WHITEPAPER

Demystifying Threat Intelligence

Author Adam Schoeman

October 2014

SECURE:DATA

CONTENTS

Abstract

As with so many other technologies throughout the history of information security, Threat Intelligence has become a buzzword, a term and an idea that seemingly holds the key to defending against tomorrow's attacks. But what is meant by "Threat Intelligence", and how does it differ from Security or Cyber Intelligence? Is it something physical that can be packaged and traded or is it a methodology or framework on how to address the growing threat of attacks? The answers are varied and often interlinked with other buzzwords, further muddying the concept.

These varied views exist because there is no definitive doctrine on what is or is not Threat Intelligence, which has resulted in a proliferation of technologies, ideas and definitions that all fall under the "Threat Intelligence" banner. This paper interrogates the landscape looking for the commonalities and building a reference document onto which the different principles of Threat Intelligence can be plotted and then gauged.

We will look at the top-level research that has already been done into the Threat Intelligence segment and what vocabulary is used to define the core concepts. This will be followed by what has caused the proliferation of Threat Intelligence products into the market. From this, we will see how these concepts have been sculpted into product offerings and where those products fit in the information security problem. Finally, analysis will be done on whether Threat Intelligence, as defined previously in the paper, is being incorporated successfully and effectively into the information security toolbox.

History

Threat Intelligence is not a new concept as intelligence gathering processes have existed within most branches of the military for many years[1]. Broadly speaking, intelligence is the process of gathering raw data, and then transmuting it into useful information. How useful that information is depends on what the original purpose was, but what this does highlight is that intelligence is not simply just a collection of raw data; it requires some aspect of analysis. The practice has already even already been integrated into private corporations with the rise of "business intelligence" and corporate espionage, from as early as 1865[2]. But while the concept of gathering intelligence on threats supersedes the invention of the microchip, the term "Threat Intelligence" is far less mature.

On a backdrop of conflict, intelligence gathering has always related to Threat Intelligence gathering, with the specific mandate dictating how the intelligence was named (as with the case of business intelligence). But this definition proved too one-dimensional for the military when the medium of conflict changed from kinetic to non-kinetic. In 2002, researchers working at West Point, the United States Military Academy[3], published works that looked towards the next generation of defensive technologies as a source of cyber defense, but that would require complex configurations to be effective. "Cyber Intelligence" was heralded as the way to find the information that was needed to make those configurations effective so that they would be relevant to the threat actors that they were meant to defend against. While they favoured the use of Cyber Intelligence opposed to Threat Intelligence (as we do today), it was a choice centred around the time period in which the research was conducted. The focus was on cyber warfare and cyber terrorism as the threat medium and actors to the US Military's connected assets, so it would stand to reason that intelligence gathering around these activities would also wear the 'cyber' moniker. To the US Military the use of 'cyber' simply implied that the intelligence related to their connected assets, opposed to regular kinetically susceptible ones. Once again because of the military's conflict backdrop, the more verbose and correct 'Cyber Threat Intelligence' was shortened to 'Cyber intelligence' because most aspects of military intelligence relates to threats.

Diagram 1: Differences in business and military information gathering

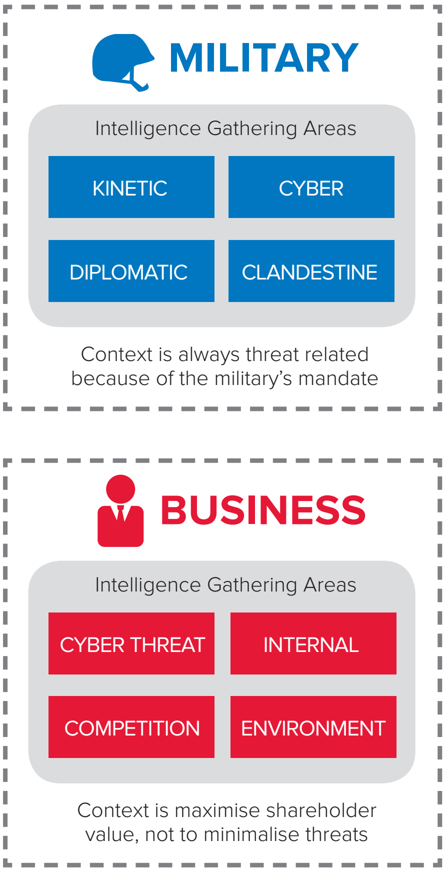
MILITARY
Intelligence Gathering Areas
KINETIC
CYBER
DIPLOMATIC
CLANDESTINE
Context is always threat related because of the military's mandate
BUSINESS
Intelligence Gathering Areas
CYBER THREAT
INTERNAL
COMPETITION
ENVIRONMENT
Context is maximise shareholder value, not to minimalise threats

There does seem to be some degree of logic in the differences between 'Threat' and 'Cyber Threat' Intelligence, if you assume that 'Intelligence' is the gathering of information and the prefix describes what that intelligence relates to. But this becomes opaque as these terms are retrofitted onto the current information security discipline. Information security is positioned uniquely in this regard because it exists in an environment that is both business and conflict orientated. The business does not exist to fend off attacks against its IT infrastructure, but for the defensive security team that is their sole mandate and it is often compared to the military because there is a single medium that is being dealt with. The West Point research had to label their IT specific intelligence gathering as 'cyber intelligence' because it was a new medium for which the military gathered information. But for information security 'cyber' or things relating to IT assets has always been the sole mandate.

Just as the researchers at West Point were looking to solve the problem of complex configurations and proposed Cyber Intelligence as the answer, the rise of advanced persistent threats (APTs) and nation-state sponsored attacks demanded fresh thinking from the information security industry. The response was a focus on intelligence gathering, specifically on the types of adversaries faced by the industry, and their modus operandi. Based on the naming logic explored previously, Threat Intelligence is a fitting description for this as the mandate is assumed to be that of connected assets (therefore it does not need the cyber prefix) and the intelligence being gathered relates specifically to threats facing the business. Threat Intelligence started to gain momentum in 2011, culminating in the February 2013 report by Mandiant[4] which unmasked Unit 61398, a computer hacking division within the Chinese military or the People's Liberation Army (PLA). The introduction of these well funded attackers dragged Threat Intelligence into the spot light as the mitigation strategy.

The appearance of what is believed to be a new or military sponsored threat could explain what kickstarted the rise of Threat Intelligence products. Attackers such as the PLA were seen as a new spectrum of attackers, which is not entirely true since their methods were nothing new, but they had been able to exploit targets which had been viewed as very secure (such as RSA and Lockheed Martin). This demanded a response from teams who were in the business of protecting their enterprise's IT assets.

If the attackers are now believed to be a part of a military outfit, then one possible response is to look to the military for guidance. As shown previously, the military follows a process of intelligence gathering in almost all situations as it tries to understand and monitor potential threats, which could have led to the rapid demand of Threat Intelligence products. Another possible cause is the natural human tendency to seek out information when faced with a decision or response. The conventional approach to decision-making says that the principal ingredients required to make a decision are a set of alternatives, a set of constraints on the choice between the different alternatives and a way to evaluate the strengths and weaknesses of each. To populate the decision-making process a large amount of uncertainty needs to be removed which is normally achieved through information gathering.

The logical naming for intelligence activities that had been established by the military and early business sciences unravels with the discovery of Unit 61398. For whatever reason, the discovery of Unit 61398 created a need for consumable Threat Intelligence in enterprise defenses, and a demand for it to be filled. The market responded with Threat Intelligence products that use the basic idea of intelligence gathering as a response to these sorts of attacks. This has resulted in a proliferation of intelligence products with slight differences in both functionality and prefixes to differentiate each particular offering. Thus we have products that are fundamentally similar, but that wears names such as 'Advanced Cyber Security Threat Intelligence Service'.

Without the logic-based definitions, it is almost impossible to classify a product based on how it is named. The use of prefixes has become a marketing differentiator instead of a functional one, so a different way of evaluating these products is required. One way is too look at what could be conceived as the perfect idea of Threat Intelligence, as designed by some of the industries biggest research firms.

The research industry view

FORRESTER

Industry researchers have been trying to classify what Threat Intelligence is since it came to the forefront. Forrester Research, in a paper titled 'Five Steps To Build An Effective Threat Intelligence Capability', describes Threat Intelligence as a framework rather than a single product or service. In the paper Forrester shows that five distinct focuses need to be combined to harness Threat Intelligence effectively, comprising of:

1. Laying the foundation
2. Establishing buy-in
3. Staffing the team
4. Establish sources
5. Derive Intel

This process is actually based on a military doctrine called the intelligence cycle. The intelligence cycle converts information into intelligence and is generally represented as having these steps:

1. Planning and direction: intelligence requirements are aligned to business requirements
2. Collection: information that meets the requirements is collected
3. Processing: information is prepared to present to analysts by using technology
4. Analysis and Production: here analysts are used to convert information into intelligence
5. Dissemination: intelligence is delivered in timelines that make them actionable

The Forrester five step plan follows the basic idea of the intelligence cycle, in that the intelligence requirements need to be defined and linked to the business requirements. After that, the processes diverge slightly as Forrester addresses two IT specific problems, namely the difficulties in implementing any new IT project and staffing. Establishing buy-in for the project is an exercise in proving that there is a need for Threat Intelligence to the decision-makers of the company, and while it is an important part of any successful programme, it is, strictly-speaking, not unique to Threat Intelligence. The same is true for the third step, which addresses how arduous it is to find staff that possess the right set of skills, which is again a general IT problem. The 'establishing sources' step is where the Forrester Threat Intelligence framework starts to take shape. Here they look at five possible sources of information that can be pulled into the framework, originating from internal, Government-sponsored, industry, public (OSINT) and commercial sources.

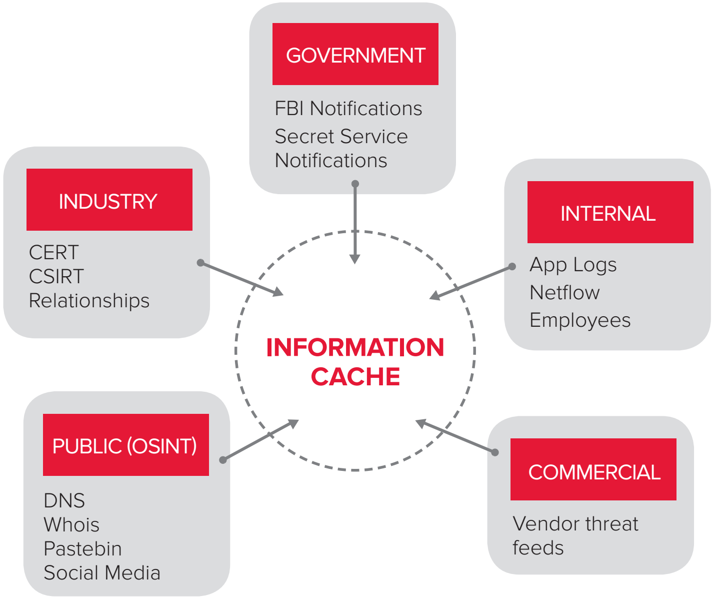

Diagram 2: Establishing Sources

Internal sources relate to any information that is generated by the business itself, in the form of application logs, netflows and data collected from employees, and are normally free of cost and relatively easy to insert, but have not been normalised or classified. Government-sponsored information normally takes the shape of data that a relevant government department has discovered and made public, such as the FBI or Secret Service relaying information on a new botnet that they have uncovered. Industry data is gathered from business partners that operate in the same threat landscape and which have a strategic symbiosis in sharing threat information, such as supply chain businesses. Industry-specific information security organisations have also been established to alert their members of possible threats, such as the The Industrial Control Systems Cyber Emergency Response Team (ICS-CERT). Public information, which is sometimes referred to as open source intelligence (OSINT) is any information that is publicly available to anyone, such as DNS WHOIS results or posts made on public forums such as social media and Paste Bin. Publicly available threat feeds, such as FireEye Track and Openbl.org are also classified into this category. The last source of information is any feed that has been purchased from a specialised vendor such as Dell SecureWorks. This data is normally highly structured, having already been transformed by some analytically driven function. Once the data sources have been defined and pulled into the organisation, step five of the Forrester plan sets out to derive intelligence from those sources. Pieces of data are linked together like evidence, building a case against a certain threat and describing how it can be detected and ultimately mitigated, which is mostly done by analysts.

The Forrester paper does a good job of laying out the lifecycle in a familiar way, and it also highlights some of the difficulties in launching a new IT project (buy-in and staffing) but it does not shed much light on exactly what a Threat Intelligence capability would be trying to achieve. The framework is very open-ended and requires each application of the framework to define what is necessary to generate value within the mandate of the organisation. The 'establishing sources' section is also of the utmost importance because it creates and limits all possible future intelligence derivations. If the information sources that have been chosen are too limiting then so to will the intelligence, but if the chosen information sources are too limited then the intelligence will be too, then more work is required to derive meaningful intelligence. The report is also not very verbose on how intelligence is derived once the final step is invoked, or it simply relies on the ability of the organisation's human assets (analysts) to find the intelligence amongst the data sources. The reliance on analysts to 'find the intelligence' also means that the staffing difficulties mentioned in step three are actually not that generic and is more critical to the Threat Intelligence process. This reliance means that the skills that are housed in the Threat Intelligence analytical division are the single most crucial moving part of the framework, and opens the whole framework up to a series of human resources related risks.

Forrester's idea of Threat Intelligence is then a framework that is constructed around high quality information sources and skilled analysts that together are able to generate actionable intelligence that relates to the organisation's threats.

GARTNER

In a paper published by Gartner, it defines Threat Intelligence as “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”[5]. Unpacking this definition we can say that Gartner sees Threat Intelligence as knowledge that pertains to an asset that is at risk from a new or existing entity. That knowledge needs to contain contextual information about the objects involved (asset, risk, probability and so forth) how it is in danger, or how it will be attacked, indicators that designate that a particular action is being taken against the asset and advice on how to respond that is actionable.

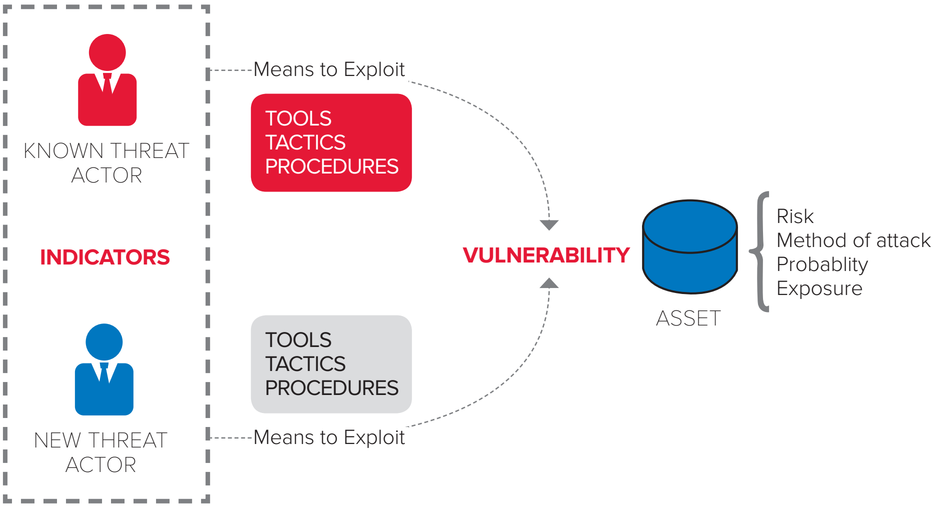

Diagram 3: Unpacking the Gartner definition

The Gartner high level view of Threat Intelligence does create a definition that could potentially encase a single product, as it refers back to a piece of knowledge and a series of criteria that that knowledge needs to fulfil. With this in mind, the possibility does then exist that a black box product could be built that would satisfy all of these points. The difficulty comes in trying to satisfy all of these criteria. Even one of the most crucial criteria, for example ‘knowledge on a hazard to a particular asset’ is difficult to package because it assumes that the knowledge is able to discern between what is and is not a potential threat to an asset. Knowledge that is Threat Intelligence then needs to be filtered based on victim implying that whatever is distributing the Threat Intelligence is aware of the environment where it will be deployed. This is problematic when the knowledge is purchased as a black box offering because it is not feasible for the provider of the knowledge, assuming they are a third party, to understand the organisation’s assets well enough to contextualise threats.

The other option is that the provider of the knowledge expects the asset owner, the organisation, to filter the general knowledge that they provide and apply it to the most relevant assets, but this positioning is at odds with the idea of a purchasable black box product. This dichotomy persists for almost all of the criteria outlined by Gartner. The actionable knowledge is dependent on the defences protecting the asset, how the assets are deployed and maintained and what its intended use is. For example, a recommendation of disabling access to an asset for a period of time based on a threat might be applicable to some situations, but if that asset forms part of a core business value chain then that advice is not actionable.

Gartner’s definition does not initially come across as a framework, but after analysing the rhetoric and criteria it is clear that a support system would need to be put in place to contextualise the knowledge that originates from third parties, not unlike a framework. So while the Gartner definition could relate to a black box product, it is unlikely that a Threat Intelligence system could be deployed that adhered to the definition, that was not part of a framework.

Diagram 4: Aspects of a threat

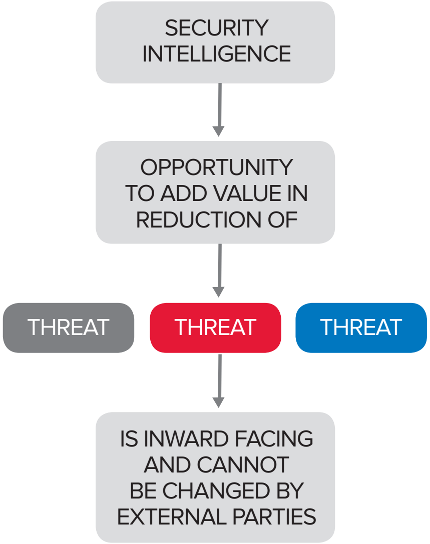
SECURITY INTELLIGENCE
OPPORTUNITY TO ADD VALUE IN REDUCTION OF
THREAT
THREAT
THREAT
IS INWARD FACING AND CANNOT BE CHANGED BY EXTERNAL PARTIES

THE SANS INSTITUTE

While not strictly a research group, The SANS Institute does publish various research-driven reports via its Reading Room project. It is also a leader in the information security training space. While it has not published research that is directly related to defining Threat Intelligence, in 2009 it released a two-part series named 'Security Intelligence: Introduction' where it explored the concept of Security Intelligence. The Institute goes on to equate Security Intelligence as a process of risk reduction, sighting that the Threat aspect of a risk (where a risk is a mixture of a Vulnerability, Threat and Impact) is one of the main places where intelligence can make a difference. The idea of Security Intelligence as a uniformed rollup of potential information sources is appealing because the thinking is that the more information that can be gathered the more informed future decisions would be. But this once again supports the idea that it is not possible to implement a black box security intelligence system because vulnerabilities as a class are internally mutable but not externally. In other words, an external entity does not have the ability to alter the existence of a vulnerability[6]. An external threat feed, will also not be able to change the state of vulnerabilities within an organisation, relegating any external Security Intelligence system to focus on threats.

The research from The SANS Institute follows a similar line of thinking, shifting focus from the greater security risk equation, and zoning in on the threat aspect. It defines a threat as an adversary possessing three things: the intent to be a threat, the opportunity and the capability to act against the asset, which is similar to the criminal law triad of means, motive and opportunity. Intelligence gathering activities would then have to cover the these three aspects to construct an informed view of the threats being faced by an organisation, but the same problem that applies to vulnerabilities, that they are inward looking, applies to the three aspects. To fully understand if a potential adversary as the opportunity, capability or intent to attack an asset, the asset itself needs to be understood. This is again very difficult to achieve from a black box point of view, where the system has no knowledge of the environment where it is deployed, or that its feed is servicing.

Conclusion

Looking through the research as it has been presented in the past, there are many different definitions and models that have been attributed or grouped under the banner of Threat Intelligence. Of the three research pieces that were analysed, Gartner was able to condense Threat Intelligence down to a simple definition, while Forrester pitched the concept as a continuous framework. The SANS Institute's contribution is markedly more vague than Forrester and Gartner, but it is able to further enumerate the pieces that make up Threat Intelligence. But there is one constant undertow that exists from all three of the research sets: Threat Intelligence simply cannot be deployed in a way that adds value as a black box system.

The largest limitation to this sort of deployment is that a threat feed that has been built to be consumed by any end user or system, needs to be generic to scale, and therefore cannot incorporate information that is specific to the point of deployment[7]. As the research has shown, it is in that localised knowledge where contextual Threat Intelligence can be wrought, or rather where Threat Intelligence can be contextualised in a way that makes it relevant to the end user. Without the linkage into local data to provide that contextualisation, an information feed cannot be described at Threat Intelligence.

Products

From a research and theoretical standpoint, it is possible to come to some common understanding as to what is and is not Threat Intelligence. But while there is value in trying to define it from such a vantage point, it doesn't change the fact that there are offerings on that market that have been deployed and that refer to themselves as Threat Intelligence products or services. Instead of comparing individual products to the Threat Intelligence definitions that this paper has established, we will undertake a process of collecting products into groupings that exhibit similar criteria, and then comparing those product groupings to the foundation. Looking at the market for Threat

Intelligence products, there are three groupings that can encapsulate a large portion of the market, namely: Feed driven, Research driven and Platform driven.

FEED DRIVEN PRODUCTS

Products that can be classified as 'Feed driven' were one of the first categories of Threat Intelligence to appear on the market. This is logical because it is relatively easy to convert logs that have been generated by traditional security controls (such as endpoint antivirus installations, network IPS devices and web proxies) into a feed. A Feed driven product is predominantly categorised by the way in which it delivers Threat Intelligence to the point of consumption, i.e. through a feed. This can take many forms and spans numerous iterations of complexity, but the overriding theme is that of delivering a feed.

For the most part, the feed provider gathers information through an array of collection points (often referred to as sensors) and transforms that information into a consumable feed which can then be delivered to the point where it needs to be deployed. These three steps of Gather, Transform and Deliver form the basis of the Feed driven category, but also allow for large deviations between different individual products. Vendor specific methodologies come into play here, specifically on how the data is gathered and the processes that are used to transform it. The Delivery step has been commoditised, as it is the defining characteristic of these forms of products, and has been optimised for rapid deployment and digestion. This leaves the first two steps as the only way of evaluating Feed driven products.

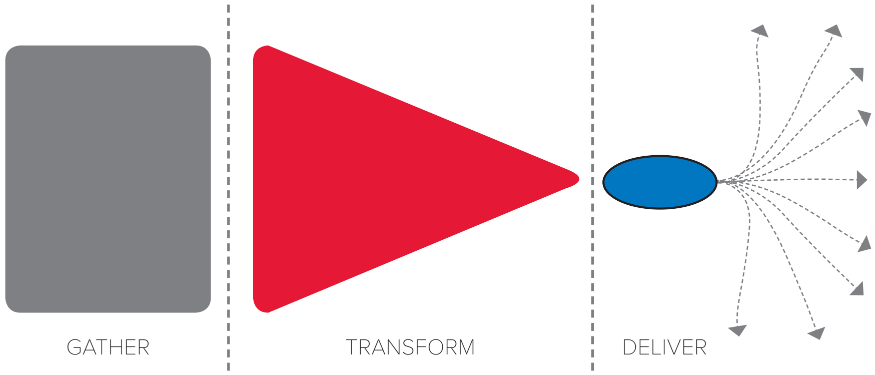

Diagram 5: Gather, Transform, Deliver

GATHER

Gathering of the information is fairly self-explanatory; it is the part of the chain where raw information is captured or collected, forming a base onto which analysis efforts can be applied. The information can come from anywhere but is generally collected by a large array of sensors that have some exposure to possible attacks. An example of this would be a security company that sells firewall devices, and repurposes the device to pass on to a central system the IP addresses of hosts that are trying to connect to something that is not explicitly accepted by the firewall's rule base. In this example the firewall is the sensor and the IP addresses are the information that the sensor is gathering. The sensor can take numerous forms, but is normally linked to the type of data that is being captured, so that if the analysts were looking for data that related to a new vulnerability in PHP global environment variables, it would be logical to deploy a honeypot that exposed this functionality to the internet and that was configured to collect information relating attempted exploits. But often the sensor infrastructure is chosen based on what is available rather than what would be efficient based on the data needs. This might be the case of security product companies that market their sensor infrastructure as containing millions of nodes and collecting billions of data points who, as in the first example, reuse their existing security control footprint (antivirus, IPS, web proxies) to collect and transfer data back to a centralised point.

TRANSFORM

At a very broad point of definition, the Transform step is the process of manipulating the collected data in a way that makes it more valuable to its target audience. Sometimes this can simply mean reordering the data (according to timestamp or by targeted port number for example) in such a way that it is more usable to the final consumer, or it could mean adding data to give additional context, grouping data in such a way that it creates a unique linkage between previously unconnected data points. There is a lot of room for an intelligence provider to add value in this process as the Transform step is as complex as the provider wants it to be. One extreme case is a provider that takes the data that it collects, pulls out the source IP addresses and publishes that list as a register of known bad source IP addresses. The other extreme option is that the Feed provider assigns an analyst to research every detail about every piece of information that is collected. There is no right or wrong approach here, simply the problem space that the feed is trying to address, but there is a lot of latitude to manoeuvre within this step. This is the step that differentiates one Feed driven product from another.

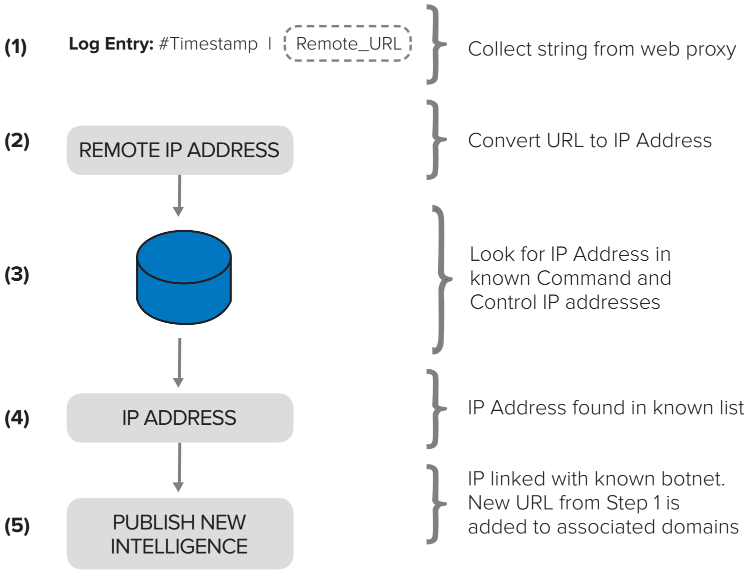
(1)
Log Entry: #Timestamp I Remote_URL
Collect string from web proxy
(2)
REMOTE IP ADDRESS
Convert URL to IP Address
(3)
Look for IP Address in known Command and Control IP addresses
(4)
IP ADDRESS
IP Address found in known list
(5)
PUBLISH NEW INTELLIGENCE
IP linked with known botnet. New URL from Step 1 is added to associated domains

Diagram 6: Grouping data

DELIVER

Due to this step being commoditised, there is very little room to deviate from the market-defined template. One such template is STIX (Structured Threat Information Expression) and TAXII (Trusted Automated Exchange of Indicator Information) from the not-for-profit Mitre Corporation. STIX outlines a standardised language to represent structured cyber threat information and intends to convey the full range of potential cyber threat information and strives to be fully expressive, flexible, extensible, automatable, and as human-readable as possible. TAXII is a set of protocols that defines a method for exchanging Threat Intelligence, and is the preferred method of exchanging STIX formatted intelligence. The use of standard templates makes it possible to interact with numerous different feed providers, without having to fundamentally change the way in which the customer consumes the intelligence.

The advantage of a Feed driven system is that it is able to collect a huge amount of data without designing and deploying dedicated sensor architecture. But this information dragnet approach also has its disadvantages; primarily that data is collected unambiguously and without any data collection model in mind. Because the data collection decision is based on what is already in place opposed to what is needed to collect the right data, the possibility exists that irrelevant data will be analysed during the Transform step, wasting time, or worse yet, tainting the results.

RESEARCH DRIVEN PRODUCTS

The second suite of intelligence products currently available in the market are those that rely on the use of analysts to distil information into a research report that can be delivered to a specific audience. These products are research driven because they are built on the notion that human analysts will apply rigorous research on the information that they retrieve, generating value for whoever is the target audience. These products can be compared to their Feed driven cousins in that they share the same basic three steps, but while those Feed driven products are defined by the step, Research driven products only display lose alignment. The reason for this is that Feed driven products need to be derived as autonomously as possible because of the amount of data that is being gathered, transformed and delivered. The research approach on the other hand is more specialised, looking at a specific set of data and delving deeper and deeper into it or adding supporting information. So while on the surface the Research and Feed driven models are similar, they are in reality very different.

GATHER

The strong reliance on analysts means that these sorts of products can afford to erect specific sensor architectures because of the need to analyse comparatively less data. This could end up being more expensive compared to repurposing but it also means that the data acquisition process is tailored to aid the Transformation step. There also seems to be two types of data sets being harvested for this type of product: the traditional threat information as stated above, and less raw data such as technical reports, white papers and presentation.

TRANSFORM

The right data analysts can then connect the pieces of data together and start creating meta-data. This, again, sounds very similar to the Feed driven products, but while the one extreme example from before (where an analyst was assigned to each incident) is considered to be commercially improbable, Research driven products have the resources and limited raw data to warrant such an approach. Their visibility might be less defined by the amount of data being analysed, but these products have the ability, and more importantly the time, to enumerate the data that they have collected. This results in information that is more concise and rigorously researched; adding layers of additional information to a potential threat, but that only covers a smaller spectrum of the threat landscape. For example, if the underlying data relates to third party research, then the analysts would condense a research paper into a short executive summary with some context as to how it relates to the real world.

DELIVER

Due to the output of this product being very high-level, the transformed information is usually delivered as a report or white paper. They normally have a theme that focuses on a particular threat actor,exploit or campaign, but describes it in great detail, referencing the infrastructure being used, what the risk of it is, and a series of IOC's (Indicators Of Compromise) that can be used to verify that this particular threat is not active on a network. If the original data was also research related, then a similar sort of report would be created that could be used as a summary of the research. An example of this would be an analyst who condenses the full Black Hat Security Briefings into a report that contains the most impactful aspects, how they affect the current landscape and what can be done to mitigate against them.

The advantages and disadvantages of a Research driven product is nearly the opposite of the Feed driven products. They are expensive to set up because they do not lean on an established dragnet gathering infrastructure, but they have the potential to miss large portions of data. The quality of the intelligence is directly related to how thorough that particular analyst was in terms of gathering and transforming the data. With such a strong reliance on human intervention it is unlikely that these sorts of products would ever reach a state of high maturity as defined by the Capability Maturity Model.

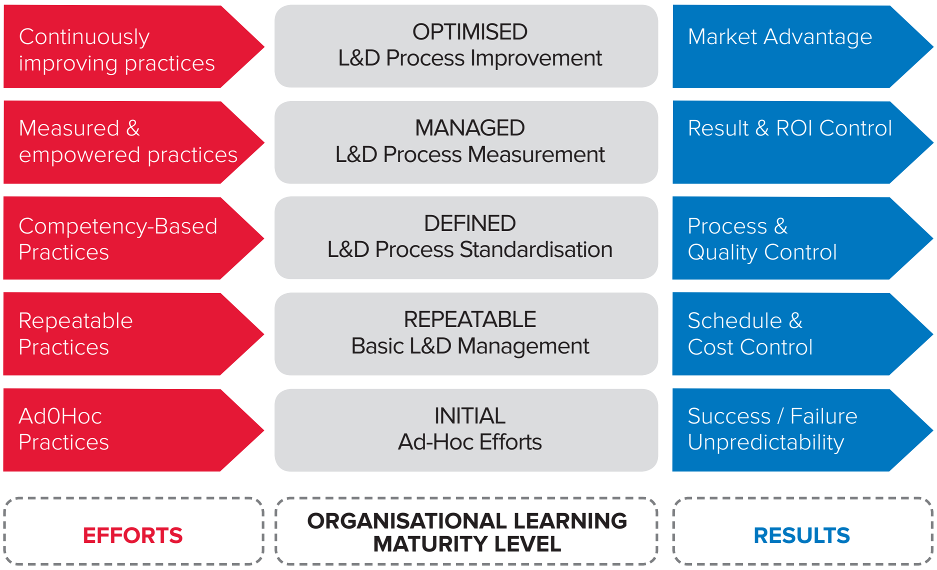

EFFORTS	ORGANISATIONAL LEARNING MATURITY LEVEL	RESULTS
Continuously improving practices	OPTIMISED L&D Process Improvement	Market Advantage
Measured & empowered practices	MANAGED L&D Process Measurement	Result & ROI Control
Competency-Based Practices	DEFINED L&D Process Standardisation	Process & Quality Control
Repeatable Practices	REPEATABLE Basic L&D Management	Schedule & Cost Control
Ad0Hoc Practices	INITIAL Ad-Hoc Efforts	Success / Failure Unpredictability

Diagram 7: Capability Maturity Model.

PLATFORM DRIVEN PRODUCTS

The last type of product is radically different to the other two in that they do not really produce Threat Intelligence, but rather provide a platform with which Threat Intelligence can be easily managed and shared. Being a platform there aren't any definable steps in delivering the information since the platform is always available and any information on it needs to be added to it by the end user. These sorts of products are in reality more than just different when compared to the other types, but are really something altogether different. These Threat Intelligence Platforms (TIP) exist to make the exchange of threat information easier and store data in such a way that it is intuitive for a Threat Intelligence analyst.

The comparison below aims to give an idea of where some of the current industry vendors positions their intelligence products. Some of them have offerings that straddle more than one particular silo, as indicated by the X, but the majority of those multi-silo products still have one classification that they are particularly strong in. This is highlighted with a bold **X**.

	FEED DRIVEN	RESEARCH DRIVEN	PLATFORM DRIVEN
Threat Connect	X		**X**
iSight Partners	**X**	X	X
RecordedFuture	**X**		X
McAfee Research	**X**	X	X
Symantec Deepsight	**X**	X	X
IBM X-Force	X	**X**	
Dell Secureworks		**X**	

Comparing the Products to the Research

By looking at what some of the large research firms within the industry have said about Threat Intelligence we have managed to draw some foundational characteristics of what Threat Intelligence is, but also what features a Threat Intelligence product should strive for. Now that we also have grouped the most prevalent Threat Intelligence products into qualifiable classes we can compare each of them to the view produced by the research firms. The Platform driven products will not be evaluated because they do not provide Threat Intelligence, but rather offer a way to access and use it.

Gartner's simple-to-apply definition for Threat Intelligence has the following four key points, with which we can evaluate the different Threat Intelligence products currently available on the market:

1. Evidence-based knowledge...
2. including
   i. context,
   ii. mechanisms,
   iii. indicators,
   iv. implications
   v. and actionable advice...
3. about an existing or emerging menace or hazard to assets...
4. that can be used to inform decisions regarding the subject's response to that menace or hazard.

Feed driven products are certainly evidence based-knowledge because they start off as raw data (evidence) being gathered from some sensor. The second point is not that simple however, because it defines what that evidence should contain, and is partly open to interpretation. Context (i) is difficult for Feed driven products because they are constructed in such a way that they can be consumed by lots of different end points, allowing the service to scale. But by designing a product to scale, very little detailed assumptions can be made with regards to where the Threat Intelligence will be used, meaning that no localised knowledge can be built into these products. Mechanisms (ii), indicators (iii) and implications (iv) can be encapsulated in a threat feed as they are attributes of the attacker, and can be collected without taking into consideration the Threat Intelligence consumer. But actionable advice (v) is again similar to context (i) because without the ability to add localised knowledge to the intelligence it is not possible to advise on how best to mitigate a potential threat. Point 3 is an excellent example of the missing link when it comes to Feed driven products, as it references the two parts that Gartner highlights as being important: the existing or emerging hazard and the asset. The Feed driven products are able to collect information about hazards, but they have no way of knowing anything about the assets, and therefore are unable to offer any intelligence that uses knowledge of the local assets. The missing information on the local assets means that in Point 4 the feed could be used to help the decision making process with regards to response, but that decision-makers would need additional information pertaining to the local assets as well as the feed data, defeating the purpose of an information gathering service such as a threat feed.

Applying the Gartner definition to Research driven products highlights exactly the same issue as found with Feed driven products: because research driven products are also designed to be delivered to more than one customer. They therefore cannot take accurate local-asset information into account, settling instead for best-guess or generalist assumptions. These products are able to generate a great deal of relevant information on threats (or hazards as Gartner calls them) but cannot link that back to how actual end point architecture would be affected.

Both the Feed and Research driven product sets are able to satisfy certain sections of the Gartner definition by collecting information on hazards, but neither can link that intelligence back to the end user in a way that doesn't require some additional transform work by the consumer. According to Gartner, they are parts of a Threat Intelligence offering, but cannot be referred to as Threat Intelligence by themselves.

The Forrester research is less straightforward than the Gartner definition, presenting the concept of "Threat Intelligence as a framework" that is difficult to evaluate products against. Forrester uses a framework that leans heavily on gathering the right data and applying the right set of analyst skills to that data in order to generate intelligence. Given these requirements for Threat Intelligence it would be up to the end-user to choose a provider that they feel has collected the right data, or who has access to the type of data that they feel would be relevant to their threat landscape. While this doesn't disqualify either of the products, it does shift the responsibility of accurate Threat Intelligence over to the consumer, creating a dichotomy where a party has all the responsibility but limited information to make an informed decision. This is less than optimal.

Forrester also notes that once the data is collected it is up to the analysts to churn through it and pull out the relevant intelligence. Considering the difficulties of staffing a full time Threat Intelligence analyst team it would seem desirable to purchase one of the products that already offers this, but the same dichotomy as with information gathering applies: the consumer of the Threat Intelligence product needs to evaluate the competency of that product's analysts and determine whether their skills and methodology are relevant to the type of intelligence that the consumer would be interested in. Strictly speaking both Feed and Research driven products can collect the right information and house the relevant skills needed to generate Threat Intelligence, but the potential pre-work demanded from the consumer is substantial. While 'actionable' doesn't appear in the Forrester framework as a discrete step, it does address it in the text stating that through the correct combination of data and analysts, actionable Threat Intelligence can be generated. But for intelligence to be actionable it needs to have knowledge of the assets that are deployed by the consumer. Unfortunately, due to the way these products have been designed it is not possible for them to package that degree of knowledge without compromising their ability to scale. Once again, it is the product's lack of localised knowledge that prevents it from adhering to Forrester's framework for Threat Intelligence.

The Gartner and Forrester research models present different aspects of what makes a certain source of Threat Intelligence valuable. When the two models are used to evaluate current commercial product sets however, none that we've seen are able to satisfy either research house's "actionable" or "relatable" criteria. This is further evident in the SANS Institute's white paper definition of a threat (something that has the intent to be a threat, the opportunity and the capability to act against an asset) which ties back to an asset, and therefore localised knowledge. Without knowledge of what is or could be attacked, how can intelligence relating to it be gathered?

Conclusion

The name Threat Intelligence has undergone many transformations from military doctrine and experience before being repurposed by the information security industry. But while the exact name has changed, picking up and losing prefixes and suffixes along the way, it can still be defined as a way of gathering information relating to adversaries. This could be the adversary itself, what they have targeted in the past or what they are likely to target on a specific network, but it is centred on the notion of gathering information to support decision-making. By investigating what the large research houses have distributed with regards to Threat Intelligence we were able to derive an approach that could be used to evaluate a potential Threat Intelligence product. After factoring the current product landscape into three groups (Feed, Research and Platform driven) two of them were evaluated against our approach. The last, Platform driven products, was excluded because it does not provide intelligence, only a way to house and share it. In evaluating the Feed and Research driven products it becomes clear that while they had the potential to add value, such as offering an outsourced information gathering and analyst function, there is a serious lack of contextualising around localised knowledge, limiting the ability of any intelligence derived from these products to be actionable.

Threat Intelligence products have evolved rapidly, creating offerings that have huge visibility and that undergo heavy data transformation before reaching the end user, but there is still a significant piece missing: localised knowledge of the target environment. It could be possible to overcome this limitation on the end user side through rigorous evaluation of Threat Intelligence products before purchase, and then using internal analysts to mutate the incoming intelligence to better suit the consumer architecture. The cost of this would be significant however. An alternative would be for a consumer to have direct access to a Threat Intelligence provider's storage and transform backend so that they could pull out intelligence based on their localised knowledge. This is unfortunately not a problem that can be solved with the current set of products, because they rely on the production of a generic service that can be consumed by numerous end users, which fundamentally clashes with the product's ability to gain localised knowledge.

References

[1] Bayly, C. A. (1996). Empire and information: intelligence gathering and social communication in India, 1780-1870 (Vol. 1). C. A. Bayly (Ed.). Cambridge University Press.

[2] Devens, Richard Miller. (1868). Cyclopaedia of Commercial and Business Anecdotes: Comprising Interesting Reminiscences and Facts... of Merchants, Traders, Bankers... Etc. in All Ages and Countries... D. Appleton.

[3] Tinnel, Laura S., O. Sami Saydjari, and Dave Farrell. (2002). Cyberwar strategy and tactics. IEEE workshop on information assurance, United States Military Academy, June 2002.

[4] Mcwhorter, Dan. (2013) APT1: Exposing One of China's Cyber Espionage Units. Mandiant.com (2013).

[5] McMillan, Rob. (2013) Definition: Threat Intelligence. Gartner Publications, May 2013. Online: https://www.gartner.com/doc/2487216/definition-threat-intelligence

[6] Cloppert, Mike. (2009) Security Intelligence: Introduction (pt 2). SANS Digital Forensics and Incident Response Blog, Jul 2009. Online: http://digital-forensics.sans.org/blog/2009/07/23/ security-intelligence-introduction-pt-2/

[7] Schoeman, Adam. (2014) Amber: A Zero-Interaction Honeypot with Distributed Intelligence. ISSA-13.