priority list.

And this process is in no way unique to vulnerability management. With a powerful threat intelligence capability, this level of content can constantly be gathered, analyzed, and used to inform a risk-based information security strategy. You'll be in a position to identify the most significant threats to your organization at any point in time, and allocate your resources accordingly.

Best Practices for Utilizing Threat Intelligence

Make a Business Case and Measure ROI

At this stage, you may well be convinced of the value contextualized threat intelligence can bring to your organization. Unfortunately, convincing other decision-makers could be a challenge, particularly when it comes to developing a business justification and measuring ROI.

In our earlier example of integrating threat intelligence with existing SIEM technology, the metrics speak for themselves: A clear gain in effective use of analyst time, and a dramatic reduction in the impact of "alert fatigue" in security operations.

When it comes time to develop your own business case, then, this element of threat intelligence almost takes care of itself. Not only can you point to an impartial study demonstrating the value of a powerful threat intelligence facility, you can ultimately utilize the same metrics post implementation to evidence efficiency gains.

But when it comes to monitoring and analysis, ROI may be more difficult to prove. After all, unlike the SIEM test, which could be measured in a purely *quantitative* manner, enhancements in monitoring and analysis yield predominantly *qualitative* benefits.

Simply put, clear thinking from the outset is what enables the development of a powerful threat intelligence capability. And ultimately, whether you choose to focus predominantly on easily trackable metrics or more esoteric benefits will entirely depend on your organization's unique requirements.

To that end, we strongly suggest you follow these two key steps:

1. Know your goals (and prove you can meet them).
    - The more clearly you define and measure the specific areas in which you believe threat intelligence will advance your security profile, the more likely you are to be successful. Don't be afraid of being very specific at the outset in order to ensure you maximize value in just a few key areas.

2. Don't just find a provider, find a partner.
    - To develop your threat intelligence capability, you'll keep adding new goals as you begin to reap the benefits of those laid out initially. A threat intelligence provider who is invested in the success of your efforts, and works with you to uncover new potential use cases, is of far greater value to your organization than a vendor who simply sees your organization as another paycheck.

Best Practices for Utilizing Threat Intelligence

Conclusion

As with most areas of security, the difference between good and bad threat intelligence is a yawning chasm.

In an ideal world, your threat intelligence capability would consistently provide relevant, contextualized alerts that directly inform proactive and reactive security measures. On the flip side, the worst case scenario would be a platform that perpetually flooded security operations staff with outdated, irrelevant, and otherwise unusable alerts, leading to a bad case of the "alert fatigue" blues.

Here's the thing. Most organizations looking to implement threat intelligence for the first time believe that threat *feeds* are the way to get started. To that end, they simply implement a basic threat intelligence platform (TIP), add a few open source feeds, and jump right in. And if you're looking to jump right into the nightmare scenario described above, that's exactly what you should do.

Simply put, there's no faster way to convince your security operations staff that threat intelligence is a waste of time and resources than to force them to use nothing but open source threat feeds. Because as we've already explained, without context threat intelligence is really just threat data, and busy operational staff just don't have the time to manually discount thousands of false positives.

Accept that threat intelligence doesn't *have* to be hugely complicated. If you're clear about your objectives from the start, and you take the time to identify the right technologies and providers to help you reach them, the daily reality of utilizing threat intelligence can be remarkably simple.

In order to achieve this, whatever your implementation of threat intelligence looks like, it must deliver in four key areas:

1. Integrates with and enhances existing technologies.
2. Scours the technical sources but also the open and dark web, for threats, converting foreign language alerts into a useable format.
3. Provides fully contextualized alerts in real time with no false positives.
4. Consistently improves the efficiency and efficacy of your security operations.

Many so-called threat intelligence solutions deliver little more than freely available threat data, which as we've already explained will most likely hinder your security operations more than it helps. So when you're meeting with providers to identify the perfect solution, make sure you challenge them on these points, and ask them to prove any claims they make with hard facts.

And most importantly, take the time to find a provider who will act as a true partner in your threat intelligence journey. After all, they're supposed to be the experts, so who better to help you maximize the value of your threat intelligence capability?

About Recorded Future

Recorded Future delivers threat intelligence powered by machine learning, arming you to significantly lower risk. We enable you to connect the dots to rapidly reveal unknown threats before they impact your business, and empower you to respond to security alerts 10 times faster. Our patented technology automatically collects and analyzes intelligence from technical, open, and dark web sources to deliver radically more context than ever before, updates in real time so intelligence stays relevant, and packages information ready for human analysis or instant integration with your existing security systems.

@RecordedFuture | www.recordedfuture.com

About Brookcourt Solutions Ltd

Brookcourt Solutions Ltd is an award-winning organisation; committed to providing independently sourced, leading edge protective technology solutions. Working together to ensure your organisation stays safe, confident and empowered in today's challenging Cyber Security and networking environment.

We're not just a reseller, we approach every client strategically – understanding the specific business landscape, infrastructure and associated threats.

We are an adept team – we won't throw bodies at a problem – we provide the best people for the job. Our analytical and independent approach to your specific circumstances means we deliver carefully throughout, purpose built solutions to ensure your business is protected and your network is consistent, intelligent and efficient.

@TweetBrookcourt | www.brookcourtsolutions.com

Recorded Future