MWR INFOSECURITY

CERT-UK

CPNI
Centre for the Protection of National Infrastructure

Threat Intelligence: Collecting, Analysing, Evaluating

Authors:

David Chismon

Martyn Ruks

MWR would like to acknowledge the help and support of CPNI and CERT-UK in researching this topic and producing the accompanying products

Contents

Introduction

Threat intelligence is rapidly becoming an ever-higher business priority. There is a general awareness of the need to 'do' threat intelligence, and vendors are falling over themselves to offer a confusingly diverse array of threat intelligence products.

Figure 1: Google results for "threat intelligence" from different years

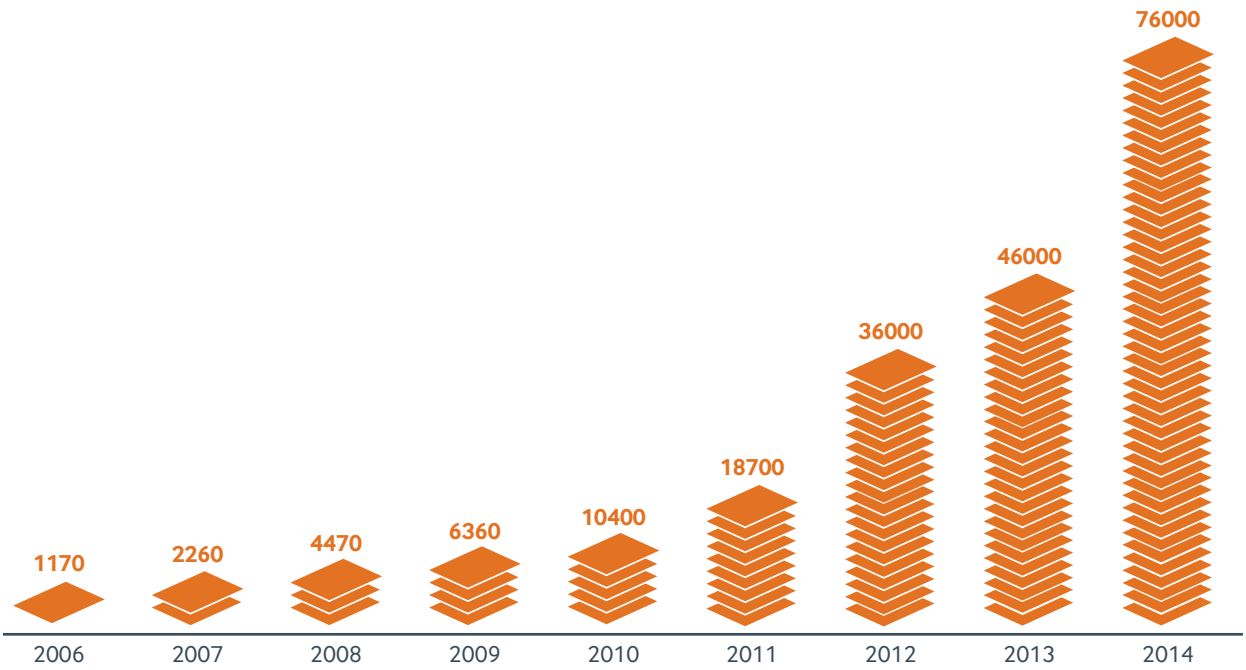

The promise of threat intelligence is alluring. It should help organisations to understand and manage business risk – to turn unknown threats into known and mitigated threats, and to improve the effectiveness of defence. After all, targeted attacks need targeted defence. If analysis is performed correctly, the products of threat intelligence can be genuinely useful to a business, providing real benefits at all levels, from on-the-ground defenders to the board.

However, threat intelligence is currently very loosely defined, with little agreed consensus on what it is and how to use it. There is a risk that in the hurry to keep up with the threat intelligence trend, organisations will end up paying large amounts of money for products that are interesting but of little value in terms of improving the security of their business. 'Doing' threat intelligence is important – but doing it right is critical.

To address this, MWR InfoSecurity reviewed the area and designed a framework for threat intelligence that can be scaled to different sectors, sizes of organisation, and organisational goals. The paper is the product of literature reviews, internal experience, and a large number of interviews with people involved in threat intelligence and related fields across a range of organisations.

What is Threat Intelligence?

What is Intelligence?

Intelligence is regularly defined as information that can be acted upon to change outcomes. It's worth considering traditional intelligence before exploring threat intelligence, as in many ways the latter is simply traditional intelligence applied to cyber threats.

Since Donald Rumsfeld's DoD briefing in 2002, the concept of 'knowns' and 'unknowns' tends to appear regularly in discussions on the subject of intelligence. An 'unknown unknown' is a threat or risk that we don't know we don't know about – in other words, we have no idea that the threat even exists. For example, we are completely unaware that someone is waiting outside the office to attack the CEO. A 'known unknown' is something we know that we don't know: perhaps we've been told that the CEO is going to be attacked outside the office, but we have no details as to who, why, when or how.

One description of threat intelligence is the process of moving topics from 'unknown unknowns' to 'known unknowns' by discovering the existence of threats, and then shifting 'known unknowns' to 'known knowns', where the threat is well understood and mitigated. For example, once we've been told the CEO is going to be attacked outside our office, we find out who the attackers are and what weapons they're carrying; and then inform the CEO so that travel plans can be changed – or the attackers arrested before the incident takes place.

Figure 2:

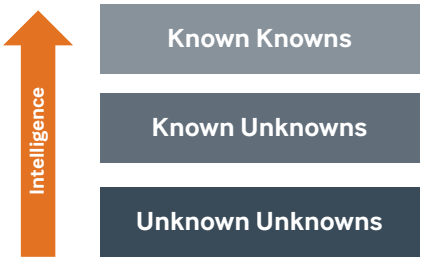

Understandably, the aim is to have the majority of risks in the 'known knowns' category, while developing some current 'known unknowns' and allowing as few threats as possible to remain as 'unknown unknowns'. However, this is a considerable challenge in traditional intelligence and equally so when applied to cyber threats. The Butler Review of Intelligence on Weapons of Mass Destruction noted a limitation of intelligence, in that it is often incomplete and seldom obtains the whole story – as intelligence inherently seeks to gain knowledge of things that others are working to obscure[1]. Furthermore, the report commented that, "The necessary protective security procedures with which intelligence is handled can reinforce a mystique of omniscience."

It could be argued that the NDAs (non-disclosure agreements), marketing and sheer price of cyber threat intelligence can contribute to the same perception of omniscience by its consumers.

Different Definitions

In the world of information and cyber security, threat intelligence is a young field and there are large numbers of threat intelligence vendors and advisory papers that describe very different products and activities under the banner of 'threat intelligence'. As with traditional intelligence, a core definition is that threat intelligence is information that can aid decisions, with the aim of preventing an attack or decreasing the time taken to discover an attack. Intelligence can also be information that, instead of aiding specific decisions, helps to illuminate the risk landscape.

However, the nature of that information can vary greatly, often with almost no commonality or comparability among the various threat intelligence offerings. Prices for similarly positioned (but very different) offerings can also vary wildly, with 100-fold variations in product pricing from different providers – even when the products claim to meet the same need.

The products and services sold as threat intelligence can vary enormously in their scope, usability, aims and content. For example, at a high level, some products come in the form of prose that explains developments in a particular area, while at a lower level, others might be a stream of XML-formatted indicators of compromise, such as IP addresses or binary hashes.

Even within similarly placed sources, such as feeds of indicators of compromise, there is very little overlap between competing products. Recent research suggests that in three popular feeds of flagged IP addresses, containing more than 20,000 IP addresses in total, there was as little as a 1% overlap[2]. This suggests that either attackers are using huge numbers of IP addresses and even well-known feeds see only a small part of the picture, or only a minority of IP addresses contained within the feeds are of intelligence value. It's likely that the truth is a mixture of both explanations.

As market demand for threat intelligence grows, with a large number of organisations either interested in products or actively building programmes, some vendors are offering existing products – or subtly reworked versions of existing products – as 'threat intelligence'. At the more cynical end of the spectrum, it's been suggested that threat intelligence is at a threshold where it could become either useful, or simply antivirus signatures by another name... and at a higher price[3].

Subtypes of Threat Intelligence

Any information about threats that could inform decisions is arguably threat intelligence. This broad definition obviously covers a huge variety of sources and information, from watching a TV news report about how attackers are exploiting a flaw, to a quiet drink with a friend at a competing organisation who mentions they are seeing more phishing with PDF documents. Organisations that make good use of these relatively abstract sources will often be more resilient and aware of threats than organisations that make poor use of expensive products.

With so many different sources falling into the category of threat intelligence, it can be useful to have subdivisions to focus effort and better manage the information. For example, a prose report of national activity is not comparable to an IP address and cannot be actioned in the same way.

Identifying subtypes of threat intelligence can be based on who consumes the intelligence and what it aims to achieve. We propose a model that breaks down threat intelligence into four distinct categories based on consumption. Each area is discussed in depth in later sections, but the following is a summary of the four categories:

Figure 3: Subtypes of threat intelligence

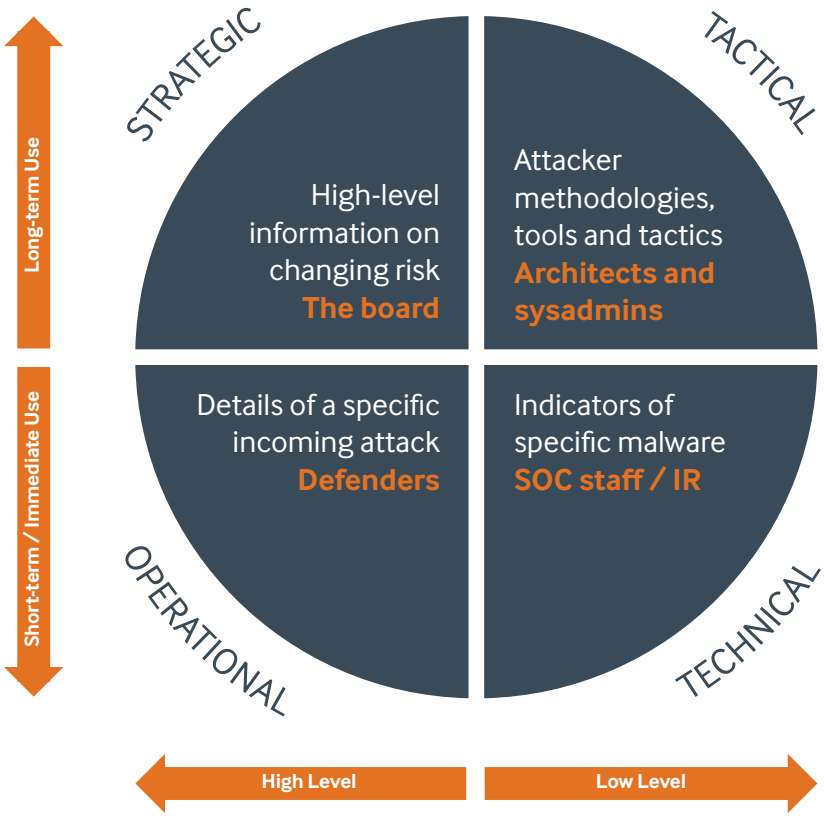

Strategic Threat Intelligence is high-level information, consumed at board level or by other senior decision-makers. It is unlikely to be technical and can cover such things as the financial impact of cyber activity, attack trends, and areas that might impact on high-level business decisions. An example would be a report indicating that a particular government is believed to hack into foreign companies who have direct competitors within their own nation, hence a board might consider this fact when weighing up the benefits and risks of entering that competitive marketplace, and to help them allocate effort and budget to mitigate the expected attacks. Strategic threat intelligence is almost exclusively in the form of prose, such as reports, briefings or conversations.

Operational Threat Intelligence is information about specific impending attacks against the organisation and is initially consumed by higher-level security staff, such as security managers or heads of incident response. Any organisation would dearly love to have true operational threat intelligence, i.e. to know which groups are going to attack them, when and how – but such intelligence is very rare. In the majority of cases, only a government will have the sort of access to attack groups and their infrastructure necessary to collect this type of intelligence. For nation-state threats, it simply isn't possible for a private entity to legally gain access to the relevant communication channels and hence good operational threat intelligence won't be an option for many. There are cases, however, where operational intelligence might be available, such as when an organisation is targeted by more public actors, including hacktivists. It is advisable for organisations to focus on these cases, where details of attacks can be found from open source intelligence or providers with access to closed chat forums. Another form of operational threat intelligence that might be

available is that derived from activity-based attacks: where specific activities or events in the real world result in attacks in the cyber domain. In such instances, future attacks can sometimes be predicted following certain events. This linking of attacks to real-world events is common practice in physical security but less commonly seen in cyber security.

Tactical Threat Intelligence is often referred to as Tactics, Techniques, and Procedures (TTPs) and is information about how threat actors are conducting attacks. Tactical threat intelligence is consumed by defenders and incident responders to ensure that their defences, alerting and investigation are prepared for current tactics. For example, the fact that attackers are using tools (often Mimikatz derivatives) to obtain cleartext credentials and then replaying those credentials through PsExec is tactical intelligence that could prompt defenders to change policy and prevent interactive logins by admins, and to ensure logging will capture the use of PsExec[4]. Tactical threat intelligence is often gained by reading white papers or the technical press, communicating with peers in other organisations to learn what they're seeing attackers do, or purchasing from a provider of such intelligence.

Technical Threat Intelligence is information (or, more often, data) that is normally consumed through technical means. An example would be a feed of IP addresses suspected of being malicious or implicated as command and control servers. Technical threat intelligence often has a short lifetime as attackers can easily change IP addresses or modify MD5 sums, hence the need to consume such intelligence automatically. Technical threat intelligence typically feeds the investigative or monitoring functions of a business, by – for example – blocking attempted connections to suspect servers.

One Source, Multiple Intelligence Types

While a single source tends to provide intelligence of only one specific type – for example, a data feed that is useful only as technical threat intelligence – many useful sources can provide multiple types of intelligence that can be analysed and turned into different products for effective consumption.

An increasingly common practice is for private organisations to publish white papers on attack groups or campaigns. A single document can contain almost all types of intelligence. For example, the fact that hackers believed to be working for a particular nation state have been attacking a specific industry sector is strategic intelligence. The details of their modus operandi, tooling and capabilities is tactical intelligence and can inform defences, while the list of MD5/SHA-1 hashes of binaries that often appears in appendices is technical intelligence that can be used for investigation.

Few, if any, of these reports contain operational threat intelligence as, by the definition given in this paper, the report would need to contain details of a specific impending attack.

How Do You Build and Evaluate a Threat Intelligence Programme?

The Threat Intelligence Cycle

An effective threat intelligence (TI) programme will have a number of areas of focus. The breakdown of threat intelligence into specific functions is more scalable, as staff are likely to be better skilled at specific aspects of intelligence. Individual parts of the cycle can be focused on and developed, while it will be easier to track insufficient results from the programme to specific weaknesses.

An oft-quoted model is the 'intelligence cycle'. The steps in the cycle are as follows[5].

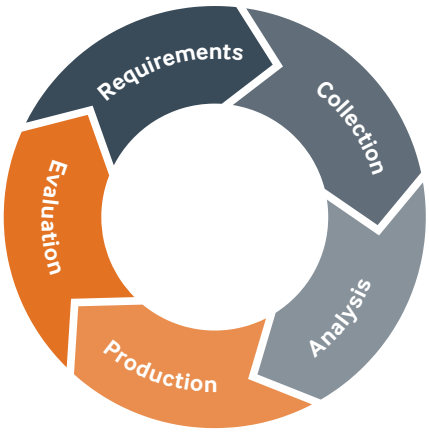

Requirements: A step that is often overlooked is also the key to a successful programme. Decision-makers need to identify what they specifically want to know and what the TI programme should be telling them. For example, a requirement might be: "Inform us of all publically known, widely exploited vulnerabilities within one day of them becoming known." This can also be referred to as 'tasking'. For example, if a company were considering a partnership with an organisation from country X, the TI team could be tasked with understanding whether country X is known to abuse such relationships, and what technical tools and tactics have been used to do so. Requirements can also be more demanding of TI teams, such as, "Obtain details and samples of the majority of criminal outfits' remote access toolkits for our forensic teams." TI teams need to work with decision-makers to agree on requirements that are not only feasible but, crucially, that will supply products on which the organisation will be able to act.

Collection: The step that can dominate much of a TI budget is collecting the information or data that is expected, once analysed, to fulfil the requirements. The information can come from a large variety of sources, such as news feeds, paid-for services or feeds, forums, white papers, or even human sources. Almost all paid-for threat intelligence from vendors comes under this category and will require some form of analysis. Understanding which sources are likely to produce the desired information, to be reliable and to provide information that can be consumed in a timely manner, is not a trivial process, and it is far better to 'pour a measure of spirits than to try sipping from a fire hose'. Collection for specific subtypes of intelligence will be discussed in later sections.

The value of collecting from human sources should not be underestimated. In traditional intelligence, covert sources are normally tapped to provide intelligence, but in threat intelligence the focus is on sharing information through relationships with peers in other companies in the same (or potentially other) market sectors. The ability to have a quiet catch-up with a peer and ask whether they saw increased activity once they became involved with country X can provide highly useful information. However, it's important to do so in a way that doesn't tip off a competitor to unreleased business plans. Trusted forums and relationships can help an organisation share information safely – and also help others to trust the information received.

Analysis: Turning data into information that can be actioned often requires analysis[6]. In some cases, analysis will be relatively simple, e.g. parsing a feed into a firewall deny-and-alert ruleset. In other cases it will require extracting the relevant information from a larger work, such as a report, and understanding which elements apply to the organisation's assets. An important role for the analyst is to look for opportunities to create new types of intelligence through synthesis from current intelligence. For example, an analyst might spend time reading through white papers to extract indicators of compromise, and also identifying operational intelligence that can be given to network defenders. Or, after reading such papers and other sources, the analyst might identify trends that can be drawn together into a strategic intelligence product for higher management. An interplay between collection and analysis often occurs, where analysts realise that the collection is not producing the required raw material; or perhaps that different information needs to be collected for appropriate analysis. Collection can then be altered and analysis continued.

Production / Dissemination: In this stage, an intelligence 'product' is created and disseminated to the customers (senior executive officers, network architects, defenders, etc.). The product will vary, depending on the subtype of intelligence and the customer. For example, it might require a three-line report to the board, a white paper to defenders, or simply an approved rule added to defence hardware.

Evaluation: Another frequently neglected phase of threat intelligence (if modelled on traditional intelligence) is the evaluation of the intelligence product to ensure it meets the original requirements. If the requirements have been met, then the product can further feed the requirements to help develop new, deeper requirements that build upon the intelligence product – and the intelligence cycle can repeat. If the produced threat intelligence does not meet requirements, then it suggests a failure at some point, and the cycle model can be used to establish where the failure occurred. Were the requirements unrealistic? Did the collection use the wrong sources? Was the data contained within the sources but not drawn out during analysis, or did the final product not contain the intelligence gained?

A Modified TI Functional Flow

In 1996, the United States Senate Select Committee on Intelligence published a study on how the intelligence community might look in the 21st Century if it were redesigned from scratch. This study proposed a functional flow for intelligence that can be used as the basis for a mature, scalable TI programme, as shown in figure 4.

Although similar to the threat intelligence cycle, there are some subtle differences. The functional flow differentiates between intelligence management and execution, and this distinction can be useful when building and managing an organisation's teams. Requirements remain the cornerstone and a good entry point into the cycle. Requirements drive collection and analysis management, with resources balanced between them as necessary – and open to changes as the cycle progresses. Collection feeds analysis, but analysis also informs and modifies collection to ensure the necessary data is gathered. The products are then evaluated against the requirements, which helps to set the requirements for the next cycle.

An important departure from the traditional threat intelligence cycle is that resources can be used to develop systems and capabilities of potential use to both collection and analysis, based on advice from the collection and analysis functions. Applied to a TI programme, this might mean that new feeds are required, or new systems to parse and process the feeds, or perhaps there is a need to develop analytical engines. As an example, analysts might raise the fact that some elements of the information they are collecting are not currently being analysed and acted upon. It might be wise to modify the requirements to include that information, or the issue could simply be that analysts lack sufficient staff or systems to analyse the collected data effectively.

Figure 4: A modified threat intelligence functional flow

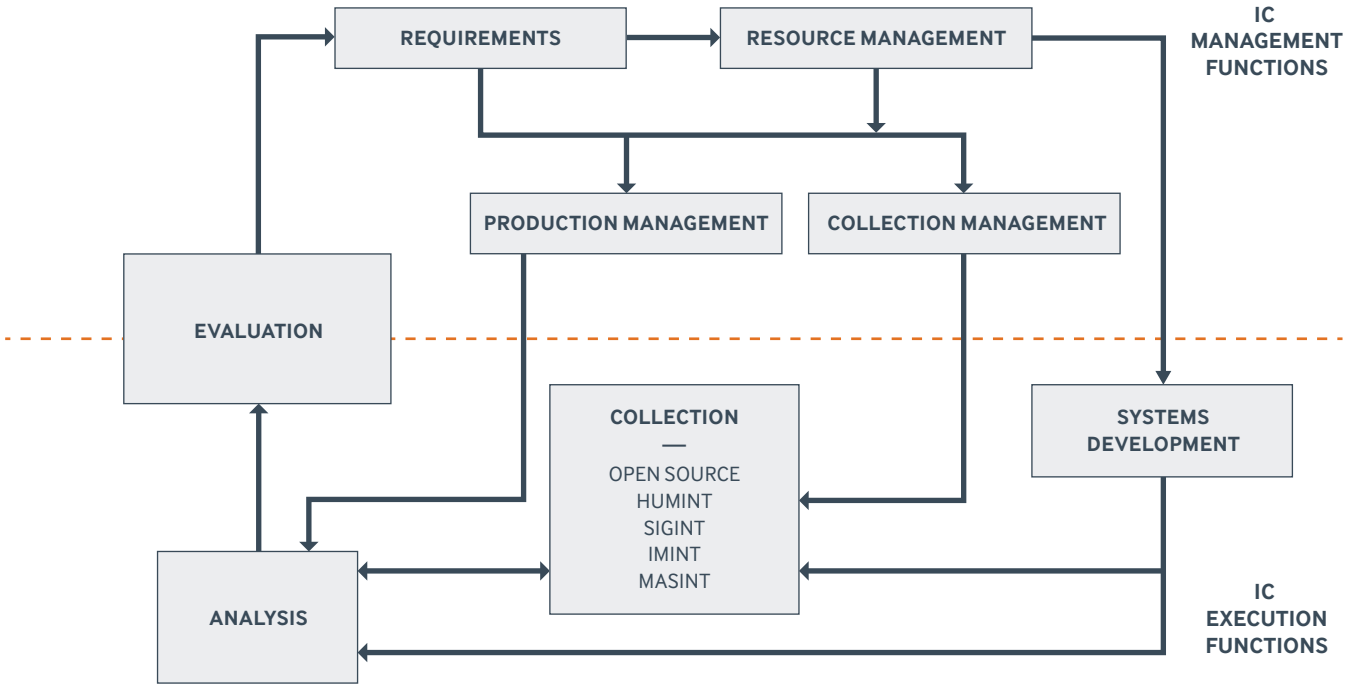

Adapted from: http://www.gpo.gov/fdsys/pkg/GPO-IC21/html/figure1a.gif

How to Build a Threat Intelligence Programme

As previously stated, it's crucial that threat intelligence is 'requirements focused', with the requirements phase of the threat intelligence flow defining the questions that need to be answered. Since the definition of both traditional and threat intelligence is information that can be acted upon, it's only logical that organisations should also ensure they will be able to act on the answers they seek. Resources and tasking will be required by both the threat intelligence function and whoever intends to act on the resulting intelligence. There is little point, for example, in obtaining a list of MD5/SHA-1 hashes if the organisation has no ability to search for binaries with those hashes on its network or hosts.

Once requirements have been decided, the next step is to identify the sources from which information and data will be collected, along with the analysis necessary to produce actionable threat intelligence.

How Not to Build a Threat Intelligence Programme

The majority of TI programmes that are failing to provide meaningful intelligence and business value have factors in common when it comes to how they were built. Typically, senior management decided that a threat intelligence team was necessary, a decision based on interactions with peers, writings in the field or even vendor pitches. Rather than the requirements driving the establishment of teams, the perceived need simply to have a team drove the whole process. It's not unknown for senior staff to muse, "We don't know what threat intelligence is, but we know we need it."

In the absence of clearly defined requirements, these teams, once created, search for something to offer as threat intelligence, and often end up simply consuming whatever vendors are selling. This can be defined as 'collection-focused threat intelligence' that seeks to consume feeds – or whatever is in vogue – in the hope of extracting meaning, and it rarely offers significant benefits to the organisation.

Official or Unofficial?

Many organisations have, or are currently building, dedicated threat intelligence teams with full-time staff members and a budget for hardware and software to manage the intelligence feeds. However, other organisations are benefiting from threat intelligence without any dedicated staff or specific budget and, in some cases, might not even be aware that they are effectively 'doing' threat intelligence.

Returning to this work's definition of threat intelligence as information on threats that is actionable, organisations can obtain and act on information without any full-time staff dedicated solely to that purpose. A number of organisations have engaged staff members who, by merit of their reading around the subject and staying abreast of developments in security, as well as participating in forums such as CiSP[7], have a well-developed understanding of the threats to their business and developments in attacker methodology. In small organisations, these individuals are often the ones who directly act on the information – by changing group policy to prevent a certain type of attack, or adding a block rule to a firewall. Although there is no specific process or team, they are setting requirements (they seek awareness of threats to their business), collecting information (reading blogs, twitter, forums, etc.), analysing information (realising they are running services that are vulnerable) and acting on that information (patching services). A subconscious evaluation then occurs, where they realise that certain blogs are better than others or certain companies produce reports that are more directly useful.

In cases where unofficial threat intelligence is already taking place in the business, staff members should be encouraged.

The further development of threat intelligence should focus on supporting such efforts, with management identifying the areas in which to put resources (including both money and allocated time) to develop the function further.

Where organisations desire an official threat intelligence function, it is important to staff it with team members who have the right mentality to seek out the information, as well as a level of technical and business understanding to be able to draw the right conclusions – and then apply their findings to the business's assets. Some aspects of the threat intelligence team's work, such as certain types of collection, can also be a good place to start more junior members on their security careers. It will expose young employees to interesting aspects of the team's activities, provide experience of distilling technical information into products for more senior audiences, and utilise their probable familiarity with such sources as twitter, blogs and white papers.

Ex-Bonds or Joe Public?

A common debate among organisations interviewed for this work was the merit of staffing threat intelligence teams with people who have traditional intelligence experience. Opinions were divided, with some organisations believing it to be an important factor in the success of the team, and others not thinking it particularly necessary. Some aspects of building a threat intelligence team or function in the business can benefit from an understanding of traditional intelligence and reporting; however, as with most recruitment, it probably comes down to the individual's specific skills and experience, rather than just where they have worked.

Need to Share

In the world of traditional intelligence, 'Need to Know' is a well-established security principle. By restricting information to those who genuinely need it, you reduce the data stolen when an individual's access (or a specific computer) is compromised. In today's world of effective and motivated attackers, often with nation-state funding and resourcing, such security principles are highly important when it comes to limiting information loss.

However, in the world of threat intelligence there is an equally important 'Need to Share' principle. All subtypes of threat intelligence, if shared, will aid other organisations in defending against attacks. By establishing sharing communities and relationships, everyone can benefit from each others' intelligence. A company can be damaged when a rival business's computers are hacked, since the information stolen can often be used against other organisations in the same sector; and if a nation state is keen to support its own companies by such means, the impact of information theft will end up hurting all companies in the competing UK market.

Furthermore, many attacks do not target a single organisation in isolation, but rather target a number of organisations – often in the same sector – and hence discussion and understanding of attacks can be valuable to all related businesses. As entire communities are attacked, those communities need to defend: the aim is to raise the bar and constantly increase the cost to attackers.

How to Share

The various types of threat intelligence will need to be shared in different ways[8] (more detailed advice is given later in this document). However, effective sharing requires trust, as the shared information might be sensitive – for example, revealing that you have been attacked. Trust is also important on another level, as it is generally unwise to allow threat actors to learn what you know about them, lest they change their methods. The attackers might have realised that their tools aren't 'phoning home', but this doesn't mean they know how you're managing to stop them, and hence what they need to change.

For these reasons, closed and trusted groups can enable deeper sharing than would otherwise be possible. The groups can take many forms: for example, there are information exchanges for different industries run by parts of the UK Government and there is the online CiSP portal, which ensures that members are legitimate individuals at approved organisations. Various industry sectors have groups that share information, sometimes via a forum, sometimes simply by means of an email list. There are also less official groups, such as those set up on general online forums. The more a group can trust its members and the security of information within the group, the more effective the sharing tends to be. Organisations are recommended to seek out such groups and, if none exist, to consider creating them. Supporting these groups by encouraging staff to contribute is also important.

Some of the most useful sharing, however, can come from trusted personal relationships with similarly placed people at other organisations. This is obviously not scalable and it can take time to build the necessary trust, while the sharing needs to be mutually beneficial in order to succeed. Nevertheless, the value of such relationships should not be underestimated, and should even be directly supported. Attendance at networking groups and information exchanges can prove useful, but there are also small ways to help develop these productive relationships – such as allowing threat intelligence team members to charge meals out with counterparts as a legitimate business expense.

What Prevents Sharing?

There are two common reasons cited by organisations for not sharing threat information with others. One is the belief that they have nothing worth sharing; and the second is that their competitors might use the information against them. In some sectors, even a rumour of compromise can influence purchasing decisions or market valuations.

The concern about a lack of information to share might be valid if the organisation is fortunate enough not to be under attack. However, this is an increasingly rare situation and it's likely that, if looked for, there would at least be signs of attempted attack to share. Signs of attempted compromise can be particularly useful threat intelligence, as even if an attacker didn't – for example – successfully compromise one organisation through SQL injection, it might well be luckier with another in the same market sector. It can therefore be incredibly useful to share instances where defences have proved successful, as others can consider implementing those same defences.

Concerns about the risk of revealing a weakness to one's competitors is natural and it would be wise for companies to ensure they do not reveal information that could have a negative business impact unless the benefits are clear. Ideally, trust needs to be built up, either within groups or with specific individuals, to engender confidence. An organisation needs to feel confident that a competitor's defenders will act on the information given without revealing anything sensitive that might be misused by their colleagues – for example, the sales force. The continued benefits of playing by the rules should outweigh the single-instance benefit of betraying trust.

Organisations also need to be able to trust their own employees involved in the information-sharing arrangements. The ideal employees for such activities are individuals with a high degree of personal integrity and sufficient social skills to avoid the risk of oversharing; plus, of course, it's worthwhile selecting employees who are unlikely to leave the organisation in the near future.

In some industries, even the faintest whiff of suspicion that a company has been compromised is likely to influence buyers to go elsewhere. In these cases, organisations might do well to use trusted third parties to anonymise and distribute the information, so that communal benefit can be gained with minimal reputational risk. CERT-UK is able to act as a third party, as are private sector organisations – for example, companies involved in the Cyber Incident Response (CIR) service[9].

Vulnerability Assessment and Threat Intelligence

Some organisations include vulnerability assessment within the scope of the threat intelligence function. The threat intelligence function might even have grown out of the team that manages vulnerabilities. This can make sense as, in both cases, a team is tasked with finding information on the wider internet, analysing the information to decide whether it applies to the business, and then acting upon it. Organisations can even be tempted to regard a vulnerability notification as 'threat intelligence'.

The distinction between vulnerability information and threat intelligence is subtle. That a vulnerability exists in a product used by the organisation is important information, and requires action, but it's not information about a particular threat. However, information that a particular attack group is exploiting a known vulnerability, such as was seen shortly after the Heartbleed security bug was released[10], is tactical threat intelligence.

Whether or not the same team handles vulnerability assessment and threat intelligence is up to the individual organisation, but care should be taken to avoid blurring a team's aims to the detriment of its function. Vulnerability assessment should be an on-going, business-as-usual function to detect known vulnerabilities that could have arisen through missed patching or misconfiguration. Threat intelligence should be responsive to evolving requirements – with clear tasking.

Interaction between threat intelligence and vulnerability assessment is often desirable. If, for example, the threat intelligence team identifies that a particular vulnerability is being actively exploited, especially when there are indications that exploitation is occurring within the organisation's own industry sector, it should trigger an out-of-band vulnerability assessment to ensure that any such attack on the organisation will fail. In this example, monitoring teams should be advised to look for indications of exploit attempts, as this could reveal an attacker's intentions – highly useful information.

Collecting, Using and Sharing

The four subtypes of threat intelligence proposed by this paper are very different in terms of their collection, analysis and consumption. This section will address each in turn, providing guidance on how to collect, use and share the intelligence to best effect.

Strategic Threat Intelligence

Definition

Strategic threat intelligence is consumed by high-level strategists within an organisation, typically the board or those who report to the board. Its purpose is to help strategists understand current risks, and to identify further risks of which they are as yet unaware. It deals in such high-level concepts as risk and likelihoods, rather than technical aspects; and it is used by the board to guide strategic business decisions and to understand the impact of the decisions that are made.

The intelligence is often in the form of prose, such as reports or briefings – either at meetings or one-to-one with senior management and board members. It has a high-level and business focus that is used to guide strategy.

How to Set Requirements

C-level executives (CEO, CFO, CIO, etc.) and the board require a level of understanding as to which decisions might be linked to cyber risks. In lieu of this understanding, threat intelligence team members need to ensure that they are themselves aware of the sorts of decisions being made, and proactively advise senior management and the board.

Decisions that might have cyber risk implications should be used for setting requirements. These decisions could be related to adversaries, such as when entering foreign markets, partnering or supporting ideological groups, making ideological statements / taking ideological positions, purchasing or being purchased by foreign organisations, or setting up foreign offices. Alternatively, the decisions might be related to information exposure, such as strategic directions that affect how information is stored and used within the organisation – for example, outsourcing an IT function. It is important for boards to understand that such decisions can affect risk. The threat intelligence function can then be tasked with finding answers to help the decision-making process.

Setting appropriate requirements is crucial to a good outcome in all forms of intelligence, but this is particularly true of strategic intelligence. If the threat intelligence team is tasked with the following: "We are going into business in this new country; tell us their capability and which groups will attack us, and how," it will not result in a useful intelligence product. The requirement blends a number of different intelligence types and, significantly, it is simply not possible for a non-government entity to legally gather much of the operational information required to discover whether groups would attack. To do so would probably require communications interception and human sources within the attack teams. There are similar issues when it comes to identifying an attack group's capability and further challenges in deciding what the information actually means to the board.

A better requirement would be, "We are going into business in this new country. Do we believe that is likely to result in attacks, what are the typical outcomes of those attacks, and what would be the cost or effort required to appropriately defend against such attacks, should we choose to?" A requirement phrased in such a way allows a threat intelligence team to prepare realistic advice based on what can be ascertained, rather than seeking difficult or impossible-to-obtain answers, with the probable result of purchasing suspect and hard-to-corroborate information from a provider.

Those involved in threat intelligence will need to work with the intended recipient to 'drill down' on what, exactly, they need to know – and with how much confidence – before progressing to the next stage of the TI cycle.

How to Collect

As strategic threat intelligence is high-level, the majority of the collection sources will be high-level as well. They are likely to involve:

High-Level Geopolitical Assessment

Trends in country strategies, ambitions, priorities and other high-level information can help inform strategic analysis. Is usually coupled (at analysis stage) with observations of malware or attacks thought to be related to the country, to create a picture of cyber activities.

Information to feed analysis will therefore come from high-level sources. These might include an analysis of policy releases by nations or groups of interest, news stories in domestic and foreign press, and news stories in subject-specific press, such as financial papers. Articles published in journals by high-ranking persons in the nation or group of interest can also provide useful indications of intent or capability.

Much of the information needed for analysis can be collected from what is commonly called open source intelligence (OSINT), in other words searching publically available or 'open' sources. It has been reported that in mixed-source reports (i.e. information from a number of different sources, including both OSINT and secret sources), open source intelligence regularly provides 80% of the content[11]. This can be a highly rewarding area of collection and should be actively pursued[12]. For deeper insight, organisations are advised to ensure they are not limiting searches to their own language or with a bias towards their own language[13]. Using search engines from the nation of interest and the increasingly powerful translation engines provided by the likes of Google and Microsoft Bing can enable searching and collection of news stories, articles, policy, etc. directly from the nation or foreign group of interest. Even the foreign language versions of Wikipedia can contain far more relevant information than does the English language version.

Staring Into the Abyss

Organisations should be aware of the searches and investigation techniques that can be detected by an attacker, as particularly astute attackers might well be looking for indications that their activities have triggered an investigation. For example, attackers can monitor VirusTotal (www.virustotal.com) to ascertain when malware they have created has been uploaded – suggesting that someone is investigating that malware[14]. Even visiting a website can tip off the owners that someone has visited a certain page, and it's not unknown for interesting-sounding pages to be created, simply to provide an alert when they are accessed. Technical staff who understand privacy and digital footprints are best placed to create guidance for acceptable (and unacceptable) investigation techniques, with the aim of helping to train investigators.

Collection (and analysis) of strategic information can be challenging and it does require a socio-political mindset rather than a technical one. With such a huge number of sources available, identifying those that are useful – and reliable – can be problematic. How do you find the military journal in which a senior commander once published an article on what he saw as the future of cyber-enabled conflict? How do you then establish whether that commander's viewpoint represents a trend or intention, or just one man's dream? Hence many organisations prefer to purchase analysis from strategic intelligence providers.

These providers attempt general collection and analysis to create products they feel will be useful to a large proportion of their clients. However, since they are producing a relatively broad product, the purchased analysis must be treated by an organisation as collected information (i.e. not as analysis), which is then itself analysed by the organisation's own threat intelligence team. Another issue for analysts to consider is the reliability of the information. Strategic intelligence is hard to 'do' well and some vendors have not been above selling unreliable or poorly verified intelligence – and then citing the need to protect their sources if a client challenges the collection or analysis. Careful analysis of these products is therefore important.

Security Industry White Papers

A major source of information to help inform strategic analysis comes in the form of white papers and blog posts covering particular attack campaigns or threat actors. An increasing number of such papers are being released and the information can help to build a picture of attack groups and their targets. Reports typically lend themselves to tactical and technical analysis, yet can also contribute to strategic intelligence and are therefore worth including in the collection process.

Human Contacts

Human contacts can be extremely useful when collecting for strategic intelligence. Contacts at similar organisations, or organisations in other sectors that have been in similar situations, can provide valuable information on attacks and threats. This can be seen as the receiving side of 'Need to Share'.

The depth of information provided by contacts will no doubt be proportional to the level of trust in that relationship, and so these relationships are worth building and maintaining (see 'Need to Share'), even when no information is currently being sought. Information should be treated sensitively and, unless there are specific reasons to do so, it is often better not to attribute the information received to particular individuals. If it is necessary to identify the source, then that information should itself be reliably protected. This is important in engendering bilateral trust, so that individuals will feel inclined to help your threat intelligence team build a picture of the threats.

Peers, Not Agents

Human sources are a traditional focus of intelligence, and are often considered when mapping intelligence theory onto threat intelligence. However, it's important for organisations to take care when using human sources, and not be drawn into 'playing spies'. Attempting to cultivate either human sources in an attack group, or those who can inform on an attack group, is – for a private organisation – ethically dubious, to say the least. It also risks interfering with existing investigations. Organisations are strongly advised instead to focus on human sources in the form of peers, friends and contacts in relevant organisations, with whom it's possible to build mutually beneficial sharing relationships.

How to Analyse

Strategic analysis is a long-established but complex field – and it can be a highly challenging one, since it's rare to work with absolutes. Instead, it's more usual to deal with trends, observations and perceived intentions[15]. Analysis for strategic intelligence purposes requires more expertise than in other areas of intelligence and probably a wider range of collected information, of varying levels of relevance and reliability.

Analysis and collection are likely to be tightly linked, with lines of inquiry and trends identified and then tested by collecting new information.

Organisations will often hire people with expertise in traditional intelligence or socio-political analysis and then teach them the cyber components necessary to perform effective subject-specific analysis. Alternatively, technical staff in the threat intelligence team can be trained in analysis; however, this latter approach tends to require a great deal of reading and understanding of the sociological and political background.

Meta-analysis is often a useful component of strategic intelligence, whereby results from a range of analyses are combined and reconsidered in an attempt to yield new intelligence. This can be particularly useful with, for example, technical white papers that come through the team. By analysing the white papers, trends might be identified, such as a particular piece of malware that is increasing in complexity and code quality with each iteration – suggesting active investment and development by the responsible group. In this case, developments in other areas of the group's capability are likely, and potentially in its target selection, and collection can seek evidence to support or disprove this theory.

Attribution in cyber attacks is often difficult[16]. Hence, while they can sometimes prove useful, any stated attributions in a report should be regarded with some caution. The reported range of industries to be targeted should also be treated cautiously, unless the methods for ascertaining victims is open to scrutiny. On multiple occasions, MWR InfoSecurity has investigated attacks on clients that were very likely to be part of campaigns reported by others, yet the victim's industry did not appear anywhere on the lists of targeted industries. This suggests that, in many cases, reports of threat actor activity have a limited view when it comes to the extent of the campaign.

Production and Use

The threat intelligence team should be working to tight requirements in terms of what to produce. Strategic intelligence is best used by high-level decision-makers, who will be consuming a great deal of information as part of their decision-making process, hence the product will generally need to be short and concise. In some cases, it might be no more than a couple of lines.

Informal strategic intelligence requests should also be expected and supported. Once the team has become experienced, it is more likely to be asked for analysis and comment informally – either by the board or by the security function of the organisation. In such cases, the product of the intelligence might simply be an email.

Products are typically focused on business impact and risk, while a discussion of technical details is best avoided as it's rarely useful to the board. Where the accuracy of the information can't be guaranteed, this should be indicated to the product's consumer and, where appropriate, a confidence level in the information given. Consumers will also need to understand what is, and isn't, a realistic request or task.

Confidence is Key

In intelligence analysis there will rarely be certainty. In many cases, analysis will have to focus on a small number of sources, or even a single source of potentially questionable quality. Hence communicating the confidence of a statement is of key importance, with an agreed language consistently used by those producing threat intelligence reports – and understood by those reading it. Organisations are advised to maintain an internal document that explains exact definitions of such terms as "we know", "we suspect", "we believe", "high possibility", "may", and so on. Staff with previous experience in traditional intelligence are likely to have an advantage in helping to design such vocabularies. As an example, page 6 of the 2014 'Targeting US Technologies' report by the US Defense Security Service (DSS)[17] gives descriptions of how confidences are derived before using the terms. For "High Confidence" statements, phrases such as "well-corroborated information from proven sources", "minimal assumptions", or "strong logical inferences" generally indicate that the DSS based its judgements on high-quality information, and/or the nature of the issue made it possible to render a solid judgement.

How to Evaluate

Strategic intelligence should be evaluated as to how well it supports senior decision-makers: is it accurate, impactful and timely?

Accuracy can be difficult to assess in absolute terms, as we might never fully understand a remote situation, but it's usually possible to assess a product in terms of the team's stated belief in its accuracy. (If the threat intelligence team believes it's 99% accurate in all reports and the last three reports have proved to be inaccurate, however, there could be issues with the team's faith in its own work.) As for the product's impact, the question is how useful the product is in supporting decisions and how directly it matches the stated requirements. Timeliness is simply whether the information is delivered to the consumer quickly enough – and in a useable form.

How to Share

Strategic threat intelligence itself is rarely shared, as the details could well reveal the organisation's plans. Generic strategic intelligence, meanwhile, is unlikely to be of much use to other organisations.

Instead of sharing strategic intelligence, organisations are advised to focus on sharing other types of intelligence. The threat intelligence team at another organisation can then analyse what is shared to turn it into strategic intelligence relevant to its own business.

Some sensitivity will still be required. For example, if your organisation is operating in country X and it receives an attack believed to be conducted by country X, that could be very useful information for other companies operating in, or intending to operate in, that country. Sharing that information will not leak strategy but might, if shared incorrectly, lead to other problems: for example, political issues with the government of the country in question. Sharing needs to be carefully evaluated to ensure that such risks are mitigated.

Operational Threat Intelligence

Definition

Operational threat intelligence is actionable information on specific incoming attacks. Ideally, it informs on the nature of the attack, the identity and capability of the attacker – and gives an indication of when the attack will take place. It is used to mitigate the attack: for example, by removing attack paths or hardening services.

How to Set Requirements

Consumers of operational threat intelligence naturally desire intelligence on all groups that might attack them (with corresponding details of when and how they will attack). However, it is important that organisations focus on operational intelligence that can feasibly be obtained, as in-depth information on nation-state attackers is not a realistic requirement for private companies.

Collecting operational intelligence requires penetrating the attacking groups or their communications, and requirements should be limited to groups where this is possible. Some organisations might find they are targeted by groups that communicate relatively openly about their intended attacks. These are likely to be ideologically motivated groups, rather than financial or espionage-focused groups that typically communicate using far more secure means.

Requirements should therefore be based around producing intelligence on specific groups, supported by consultation with the threat intelligence team to ensure the requirements are reasonable.

How to Collect

Collecting operational intelligence in traditional domains will include such activities as recruiting human sources within groups, and the compromise of the groups' communications. However, operational threat intelligence for private entities is necessarily restricted, as the majority of methods of collecting such intelligence would be illegal – or at best immoral – for a private company. Organisations intending to conduct monitoring operations are advised to take legal advice before doing so. Monitoring open communications by groups is more likely to be legal than other methods, although organisations are nevertheless recommended to seek advice in these cases too.

Activity-Related Attacks

In some cases, recurring attacks could be related to real-world events, such as the activities of an organisation or those the organisation is related to, supports, or finances. This is a well-understood phenomenon in physical security, where – for example – premises are attacked in response to certain triggers, and the same can be true of cyber attacks. Analysts should collect information regarding attacks, particularly those that are seen to repeat – such as DDoS attacks – and attempt to analyse whether they can be correlated to activities or events. Indicators that the attack was about to begin should also be sought: for example, social media posts.

Chat Rooms

Some ideologically motivated groups discuss plans in chat rooms. However, groups are often aware that these rooms are monitored and hence discuss more targeted operations in private chat rooms. It can be difficult – operationally and legally – to obtain access to these rooms, meaning that many organisations will be limited to the more public rooms that are used to discuss larger-scale attacks: typically those that require a large number of participants, such as DDoS.

> Organisations intending to actively collect by, for example, participating in chat rooms or forums, might wish to ensure that such activities are conducted discreetly. This could mean using non-attributable IP addresses and preventing the leakage of other indicators.

Organisations should be aware that some chat rooms used to discuss wide-ranging attacks are in foreign languages, pushing up the cost of collection.

There are threat intelligence vendors that sell collected information from both public and private rooms, and potential buyers need to ensure that the information being purchased is both legal and relevant to their business. The temptation to have vicarious access to 'closed sources' can sometimes override good judgement when it comes to whether the information is actually useful.

Social Media

Another means of gleaning operational intelligence is to monitor social networks for mentions of your organisation in relation to a planned attack. For example, Twitter has a well-documented API that can be used to set up a streaming feed[18], where all public tweets that match specific search terms are delivered through the API – and can then be consumed and filtered by scripts. Alternatively, the feeds of specific individuals who might tweet threats against your organisation can, once identified, be followed.

Some vendors offer services that monitor social networks for mentions of your organisation, with the aim of reporting potential attacks.

How to Analyse

Collected details of attacks are worth scrutinising for signs of activity- or event-correlated attacks. In other words, analysts should attempt to identify whether they are part of a pattern related to events or activities, or to reported activities in the news. It's important to be aware that attacks could be related not to activities by the organisation itself, but to those of partner organisations or groups/individuals that are in some way linked to the organisation.

Where operational threat intelligence is not events-based, it is likely to focus on social network posts and chat room conversations. These sources will typically be high volume, with a great deal of 'noise', hence organisations are advised to develop scripts that identify messages of interest. These scripts will require evaluation and modification until they can produce actionable information. It can also be useful for analysts to hunt through the collected information manually to identify indications of attacks, and then develop scripts that ensure similar messages would be extracted in future.

Groups sometimes use either codes or simply slang that obscures meaning. In many cases, these involve simple substitution, where a slang name is used for a certain target or type of attack. Analysts will want to ensure that they keep up to date with codes and slang, and that analytical scripts and wordlists are likewise updated.

Another thing to be aware of is that individuals tend to change aliases on a regular basis and analysis needs to take that into account. This might require more advanced tracking, such as linguistic analysis, or timeline analysis (where the disappearance of one 'person' is swiftly followed by the appearance of another).

Deeper analysis can be effected by combining operational threat intelligence with other forms, for example tactical, to ensure that there is understanding of

groups' methodologies and capabilities. This can be combined into the operational threat intelligence output (the report or notification) to provide more information on the expected form and scale of the attack.

Production and Use

Operational threat intelligence can sometimes provide warning of future attacks, such as a planned DDoS at a specific time or at the same time as another event. This provides the opportunity to ensure appropriate defences are in place that will both withstand the attack, and monitor/evaluate the nature of the attack in the hope that it will leak information about those behind it. However, intelligence is rarely perfect, so operational threat intelligence products should be phrased in such a way as to take this into account – with an appropriate indication of the level of uncertainty in the product.

Often, there is no significant warning of an attack, which might be only minutes away. To deal competently with these circumstances, organisations would do well to plan for, and rehearse, reacting to operational threat intelligence in short timescales. This is likely to involve readily accessible contact details for on-call staff and service providers, with escalation paths pre-planned.

How to Evaluate

Operational threat intelligence is relatively easy to evaluate. If the intelligence was able to forecast an attack and, as a result, the attack was partly or wholly mitigated in time, then the intelligence was successful. It is more likely, however, that incoming attacks were not forecast and it can therefore be useful to conduct some root cause analysis.

In the majority of cases, the conclusion will be that collecting the necessary information to provide forewarning would not have been possible or legal. Where it would have been possible, an investigation of the collection and analysis process will help to identify opportunities for future improvement. However, on-going operational threat intelligence efforts should be strictly evaluated as, despite the alluring promise of such intelligence, in reality there are few circumstances where good, actionable information is obtained – and resources might be better focused on other types of threat intelligence.

How to Share

Operational threat intelligence can be shared with others if it will provide them with advance warning of attacks. For example, if – during collection efforts – it's noticed that the groups under observation are planning to target another organisation, then that organisation can be alerted to the threat. In such instances, it can prove difficult to find reliable contact details for the appropriate individual to warn, in which case CiSP might be a useful route. The individual could have their own account on CiSP, or other members of CiSP might well have contact information.

Tactical Threat Intelligence

Definition

Tactical threat intelligence can be one of the most useful forms of intelligence in terms of protecting the organisation. It is defined as information that concerns the tactics used by threat groups – including their tools and methodologies – and is often referred to as Tactics, Techniques, and Procedures (TTPs).

The aim of tactical threat intelligence is to understand how threat actors are likely to attack the organisation, and to map this understanding to the ways in which the attacks can be mitigated or detected. For example, the reports that many groups use Mimikatz to extract plain text credentials from compromised hosts should inform policies on how administrators perform remote administration of machines, and how accounts are configured on the domain.

Tactical threat intelligence is consumed by defenders such as architects, administrators and security staff.

How to Set Requirements

Requirements should focus on understanding the tactics used by threat groups, particularly those groups that are believed likely to target the organisation. The requirements might relate to collection events, such as "Provide tactical threat intelligence to the relevant consumer three days after the release of a report on CiSP", or they might be driven by planned maintenance, development or purchasing. For example, if a domain refresh is planned, the threat intelligence team could be tasked with providing information to the domain administrators and architects on attacks seen against domains and how they can be mitigated.

How to Collect

Collection is likely to come from mid-level sources, such as reports into attack campaigns. Tactical threat intelligence requires focusing on the tactics of threats, hence collection should focus on sources that give insight into these tactics.

Attack Group Reports / Campaign Reports

In the current environment, reports on attack campaigns or specific actors are the most commonly available sources able to provide details on tactics and tooling, and efforts should be made to collect all that are available. Keeping abreast of documents posted on CiSP that have been curated by the community is an easy and effective way to collect the majority of reports. Alternatively, for those without access to CiSP, a Git repository is available at https://github.com/kbandla/APTnotes; although it should be noted that the content of the repository cannot be guaranteed.

Malware

Analysing malware samples from groups that have attacked the organisation or similar organisations can yield information on tactics and tools. Malware can be collected from feeds (either free or paid-for) that accumulate and distribute malware, while a number of websites exist that provide malware samples. Alternatively, a number of groups conduct malware analysis and release reports, which can be collected.

Incident Reports

Reports of incidents can be useful in informing analysis for tactical threat intelligence. In some cases, these will be formally published incident reports such as appear in forums. However, informal reports can also be useful and worthy of collection. These can take the form of conversations with defenders or investigators on the nature of attacks and the trends in methodologies.

How to Analyse

All collected sources should then be analysed to extract indications of tactics. White papers and reports can be deconstructed to identify the use of particular tactics and tools. Specifically, analysts should attempt to identify:

Modus Operandi and Exploited Issues

Analysts should attempt to understand how threat actors are operating when attacking networks. For example, how did the attackers initially gain access, how did they escalate privileges, how did they move laterally in the network, how did they gain access to the data they sought, and how did they extract the data? For each step in the process, attack groups – and even individual attackers – will have patterns of behaviour. Typically, these patterns exploit common issues on corporate networks, such as flat networks with no segregation, or privileged accounts used to log into workstations. Analysts should seek to identify the issues exploited by attackers and ascertain whether those issues are present on their own networks[19].

Where technology refreshes are planned, such as new file systems, networks or domains, analysts would be wise to attempt to understand common attacks against those systems, to provide guidance to systems administrators and architects in making the systems more secure from day one.

Tools

Intelligence on the tools used by attackers can inform both detection and protection, and incident response and protective technologies should aim to identify them as clearly as possible. It is highly unlikely that crude detections such as MD5 sums will work, and so detection should focus on methods such as well-written YARA rules. Attackers commonly modify open source tools to avoid trivial detection.

Understanding the capabilities of the tools in use is also important. Analysis of malware can yield such intelligence, in other words the information they might be able to obtain – for example, Mimikatz can yield cleartext passwords. Some attack groups have been seen with tooling that exploits MS14-058 and provides local privilege escalation, giving attackers higher privileges on vulnerable systems[20]. Many attackers use

publically available Remote Access Tools (RATs), which is intelligence in itself, while others develop their own. Custom RATs should be analysed to identify the information that attackers are trying to obtain, and their capabilities in this respect, to support detection and hardening.

Analysis of tools – and different versions of those tools – can give an indication of how advanced a particular actor is[21]. A remote access toolkit written by a single author in Python, a highly accessible and easy-to-learn language, suggests a far less capable (and consequently less funded) adversary than one written in efficient C++ (a more challenging language to learn) with a complex modular and extensible framework. Any change in an attack group's tools could indicate a change in its intentions and resourcing[22].

Analysts should aim to understand what the tool can do, what it might be designed to obtain, what the tool says about the skill of its creators and users and, potentially, who wrote it – although this can be difficult to ascertain.

Communications Techniques

Analysts should attempt to understand the C2 and data exfiltration channels used by attackers, and map that information onto their organisation to understand whether it would be detected or prevented. In many cases, attackers use HTTP or simple communications methods, but others are more complicated; for example, some attackers will use DNS as a command and control channel. Details of communications should be extracted from collected sources.

Forensic Avoidance Strategies

Analysts should be seeking insight into how attackers are attempting to avoid detection in their tools and actions. Although many attackers do not make particular efforts in this regard, a number take significant trouble to avoid or delay detection. Analysts are advised to identify the tactics used and establish how defences can be adapted to overcome these strategies.

Production and Use

Tactical threat intelligence should provide advice to defenders, including network architects, domain administrators, system administrators and incident analysts. Realistically, organisations need to allocate security budget and resources carefully, and tactical threat intelligence can help this process by identifying the areas in which security investment will mitigate tactics used by genuine threats.

The products of intelligence should therefore attempt to prioritise fixes for the organisation's security and inform defenders as to how crucial it is to adapt defences – as well as the likely impact of failing to do so. In consequence, providing an easily consumable product will require a degree of understanding of the organisation's network plus, potentially, liaison with consumers during the production process. In some cases, it might even be appropriate for the threat intelligence team to provide the fixes, such as the registry keys that will need to be changed to prevent a certain type of attack.

A frequently encountered problem is that network and server operations staff typically run close to full capacity, simply to keep the organisation's infrastructure running to a 'business as usual' standard. Buy-in at senior level will therefore be necessary to ensure that the threat intelligence product is acted on, and that time and resources are committed to implementing the required changes. In some instances, it might be appropriate to postpone changes to coincide with planned future refreshes – in which case, increased monitoring is advisable in the interim period to identify attacks.

Effective tactical intelligence can also aid incident response, as if an attacker's methods of operating are understood, responders can validate their observations against what has been seen previously. Where responders are having difficulty following the attack through the network, tactical intelligence can help to indicate where the attacker might have gone or what they did next.

How to Evaluate

The evaluation of tactical threat intelligence should include an assessment of how well it feeds into the defensive processes, and whether the hardening recommended by the threat intelligence team has mitigated or allowed the detection of particular attacks.

Where successful attacks have occurred, the methodologies of the attackers should be investigated – and a conclusion drawn as to whether the organisation should have been aware of, and mitigated, the attack. For example, if the attack used a previously unseen or rare technique, then it's unlikely that collection would have been able to provide intelligence.

How to Share

Sharing tactical threat intelligence helps everyone in the community. Individuals who do share such intelligence often find it encourages others in the community to come forward with similar reports, providing yet more useful threat intelligence.

When an organisation has been attacked (regardless of whether it was successful), it is strongly advised that an incident report is released. Where possible, this should include information on tooling, tactics and methods of attack. The section 'Need to Share' covers the ways this can be done while minimising any negative impact on the organisation.

Technical Threat Intelligence

Definition

Technical threat intelligence comprises technical details of an attacker's assets, such as tools, command and control channels, and infrastructure. It differs from tactical threat intelligence in that it focuses on specific indicators and rapid distribution and response, and therefore has a shorter usable lifespan. The fact that an attacker uses a particular piece of malware would be tactical intelligence, while an indicator against a specific compiled example would be technical intelligence.

Common examples of technical threat intelligence include MD5 sums of malware or document lures, subject headers of phishing emails, IP addresses for C2 endpoints or domain names used by C2. Ideally, these indicators should come from active campaigns that are currently being experienced by other organisations. By rapidly including these indicators in defensive infrastructure such as firewalls, mail filtering devices and endpoint security solutions, organisations can seek to detect attackers – either when they first attack, or in the early stages of an attack. By searching logs of previously observed connections or binaries, historical attacks can also be detected.

A challenge frequently reported by organisations attempting technical threat intelligence is that the sheer quantity of data can quickly become overwhelming. In this case, the allocation of resources needs to be carefully considered, with the organisation perhaps becoming more selective in the data it collects, or instead deciding to build/purchase large analytics platforms to cope with the quantity of data. However, it's important that resource allocation and capability development is continually balanced against an evaluation of the benefits of technical threat intelligence. It might be found that greater benefits will come from investing in other forms of intelligence.

There is much commentary in the security community as to the usefulness of technical threat intelligence, with some arguing that it's a highly effective way of preventing and detecting compromise, while others doubt its usefulness. The latter group likens it to antivirus signatures, since attackers can trivially adapt to ensure that their tools are not recognised. There is also a concern that large amounts of data sold as technical threat intelligence lack contextual information, and hence cannot feed higher analysis and appraisal of sources.

> A key failing of technical threat intelligence is that it's relatively simple for an attacker to target a specific organisation in a way that ensures no pre-existing indicators will have been available. Modified malware, custom network infrastructure and obscured C2 communications do not require great skill or resources, but still bypass technical threat intelligence efforts.

Technical threat intelligence should be consumed in an automated fashion and placed into rulesets for network security devices and endpoint security solutions.

How to Set Requirements

Setting effective requirements for technical threat intelligence can be difficult, as it's very tempting to allow collection to drive the process. Hence requirements are often set along the lines of "Process and react to feed X". This places a lot of faith in 'feed X', as it suggests the feed is sufficient to meet higher requirements. Instead, these higher requirements should be explicitly set, for example, "Identify phishing emails being sent to other companies in our sector and assess whether they are also being sent to our staff". Another requirement might be to "Identify IP addresses that are seen in attacks on similar targets to ourselves, and ensure we are not connecting to them".

Requirements should have a retrospective, as well as a current, focus. In other words, where possible, organisations are advised to check whether indicators can be observed historically; for example, has the organisation connected to an identified IP address at any time in the past year – not just this week.

How to Collect

There are various types of data that can be classed as technical threat intelligence, with some indicators harder than others for attackers to modify in their attempts to defeat signatures[23]. This section deals with the more commonly sought types.

Indicators are commonly collected from feeds (paid-for or free), provided by third parties as a result of their investigations or derived by an organisation's internal investigators. For example, once an attack has been detected, a good deal of investigation can be done online (while ensuring the threat actor is not alerted – see 'Staring Into the Abyss' box-out on p14) to derive other indicators of attack[24].

Malware Indicators

As a large proportion of attacks involve malware, malware indicators are often sought as threat intelligence. The most commonly offered indicators are MD5 or SHA-1 hashes of binaries believed to be suspicious. However, it is trivial for an attacker to modify their malware to avoid detection: a single bit changed anywhere in the binary will result in a different hash, and so many adversaries will use open source tools and make subtle modifications to change the hashes.

Indicators such as created registry keys or file artifacts can be more useful, as they are less commonly changed by attackers. However, it is still possible for the adversary to give dropped files a random or pseudorandom component in their name (for example).

Figure 5: How hard is it for an attacker to modify malware so that a specific signature is no longer recognised?

Many reports of campaigns will contain indicators that can be consumed as technical threat intelligence. Unfortunately, these indicators will often be included in PDF reports, hence collection involves copy and pasting the indicators before formatting them correctly. It can be worth contacting the report authors to ask whether machine-consumable indicators are available.

There are a number of freely available and commercial feeds of malware indicators. Before collection, the content of feeds should be evaluated to ensure they contain actionable data, as should the volume of data – in case collection overwhelms analysis by virtue of the sheer quantity of indicators to be consumed.

Notable Free Feeds

CiSP portal, maintained by the UK Government, application process required. **www.cisp.org.uk**

Critical Stack, aggregation of freely available feeds by a consultancy. **https://intel.criticalstack.com**

Open Threat Exchange, a forum to exchange indicators maintained by AlienVault, a SIEM vendor. **https://www.alienvault.com/open-threat-exchange**

Network Indicators

A number of different network indicators can be collected as technical threat intelligence, as malware frequently needs to communicate with the attack group. Attackers will operate nodes from which to conduct attacks and will sometimes use the same node for multiple victims. An IP address that has been observed by others functioning as a C2 node can therefore be a useful indicator. However, attackers will often use different IP addresses, changing C2 nodes as they are discovered or as computers become unavailable. Malware will attempt to connect to a domain name, which can then be pointed to the IP address the attacker is currently using. Where malware is using a hardcoded domain, this can be a relatively useful indicator but it's quite common for malware to use a domain generation algorithm to avoid the need to connect to the same domain twice. In such cases, a domain name has little value as an indicator.

Another potential network indicator can be found in C2 communications; for example, the 'Havex' malware was so called as its C2 communications included the term 'havex'. Such indicators can be more useful to collect, as they require more effort for attackers to change.

Figure 6: How hard is it for an attacker to modify malware communications so that a specific signature type is no longer recognised?

As with malware indicators, network indicators can be found in white papers and reports. Again, a number of freely available and paid-for feeds exist; one feed of particular note is the freely available daily C2 list from CiSP.

Email Indicators

A large number of attacks start with a phishing or spear phishing attack containing either a document exploit or simply malware disguised as something benign, so email indicators can provide useful threat intelligence. Attackers will often ensure that emails are either targeted or semi-targeted, hence generalist feeds of spam email subjects will be less useful than details of phishing emails sent to similar organisations.

It is worthwhile contacting similar organisations in an attempt to establish relationships in which the subject headers or other indicators of suspicious emails can be shared. It is important, however, neither to share nor to receive shared phishing emails themselves, as there is always a risk they will be opened. Indicators should be extracted and only the indicators shared.

How to Analyse

The analysis of technical threat intelligence will almost always be automated or heavily automated. This is because indicators will often have a short usable time before attackers makes changes, hence rapid filtering is important. Technical threat intelligence is also typically high volume and allows little meta-analysis, so is well-suited to analysis by machine.

Conversion Between Formats

Technical threat intelligence can be transmitted in a number of competing formats (STIX and OpenIOC are popular choices[25]), and tools exist to convert one format into another for easy consumption[26]. Typically, IOC formats are XML-based and readily parsed by scripts into a format suitable for toolsets. Some technical threat intelligence is offered in formats that are native to specific tools – for example, Snort or Bro IDS – and will not require conversion. In some cases, however, indicators might simply be a list of hashes or IP addresses that require formatting.

It is therefore recommended that at least one member of the threat intelligence team is able to script or program competently, so that conversion scripts can be written for new sources as they become available.

Technical Threat Intelligence Libraries

A concept that has emerged in recent times is using a 'threat intelligence library' to store indicators and seek links between them. This approach also allows an organisation to detect attacks within logs and packet captures that have been fed in[27]. These libraries are effectively large repositories that often use so-called 'big data' technologies (such as data warehousing and graph analysis) in an attempt to draw links between types of technical threat intelligence, allowing quicker response to detected threats, as well as an historical record of received IOCs.

A number of vendors offer paid-for products in this area. The Collective Intelligence Framework (CIF)[28], meanwhile, is an open source project that focuses primarily on network indicators. It is able to consume a variety of sharing formats, and allows an organisation to query and output rules in formats suitable for network appliances[29].

Production and Use

The effective 'product' of technical threat intelligence is the ruleset developed to enable network or endpoint tools to detect identified malware. There should be a smooth process for pushing rules to devices and software, and a well-established rollback protocol. As it's not possible to vouch for individual feeds, the potential exists for a benign IP or MD5 hash critical to business function to end up on a blacklist. Thus the ability to roll back offending rules should be well understood.

Malware indicators such as MD5/SHA-1 hashes can be detected either at network ingress or on the host. Detection at network ingress will require filtering or monitoring of downloaded and emailed files. Bro IDS is an open source tool that can facilitate the extraction and hashing of binaries from network traffic[30]. Detecting indicators on hosts (either hashes or more complex indicators such as registry keys) is likely to require endpoint security tools with this feature.

Network indicators such as IP addresses and domain names can be placed in firewall 'deny and log' rules, or as rules in network-based IDS products. This will create alerts for any outbound connections to those remote endpoints. Organisations should be sure that they are added to outbound rulesets, as inexperienced staff occasionally forget that connections will be initiated by malware and hence will appear as outbound connections. More complex network indicators, such as the internal workings of C2 channels, will require network IDS or network AV products for instances in traffic to be detected.

Email indicators, such as subject headers, will require email interception. Organisations using email filtering services can add indicators to blacklists, or a suitably placed Bro IDS instance can be used to extract and filter subject lines or other indicators.

Organisations are strongly advised to use technical threat intelligence to search for historical compromise, either by giving indicators to incident responders or by using similar tools to those used to search for current compromise. Searching for historical compromise will require records of network connections, binaries and emails received.

How to Evaluate

Technical threat intelligence can be a complex endeavour – not to mention expensive, if feeds and analytical solutions are purchased commercially. It should therefore be rigorously evaluated: specifically, the number of prevented attacks that would not have been prevented by other means.

Many organisations appear to focus significant proportions of their threat intelligence effort on this one area. This can prove inefficient, as by the nature of technical threat intelligence collection, attackers will always be able to avoid detection by creating a more custom-targeted attack. Evaluation should therefore consider whether resources would be better applied to other types of threat intelligence.

How to Share

Technical indicators should be shared with other organisations wherever possible. This can be done through forums such as CiSP, trusted third parties, or via direct sharing. Where possible, indicators should be shared in a machine-readable format for which other organisations' threat intelligence analysts can write parsers or converters, if their tools do not accept that format[31].

Phishing emails are more usefully shared with similar organisations in the same sector, as they are often customised to the sector. It's therefore recommended that this information is shared with trusted third parties that have specific knowledge of the sector, via sector-specific forums, or directly with similar organisations. As mentioned previously, the phishing emails themselves should never be shared, only the indicators.

Summary

Threat intelligence is at high risk of becoming a buzzword. With so many disparate offerings and so much pressure to be 'doing' threat intelligence, organisations risk investing large amounts of time and money with little positive effect on security.

However, by taking threat intelligence back to its intelligence roots and applying the same strict principles, a far more effective strategy can be devised. As is the case with traditional intelligence, tackling cyber threats demands rigorous planning, execution and evaluation. Only then can an organisation hope to target its defences effectively, increase its awareness of threats, and improve its response to potential attacks.

Much can be learnt from studying successful threat intelligence programmes and, just as usefully, the common mistakes underlying threat intelligence programmes that fail to deliver genuine business benefits.

It quickly becomes clear that effective threat intelligence focuses on the questions that an organisation wants answered, rather than simply attempting to collect, process, and act on vast quantities of data. Yet it's vital to be asking the right questions in the first place. Hence this paper looks in detail at the cycle of setting requirements, collecting and analysing data, turning the results into a consumable product and evaluating the usefulness of that product – which then feeds back into asking 'better', more useful questions for the future.

There is also value in breaking down threat intelligence into subtypes, depending on who uses it, where it comes from, and how much business benefit it really offers. By relying too heavily on one sort – or the wrong sort – of threat intelligence, organisations risk wasting effort, while leaving themselves vulnerable to attack.

Resource and budgeting will always be an issue for commercial enterprises, and it's important to realise that the most useful sources of threat intelligence are not necessarily the most expensive. Enormous value can be gained – for example – from sharing threat intelligence with other organisations, and one-to-one human contacts can be one of the simplest, yet most effective, sources of actionable information. This paper therefore looks at the benefits to be gained from sharing threat intelligence, and how to go about it without exposing the organisation to unnecessary business risk.

Quick Wins

This section offers examples of some productive steps that can be taken with only a minimum of staff and budget. It assumes no specific current security infrastructure, such as SIEM tools, IDS tools or log aggregation and analysis. It also assumes no current official threat intelligence function within the business.

Organisational

- Identify where threat intelligence processes might be taking place unofficially, and assess how they could be better supported.

Strategic

- Work with senior management to identify current cyber threats as they perceive them. Conduct open source intelligence to determine whether those threats have been realised in the past, and set up Google Alerts or RSS Feeds to alert on new information.
- Liaise with peers in organisations in the same industry sector to determine whether there are other threats that your organisation has not yet recognised.
- With the aid of senior management, create a list of all actors (companies, campaign groups, countries, etc.) that would benefit from access to your sensitive data – or from your inability to function effectively.

Operational

- Prepare a list of names and contact details (including out-of-hours details) for the people it would be necessary to contact if your organisation received notice of an impending attack.
- If regular or repeat denial-of-service attacks are being seen, use Google to search for your organisation's name, but limited to those dates immediately preceding the attacks. The aim is to determine whether negative coverage is leading to the attacks. If not, attempt to identify other factors that might be triggering the attacks.

Tactical

- Identify organisations that are producing incident response reports and white papers on threat groups. Set up RSS Feed alerts for new papers released by these organisations.
- When a paper is released, extract from it the key tactical indicators, such as the initial mechanism of entry to the network, tools or techniques used to move around the network, and mechanisms used for exfiltration. Carry out a paper exercise to determine how susceptible your organisation would be to those techniques, and what changes are needed to reduce that susceptibility.
- Consult architects and systems administrators to identify planned refreshes of technologies, environments or key systems. Identify opportunities to feed tactical intelligence into these refreshes to mitigate attacks at the design and implementation phase.

Technical

- Obtain access to the daily C2 list from CiSP or other free feeds, and place the IP addresses in an 'alert' list on the primary firewall or IDS. Review regularly to determine whether outbound connections are being made from within your organisation and – if so – initiate incident response.

Sharing

- Identify a forum in which you already participate, or in which you can readily do so, and discuss threat intelligence with the members of that forum. For example, what they are currently doing in their organisation, and what would they like to be doing?
- Identify appropriate peers in similar organisations, preferably where there is already a relationship and reasonable level of trust. Arrange to meet to discuss your joint perception of existing threats, with the aim of developing the trust to your mutual benefit.

Functions of a Threat Intelligence Team

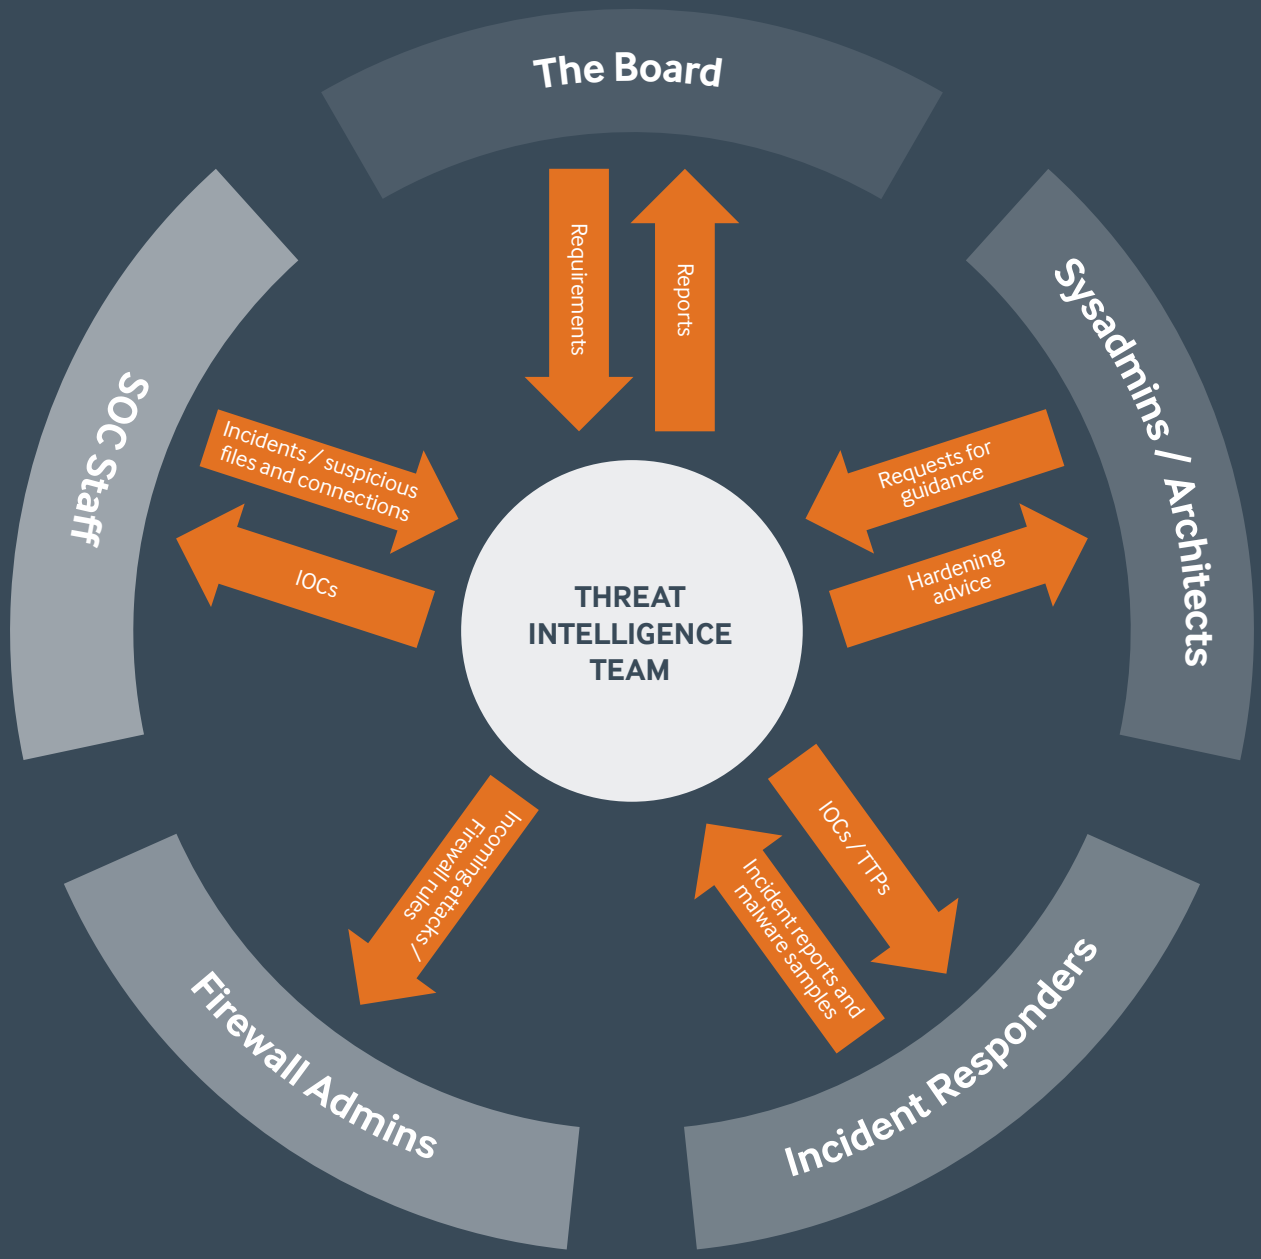

Glossary

API	Application Programming Interface – An interface for programs or scripts to interact with automatically (i.e. without a human directly involved), used for exchanging information between remote programs
AV	Antivirus
Bro IDS	An open source, highly flexible, network-based IDS – https://www.bro.org
C2	Command and Control – The mechanism used by malware to communicate with those behind it
CIR	A UK Government-run scheme for companies approved to conduct forensic investigations into attacks on government computers
CiSP	Cyber Information Sharing Partnership – https://www.cert.gov.uk/cisp/
CPNI	Centre for the Protection of National Infrastructure
DDoS	Distributed Denial of Service – An attack to render a service inoperable, conducted from large numbers of attacking hosts
DoD	United States Department of Defense
IDS	Intrusion Detection System – Software working at either computer or network level to detect signs of compromise. Typically compares activity to a list of known 'bad' activities
IOC	Indicator of Compromise – Typically a technical artifact of malware or malware communications that can indicate a compromise
MD5 hash	A 128 bit representation of an input, where the same input always produces the same output, but output (the hash) cannot be reversed to discover input
Mimikatz	An open source tool favoured by many attackers that, among other things, can allow extraction of passwords from Windows systems – https://github.com/gentilkiwi/mimikatz
NDA	Non-disclosure agreement
PsExec	A tool from Microsoft that allows running of commands on remote machines. Used legitimately by systems administrators, but also by a number of attackers
RAT	Remote Access Tool – Malware to allow remote control of a computer
RSS Feed	RSS (Rich Site Summary) is a protocol for organising content so that new content can be detected programmatically and delivered via a feed
SHA-1 hash	Similar in concept to MD5 hash but 160 bit and considered a better algorithm
SIEM	Security Incident and Event Management – Software to allow correlation and investigation of alerts
Snort	An open source network-level IDS – https://www.snort.org
YARA	A pattern-matching tool for writing and matching malware signatures – https://plusvic.github.io/yara/

Further Reading

An Introduction To Threat Intelligence (CERT-UK 2014)
Overview of threat intelligence and different sharing formats
http://www.cert.gov.uk/resources/best-practices/an-introduction-to-threat-intelligence

10 Steps to Cyber Security (GCHQ 2015)
A resource for business to help address the 10 most important areas with regard to cyber security
https://www.gov.uk/government/publications/cyber-risk-management-a-board-level-responsibility

Guide to Cyber Threat Information Sharing (NIST)
A detailed overview of the challenges and some solutions relating to sharing threat intelligence
http://csrc.nist.gov/publications/drafts/800-150/sp800_150_draft.pdf

Effective Threat Intelligence (ContextIS 2015)
A partner work to this piece, covering protecting networks through threat intelligence
http://mwr.to/ctxti

The Threat Intelligence Cycle (Krypt3ia, 2014)
Blog post covering the threat intelligence cycle and an overview of the subject
https://krypt3ia.wordpress.com/2014/10/02/the-threat-intelligence-cycle/

OSINT (Rohit Shaw)
Introduction to and overview of threat intelligence
http://resources.infosecinstitute.com/osint-open-source-intelligence/

The Chinese People's Liberation Army Signals Intelligence and Cyber Reconnaissance Infrastructure (Stokes, Lin, Hsiao, Project 2049, 2011)
An excellent example of how much insight can be gained via open source intelligence from foreign language sources
http://project2049.net/documents/pla_third_department_sigint_cyber_stokes_lin_hsiao.pdf

Maturity Model

Maturity models can be useful tools that help an organisation to define what success looks like – and then to break it down into manageable stages. By designing a model that supports and codifies the organisation's direction, those involved in its implementation can gain clear guidance as to the specific steps to take. Maturity models also provide feedback to teams implementing change, as they can see their level of maturity in each area progressing throughout the project.

Maturity models are best designed by those within the organisation, clearly mapping the successive steps needed to reach the desired result. Given here are examples of maturity models for the four types of threat intelligence.

STRATEGIC

AREA	1	2	3	4	5
Requirements	Board and senior managers unaware of threat intelligence and the team responsible for it	Board and senior managers aware, occasional intelligence offered by team and sometimes considered. Rarely, if ever, acted upon	Threat intelligence pushed by team on big issues; board receives and considers information	Threat intelligence utilised by management in decision-making, particularly decisions that are clearly cyber-related	Threat intelligence a routine part of decision-making, with advice sought on all major decisions
Awareness of threats	No awareness of threats to organisation	Some awareness of threats, focusing on one or several of the more commonly discussed actors. No tracking of motivations or development of threats	Some awareness of threats, focusing on one or several of the more commonly discussed actors. Attempts to track trends of threat motivation, economic situation and technological developments	Deeper insight into trends of traditional cyber threats; consideration of – and some understanding of – less commonly discussed threats	Awareness of many threats, including those that aren't 'traditional' cyber threats. Robust understanding of the motivational, economic, and technological developments of those threats
Collection	None	Small number of sources consumed. A focus on 'overview' style articles or reading other people's analysis on the same topic	A focus on reputable, well-known sources of information in key areas, such as popular journals, press. Occasional use of articles reporting on situations or events, rather than overviews and reviews	A number of reputable sources in a variety of areas. Includes some articles from uncommon sources, such as foreign language press or lesser-known journals. Products of other threat intelligence types occasionally used	Large range of sources, including economic, socio-political, foreign language journals, press articles, and products of other threat intelligence types. Focus primarily on reports relating to situations, events or statements
Analysis	No analysis; any sources consumed are reported directly	Some analysis of sources and verification of content of overview articles. Some attempt made to map to general businesses	Analysis leading to insight that supports publically available reviews and commentary. Broad mapping to general businesses or organisations with occasional specificity to the organisation itself	Robust analysis of sources, leading to insight, particularly where a threat isn't discussed in the open press: i.e. surpassing publically available reviews. Intelligence is gained by mapping insight to understanding of the business	Deep analysis, leading to insight that surpasses that of review articles on key topics. Mapped to business in a way that takes into account financial drivers, structure and intentions of the organisation
Production / Evaluation	Threat intelligence not involved in strategic decisions	Threat intelligence considered but generally disregarded	Threat intelligence considered and occasionally used to support greater protections or defensive measures	Threat intelligence generally used in the implementation of decisions, such as increased security budget to mitigate a risk. Intelligence rarely changes decision or the nature of a decision	Threat intelligence occasionally changes decisions and regularly affects how those decisions are implemented
Sharing	No sharing	Individuals at similar organisations identified with whom it might be possible to establish a sharing relationship	Semi-regular meetings with individuals and some non-sensitive sharing, e.g. opinions on publically available information	Relationship developed and trust built, to a stage where it is possible to trust the information received from the sharer	Trusted relationships built and maintained with peers at other organisations; regular bidirectional sharing of information that helps both parties to understand strategic risk

OPERATIONAL

AREA	1	2	3	4	5
Requirements	No tasking to identify actvity-related attacks or groups who plan attacks openly	Broad tasking to identify whether attacks are occurring as a result of activities, or by a group whose communication channels are legally accessible	Specific tasking to investigate a group or activity-related attack	Tasking includes specific groups or events to target, and identifies desired outcomes: e.g. how long it should take to trigger a response	Requirements evolve with evaluation; efforts made to develop capabilities where there is indication of a return on investment
Activity-related attacks: note, not all organisations experience these attacks	No attempts to identify activity-related attacks	Attempts made to find an activity or event correlated to attack types	Activity-related attacks sporadically predicted	Activity-related attacks regularly predicted, but no coordinated response	Activities that result in attacks robustly understood, and appropriate monitoring in place. Response planned; success of attack evaluated afterwards
Attacks by groups communicating openly: note, not all organisations experience such attacks	No attempts to monitor groups	Attempts made to identify groups that attack and communicate using open channels	Open communication channels monitored manually by staff for indications of attacks being planned	Staff tracking agitators and attempting to consolidate changing aliases and styles	Open communication chan-nels monitored; changing aliases tracked, with scripted analyses to determine when attacks may occur
Communication with staff	No plans	Key staff identified, who will be alerted to an incoming incident	List of out-of-hours contact information for key staff is maintained and available when necessary	A rehearsal has been conducted for an incoming incident, involving all relevant staff	Regular exercises undertaken (including out of hours) for different incident types
Evaluation	No evaluation	Report prepared, identifying how many alerts were produced by operational threat intelligence and whether they were plausible	Report prepared, containing details of how robustly staff responded to alerts, and how their actions could be improved	Formal process defined for evaluating the success and failure of individual cases	Efforts robustly evaluated, with undetected attacks (where detection should have been possible) subject to root cause analysis
Sharing	No sharing	Organisations that might be attacked by same groups identified	Specific individuals at other organisations, who would be involved in responding to an attack, identified	Other organisations have the necessary details of your team members who should be alerted in the event of an attack	Other organisations have been successfully alerted, allowing them to better protect themselves as a result

TACTICAL

AREA	1	2	3	4	5
Awareness of Remote Access Tools (RATs)	No awareness	Awareness of the most common RATs used, such as name and capabilities	Samples of common RATs maintained; third-party reports used for information on RATs and matched to samples	In-house analysis of RATs used to complement and expand third-party knowledge; focus given to RATs used by groups that are expected to target the organisation	Deep understanding of the RATs used by expected threat actors that includes understanding of capabilities and robust indicators (not hashes). Basic data available on RATs of groups not believed to target industry
Awareness and vulnerability with respect to C2 channels	No specific awareness beyond the fact that malware has to communicate with attackers	Intelligence team aware of common methods such as HTTP, HTTPS, DNS	Understanding of different channels and how various types of malware use them, e.g. that the 'Havex' malware uses HTTP with tags containing the phrase 'havex'	Understanding of the likely success of attackers who use the various types of communication channel within the organisation. Triaged list of controls/monitoring implemented to prevent, detect or impede communication channels	Understanding of specific channels used by pieces of malware. Susceptibility of organisation to malware communicating via those channels fully assessed. Controls or monitoring in place wherever feasible, and remaining channels under investigation or considered an accepted risk
Knowledge of modus operandi (MO)	Basic understanding of attack flow	Understanding of the fact that different attack groups favour different methodologies. Specific examples of methodologies can be provided	Knowledgebase maintained of how a variety of campaigns that have targeted the organisation's industry functioned at each stage of attack (e.g. how privileges were escalated, etc.)	Detailed knowledgebase maintained, including cross references between reports, possibly including internal reports. Root causes for key stages analysed, e.g. were attackers exploiting an issue, using a tool legitimately that did not exceed normal behaviour, etc.	Expert-level knowledge maintained on all key campaigns / attack groups. This includes breakdown of tools used (attack-supplied and those built into the OS or available from third parties), how key stages of the attack are executed, with results mapped onto specific issues that are exploited by attackers (such as re-used passwords)
Application of knowledge of MO to organisation	No specific attempts to map attacker MO to organisational weaknesses	Key issues exploited by attackers triaged by likelihood and impact	Current state of controls in the organisation that mitigate key issues assessed	Timeline prepared for mitigating most significant or likely issues exploited by attackers. Monitoring or logging put in place for exploitation of key issues that cannot be immediately addressed	Majority of issues exploited by attack groups targeting the organisation are subject to mitigation controls; otherwise, a remediation plan or monitoring/alerting is in place
Sharing	No sharing	Commentary available from organisation's experts on attacker tooling and methodology	Malware samples or studies released by organisation	Attacker methodology studies released by organisation	Sharing of deep understanding of attacker tooling and methodology. Organisation produces public reports that add to overall understanding of attackers by tying together malware capability analysis and attack methodology

TECHNICAL

AREA	1	2	3	4	5
Requirements for indicators	No specific requirements for technical threat intelligence	Requirements are broad, such as 'consume all publically available feeds'	Requirements are specific, e.g. 'identify communication to known C2 channels from our device, using automated collection and analysis from feeds believed to contain high-value indicators'	Requirements develop with the programme, and have both immediate and longer-term goals	Results of evaluation are an active part of requirement setting and management of the process
Collection of indicators	No collection	Ad-hoc collection, e.g. from occasional reports or collaboration partners	Collection from public feeds	Collection from public feeds, and private feeds such as sharing relationships	Collection from feeds supplemented by managed extraction of indicators from white papers and reports, as close to the time of release as possible
Analysis of indicators	No analysis, or indicators manually actioned	Indicators stored in a central repository that might be ad-hoc (e.g. an Excel spreadsheet)	Indicators stored in a flexible repository that allows filtering by metadata	Indicators curated and higher-value indicators prioritised; output machine-generated for consumption by detection or investigation tools	Analysis engine for indicators handles collection, analysis, linking between different types and sources based on metadata and application of indicators to data
Application of indicators	No application of indicators to organisation	Indicators are manually actioned by a staff member, e.g. by logging onto hosts to check for registry paths or looking at firewall logs	Network-based indicators are automatically investigated by network devices	Automatic searching for host-based indictors across the whole estate, probably utilising third-party software	Indicators of all types automatically searched for in network traffic and on hosts; new indicators that become available are used to search through log data for historical signs of compromise
Evaluation of indicators	No evaluation	Monthly report prepared of how many alerts were a result of indicators from specific sources	Monthly report contains analysis of whether alerts were false positives, plus commentary on validity of the indicator producing the alert	Monthly report identifies whether verified alerts were generated as a result of an indicator that was also detected by other mechanisms (such as antivirus), or whether this would have been possible	Monthly report identifies whether verified alerts were generated as a result of an indicator that was also detected by other mechanisms (such as antivirus), or whether this would have been possible. Incidents that emerge are analysed to identify whether technical threat intelligence should have allowed detection sooner
Sharing of indicators	No sharing	Informal sharing with a limited audience, e.g. emailing a peer at another organisation	Manual sharing with a closed group, e.g. copying into an email list	Automated sharing of verified indicators with a closed group	Automated sharing with a public group of verified indicators that have been investigated – and found not to expose specific attacks against the organisation

References

[1] **The Butler Review of Intelligence on Weapons of Mass Destruction, 2004**
http://news.bbc.co.uk/nol/shared/bsp/hi/pdfs/14_07_04_butler.pdf

[2] **BlackHat US, Threat Intelligence Library – A New Revolutionary Technology to Enhance the SOC Battle Rhythm, Ryan Trost 2014**
http://mwr.to/rtrost

[3] **The Blue Pill of Threat Intelligence, Dave Aitel 2014**
https://lists.immunityinc.com/pipermail/dailydave/2014-October/000769.html

[4] **Protecting Privileged Domain Accounts, Mike Pilkington 2012**
http://digital-forensics.sans.org/blog/2012/03/09/protecting-privileged-domain-accounts-disabling-encrypted-passwords
See also Sysmon, Microsoft 2015
https://technet.microsoft.com/en-gb/sysinternals/dn798348

[5] **An Overview of the Intelligence Community: An Appraisal of U.S. Intelligence 1996, Commission on the Roles and Capabilities of the United States Intelligence Community**
http://www.fas.org/irp/offdocs/int023.html

[6] **Structured Analytic Techniques for Improving Intelligence Analysis, US Government 2009**
https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/Tradecraft%20Primer-apr09.pdf

[7] **CiSP – Cyber-Security Information Sharing Partnership**
https://www.cert.gov.uk/cisp/

[8] **Guide to Cyber Threat Information Sharing, NIST 2104**
http://csrc.nist.gov/publications/drafts/800-150/sp800_150_draft.pdf

[9] **CIR – Cyber Incident Response scheme**
https://www.cpni.gov.uk/advice/cyber/cir/

[10] **Heartbleed Flaw Said Used by Chinese in Hospital Hacking, Bloomberg 2014**
http://www.bloomberg.com/news/articles/2014-08-20/heartbleed-flaw-said-used-by-chinese-in-hospital-hacking

[11] **Open Source Intelligence: What Is It? Why Is It Important to the Military? OSS 1997**
http://www.oss.net/dynamaster/file_archive/040320/fb893cded51d5ff6145f06c39a3d5094/OSS1997-02-33.pdf

[12] **The Future of Open Source Intelligence, Corey Velgersdyk for the Elliott School of International Affairs 2010**
http://iar-gwu.org/node/253

[13] **Sailing the Sea of OSINT in the Information Age, Stephen C. Mercado**
https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/csi-studies/studies/vol48no3/article05.html

[14] **Castles in the Sky – a blog on malware analysis**
http://jcsocal.blogspot.co.uk/2012/12/malware-analysis-1-protip.html

[15] **Structured Analytic Techniques for Improving Intelligence Analysis, US Government 2009**
https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monograph

[16] **Attributing Cyber Attacks, Thomas Rid & Ben Buchanan, Journal of Strategic Studies**
Volume 38, Issue 1-2, 2015
http://www.tandfonline.com/doi/full/10.1080/01402390.2014.977382#.VOtYZd4glhM

[17] **Targeting U.S. Technologies, DSS 2014**
http://www.dss.mil/documents/ci/2014UnclassTrends.PDF

[18] **The Streaming APIs, Twitter**
https://dev.twitter.com/streaming/overview

[19] **Security Intelligence: Attacking the Cyber Kill Chain, Mike Cloppert 2009**
http://digital-forensics.sans.org/blog/2009/10/14/security-intelligence-attacking-the-kill-chain

[20] **CrowdStrike Global Threat Report 2014**
http://www.crowdstrike.com/2014-global-threat-report/ (requires sign up)

[21] **Exploit This, Gabor Szappanos for Sophos 2015**
https://www.sophos.com/en-us/medialibrary/PDFs/technical%20papers/sophos-exploit-this-evaluating-exploit-skills-of-malware-groups.pdf

[22] **CosmicDuke, F-Secure 2014**
https://www.f-secure.com/documents/996508/1030745/cosmicduke_whitepaper.pdf

[23] **The Pyramid of Pain, David Bianco 2014**
http://detect-respond.blogspot.co.uk/2013/03/the-pyramid-of-pain.html

[24] **A String of Paerls, Williams, Schultz, Esler, and Harman, Cisco 2014**
http://blogs.cisco.com/security/a-string-of-paerls

[25] **Tools and Standards for Cyber Threat Intelligence Projects, Greg Farnham for SANS 2013**
https://www.sans.org/reading-room/whitepapers/warfare/tools-standards-cyber-threat-intelligence-projects-34375

[26] **toolsmith – Threats & Indicators: A Security Intelligence Lifecycle, Russ McRee 2014**
http://holisticinfosec.blogspot.co.uk/2014/08/toolsmith-threats-indicators-security.html

[27] **Threat Library: A SOC Revolution, Ryan Trost for ThreatQuotient 2014**
http://www.isaca.org/Education/Conferences/Documents/NAISRM/145.pdf

[28] **The Collective Intelligence Framework (CIF)**
https://code.google.com/p/collective-intelligence-framework/

[29] **Federated Threat Data Sharing with the Collective Intelligence Framework, An introduction to CIF, Gabriel Iovino et al. 2013**
http://www.internet2.edu/presentations/tip2013/20130116-iovino-security-event-sharing.pdf

[30] **The Bro Network Security Monitor**
https://www.bro.org/

[31] **Guide to Cyber Threat Information Sharing, NIST 2104**
http://csrc.nist.gov/publications/drafts/800-150/sp800_150_draft.pdf

MWR InfoSecurity (Head Office)
Matrix House, Basing View
Basingstoke RG21 4DZ
T: +44 (0)1256 300920
F: +44 (0)1256 323575

MWR InfoSecurity
77 Weston Street
London SE1 3RS

MWR InfoSecurity
113-115 Portland Street
Manchester M1 6DW

MWR InfoSecurity (South Africa)
Homestead Place
Cnr 12th Ave & Homestead Lane
Rivonia, Gauteng 2128. South Africa
T: +27 (0)10 100 3157
F: +27 (0)10 100 3160

MWR InfoSecurity (Singapore)
62A Pagoda Street
#02-00
Singapore 059221
T: +65 6221 0725

www.mwrinfosecurity.com
labs.mwrinfosecurity.com

Follow us on Twitter:
@mwrinfosecurity
@mwrlabs