BANK OF ENGLAND

CBEST

CBEST Intelligence-Led Testing

Understanding Cyber Threat Intelligence Operations

Version 2.0

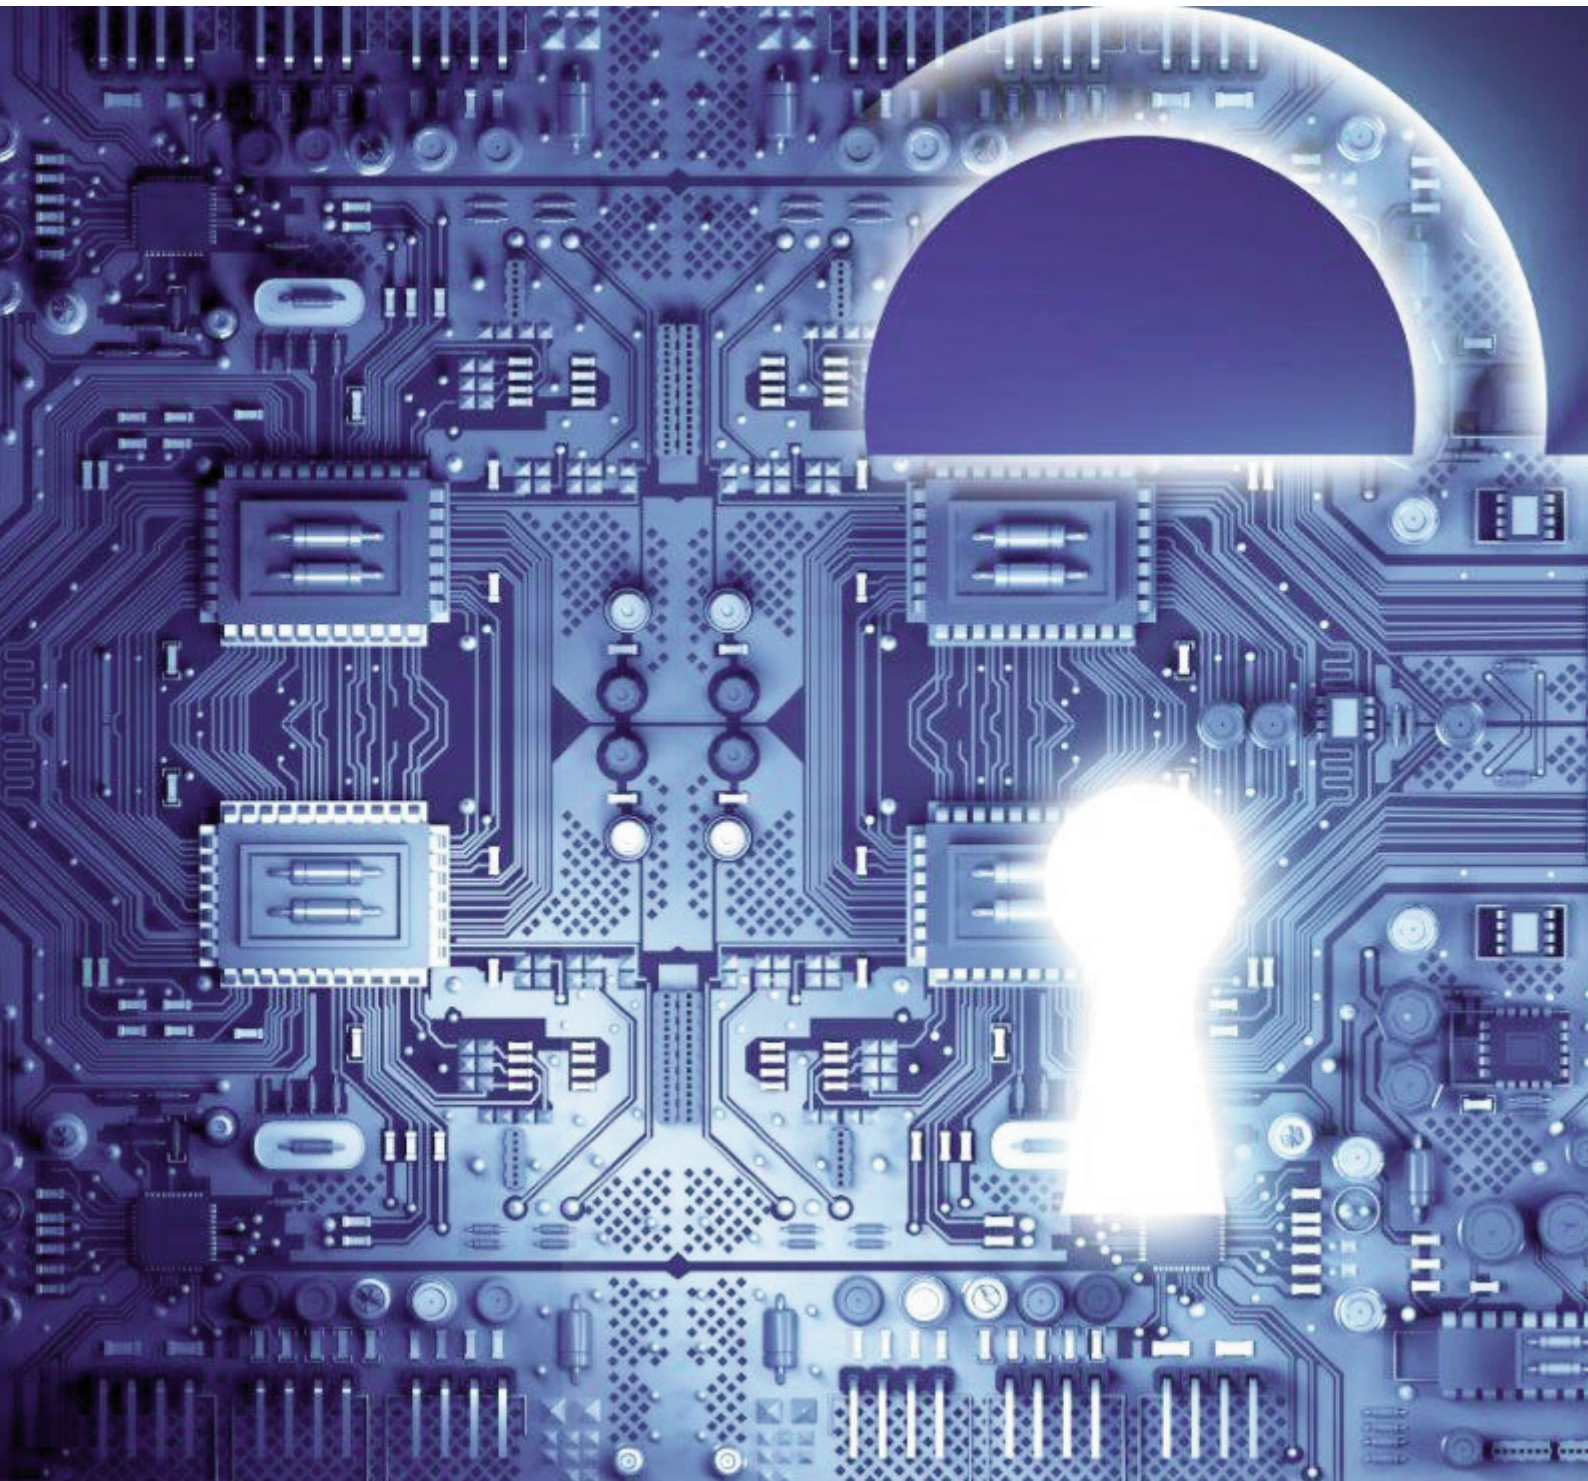

Copyright notice

Contents

Executive summary

This document defines best practice standards for the production and consumption of threat intelligence. It is intended to provide the CBEST programme with a foundation for defining and executing intelligence-led cyber threat vulnerability tests in conjunction with accredited providers of threat intelligence products and services.

After establishing some important terminology, this document presents an overview of the process underpinning a best practice threat intelligence capability and the organisation, roles and skills required for running it. It then discusses maturity models relating to the production and consumption of threat intelligence.

The key conclusions of this document are:

- in comparison to its counterpart in the government and law enforcement sector, cyber threat intelligence in the commercial environment remains a relatively immature discipline and is also the subject of much vendor hype;
- that said, in the financial sector there is a high level of information sharing (eg FS-ISAC and CISP) although immaturity remains at the strategic level;
- commercial intelligence is in a good position to exploit several decades of government intelligence best practice and it is a positive sign that intelligence providers have broadly settled on a common intelligence life cycle model;
- intelligence sharing remains at a relatively low level of maturity, much of it taking place at the tactical or operational level rather than at the strategic level;
- defining the requisite roles and skills in a threat intelligence team also remains at a relatively low level of maturity, in particular appreciating how cyber threat intelligence analysis differs from traditional information security in the way it requires skills in the humanities rather than a sole focus on computer science;
- while maturity models for threat intelligence producers are relatively simple and well-defined, consumer models are more complex and remain at an earlier stage of development;
- there are clear benefits to be gained from implementing intelligence-led cyber resilience, not just in terms of proactively managing an array of new and evolving advanced cyber threats but also the potential for improving risk management and high-level business strategy;
- as for traditional information security, adopting an intelligence-led approach has the potential to transform it from a largely reactive function that investigates, remedies, complies and reports to a more responsive function that also diagnoses, predicts, executes and influences;
- as for the future, providers and users of threat intelligence services will see increased benefits from a more robust, holistic and tailored approach to generating threat intelligence which will play a far more strategic role and help organisations develop a stronger sense of situational awareness.

Threat intelligence is a moving target and this report will only ever be a snapshot of the current state of the art. As CBEST continues to evolve the issues highlighted above, relating to the concept of operations for cyber threat intelligence, should therefore be explored in further detail.

1 Introduction

1.1 Purpose of this document

This document defines best practice standards for the production and consumption of threat intelligence. It is intended to provide the CBEST programme with a foundation for defining and executing intelligence-led cyber threat vulnerability tests in conjunction with accredited providers of threat intelligence products and services.

1.2 Terms of reference

1.2.1 Smarter adversaries

Organisations, together with their staff, customers and supply chain partners, are facing an increase in targeted cyber attacks committed by adversaries ranging from hackers and hacktivists to criminals and nation states. The goal of these attackers is to steal, compromise or destroy organisational assets that have financial, operational, intellectual, confidential or reputational value. Underpinned by intensive preliminary research on their targets, which then forms the basis for carefully crafted and targeted phishing attacks, the threat actors involved are commonly referred to as '*advanced persistent threats*' (APTs) (Daly (2009).

So professional, targeted and sophisticated are their methods, and so dynamic and fast-changing, that cyber attackers are now bypassing traditional perimeter defences that typically only react once a known threat has been detected by its signature. The question is therefore not if cyber attackers will penetrate perimeter defences but when. In many cases the enemy is already inside the wire. A delayed reaction while cyber attackers have already begun stealing, compromising or destroying assets is simply not acceptable.

1.2.2 The shift to intelligence

All systemically important organisations therefore need to raise their security game to defend themselves against 21st century cyber attackers; they cannot risk-assess their way out. In military terms this is an asymmetric war and, currently, the best result that attack targets can achieve is a draw. Speed of response and a better understanding of who is behind the attack will separate the winners from the losers.

Organisations therefore need to train harder than they will fight otherwise their adversaries will win. A key element of this training is the adoption of an intelligence-driven approach taken from traditional warfare. This has two goals:

- to prevent an attacker from successfully attacking;
- to be able to recognise and respond effectively to an attack that has already happened.

Information security practitioners already undertake a degree of intelligence work albeit after the attack has taken place. Many are now trying to improve their detection capability by identifying and sharing so-called '*indicators of compromise*', or forensic remnants of an intrusion residing in operating system and network devices.

It is now a question of becoming more proactive by moving beyond the technical details of the attack (the *what*, *when* and *where*) towards a better understanding, and attribution, of the tactics, techniques and procedures (TTPs) behind the attack (the modus operandi or *how*) and, critically, the attackers themselves (the *who* and *why*). Such intelligence places cyber threats in context and, through greater situational awareness, better informs the countermeasures. In this way, through better understanding, information security can move from reactive, '*seize and wipe*' defence to responsive, proactive, intelligence-led cyber resilience.

1.2.3 CBEST

Information security testing regimes stand to benefit from taking on intelligence-led techniques to make their tests more focused and proactive. While traditional security testing has been more than adequate for the vast majority of target environments, it does not adequately cover the new breed of professional, sophisticated and industrialised threat actors and may not therefore be suitable for systemically important organisations. To date, organisations have been loath to test their critical systems against an attack because of associated risks. Furthermore, the security testing industry has not had sufficient access to high quality threat intelligence.

In light of this, and the general recommendations of the Financial Policy Committee (FPC) on improving resilience against cyber threats, the Bank of England is in the process of implementing CBEST. This is a framework for developing intelligence-led cyber threat vulnerability tests against financial institutions' critical systems. These tests mimic the actions of groups and individuals who are perceived by Government and commercial threat intelligence providers as posing a genuine threat to systemically important financial institutions within the Critical National Infrastructure.

CBEST is supported by the Cabinet Office and in turn supports the objectives of the UK Cyber Strategy objectives (Cabinet Office (2011)), in particular:

- being more resilient to cyber attack;
- enhancing the United Kingdom's cyber security knowledge.

The goal, through intelligence-led supervision and CBEST-accredited threat intelligence services, is to test and improve banks' resilience against the highest level of cyber threat. CBEST therefore takes banks beyond the '*10 Steps*' model created by CESG, BIS and the Cabinet Office (CESG (2012)). As well as testing current resilience, CBEST reconnaissance exercises (ie the enumeration of a target's technical and organisational infrastructure) provides banks with a valuable '*attacker's eye view*' into their organisations.

1.2.4 Gauging best practice

To be effective, CBEST simulated cyber attacks must be based on realistic, threat-informed scenarios. The Bank of England is therefore forming partnerships with commercial providers of threat intelligence and security testing services to help establish a best practice approach to defining and executing the tests. Essentially the threat intelligence service providers pass threat intelligence, augmented by Government sources, on to security testers who then use it to target their attacks.

However, the increased interest in threat intelligence across industry and the media has created a significant amount of hype in the market with all kinds of vendors claiming to provide threat intelligence services. Some are newly formed specialist vendors while others have simply rebadged their existing services by replacing '*vulnerability*' and '*analysis*' with '*threat*' and '*intelligence*' (451 Research (2014)). One industry analyst has commented: '*The security vendor community has hijacked the term 'intelligence"* (Holland (2013)).

The extent to which vendors provide the kind of intelligence needed to underpin CBEST varies considerably, from elementary fact-finding about commodity threats through to information that provides sufficient understanding for mitigating a harmful event. Quality varies considerably between providers and a body of confusing and inconsistent jargon has arisen which only serves to muddy the picture further.

Therefore, at this stage of CBEST development, the Bank of England requires the answers to three key questions:

- what is best practice in the production and consumption of cyber threat intelligence?
- what are the characteristics of best practice analytical model of threat intelligence?
- what baseline criteria should be used to accredit CBEST providers?

This report provides an answer to the first question above. Its scope is intelligence relating to criminally motivated external cyber attacks. Insider threats, whether criminally motivated or accidental, are out of scope and so too are physical threats.

Two sister CBEST reports, *An Introduction to Cyber Threat Modelling* and *CBEST Services Assessment Guide*, answer the other two questions above (CBEST (2016a); CBEST (2016b)). Together the three reports are intended to provide a foundation for defining and executing CBEST tests in conjunction with accredited service providers. They will also help the greater community understand and use cyber threat intelligence. **Figure 1.1** summarises the scope and context of the three reports.

Figure 1.1 Scope and context of the three reports

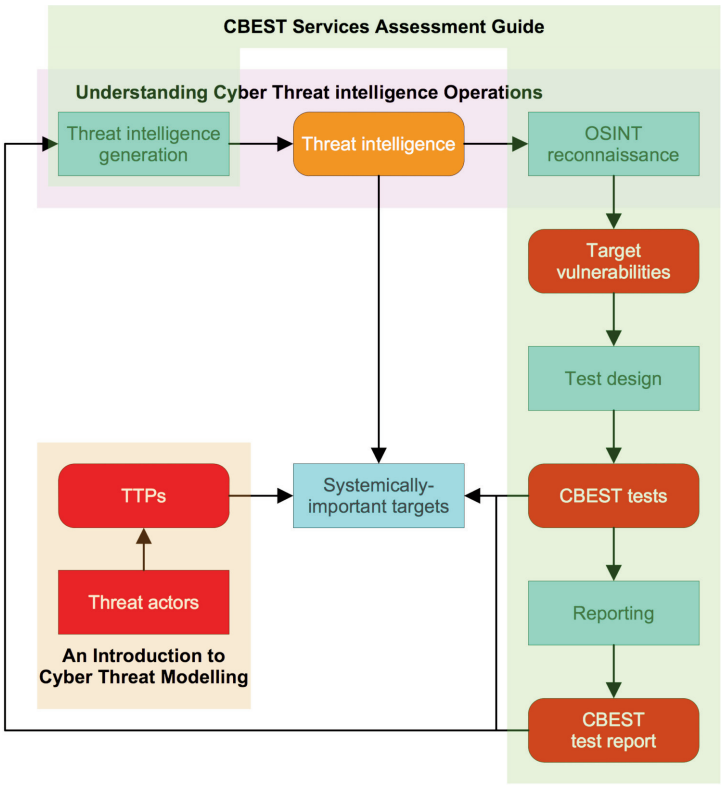

1.3 Structure of this document

The remainder of this document is structured as follows:

- Section 2, *Terminology*, establishes some important terminology as a pre-cursor to more detailed discussions on threat intelligence;
- Section 3, *Process*, presents an overview of the process underpinning a best practice threat intelligence capability;
- Section 4, *Organisation*, presents an overview of the organisation, roles and skills required for running a best practice threat intelligence capability;
- Section 5, *Maturity*, discusses maturity models relating to the production and consumption of threat intelligence;
- Section 6, *Conclusions*, presents a summary of the key points and discusses future developments in cyber threat intelligence;
- Section 7, *References*, lists sources of information used in the production of this report.

1.4 Information sources

Information for this report was gathered from online open sources and discussions with industry professionals. A full set of references appears at the end of this document. Information was also derived from various CBEST meetings and workshops attended by the representatives of the Bank of England, CREST and the Cyber Working Group during the first quarter of 2014. In 2015 the Bank of England Cyber Sector Team commissioned a review and update of this document during which various stakeholders were canvassed for their input.

1.5 Legal disclaimer

The information and opinions expressed in this document are for information purposes only. They are not intended to constitute legal or other professional advice, and should not be relied on or treated as a substitute for specific advice relevant to particular circumstances. The sponsors and authors of this document shall accept no responsibility for any errors, omissions or misleading statements in this document, or for any loss that may arise from reliance on the information and opinions expressed within it.

2 Terminology

2.1 Introduction

This section establishes some important terminology as a pre-cursor to more detailed discussions on threat intelligence.

2.2 Threat

2.2.1 Fundamentals

The dictionary definition of a threat is:

- an expression of intent to injure or punish another;
- an indication of imminent danger;
- a person or object that is regarded as a danger; a menace.

At the Bank of England Cyber Working Group meeting held on 9 January 2014 the following draft definition of a threat was drawn up:

- actions undertaken by an agent with the intention to harm, undermine, weaken, or deceive a target;
- the agent itself or its tactics, techniques, and procedures (TTPs).

Combining and refining these we arrive at the following definition of a threat:

Threat

- an expression of intent to do harm, ie deprive, weaken, damage or destroy;
- an indication of imminent harm;
- an agent that is regarded as harmful;
- a harmful agent's actions comprising of tactics, techniques, and procedures (TTPs).

It is important to note that a threat is someone or something that exploits a vulnerability in a target. The vulnerability, such as a software bug or a weak password, is not the threat itself; the threat is whoever who takes advantage of that vulnerability (451 Research (2014)).

2.2.2 Cyberspace

CBEST deals with cyber threats, or threats related to cyberspace where cyberspace can be a tool or a target for an adversary.

What is cyberspace? For the purposes of CBEST, the default definition of cyberspace comes from the UK Cyber Security Strategy:

Cyberspace

An interactive domain made up of digital networks that is used to store, modify and communicate information. It includes the Internet and also the other information systems that support our business, infrastructure and services (Cabinet Office (2011)).

This is broadly similar to other definitions of cyberspace. For example, the US Army Field Manual that covers '*cyber electromagnetic activities*' describes cyberspace as '*…the interdependent network of information technology infrastructures and resident data, including the Internet, telecommunications networks, computer systems, and embedded processors and controllers*' (US Army (2014)).

To properly understand cyberspace it is necessary to put it in context. From a military perspective, cyberspace is relatively recent addition to the four traditional operational domains of air, land, maritime and space. An additional domain that cuts across all of these is the electromagnetic spectrum (EMS). Of these six domains, air, land, maritime, space and the EMS exist naturally. Cyberspace, by contrast, is entirely man-made. Because cyberspace is man-made, it is only through continued attention and maintenance that it persists.

Cyberspace takes the form of a global network of computers located mainly on land but also across air, sea and space. The telecommunications network (the Internet) that binds cyberspace together utilizes the EMS. Cyberspace is a system of systems in that many small and diverse systems comprise the structure as a whole. This is summarised in **Figure 2.1**.

Figure 2.1 Cyberspace in context

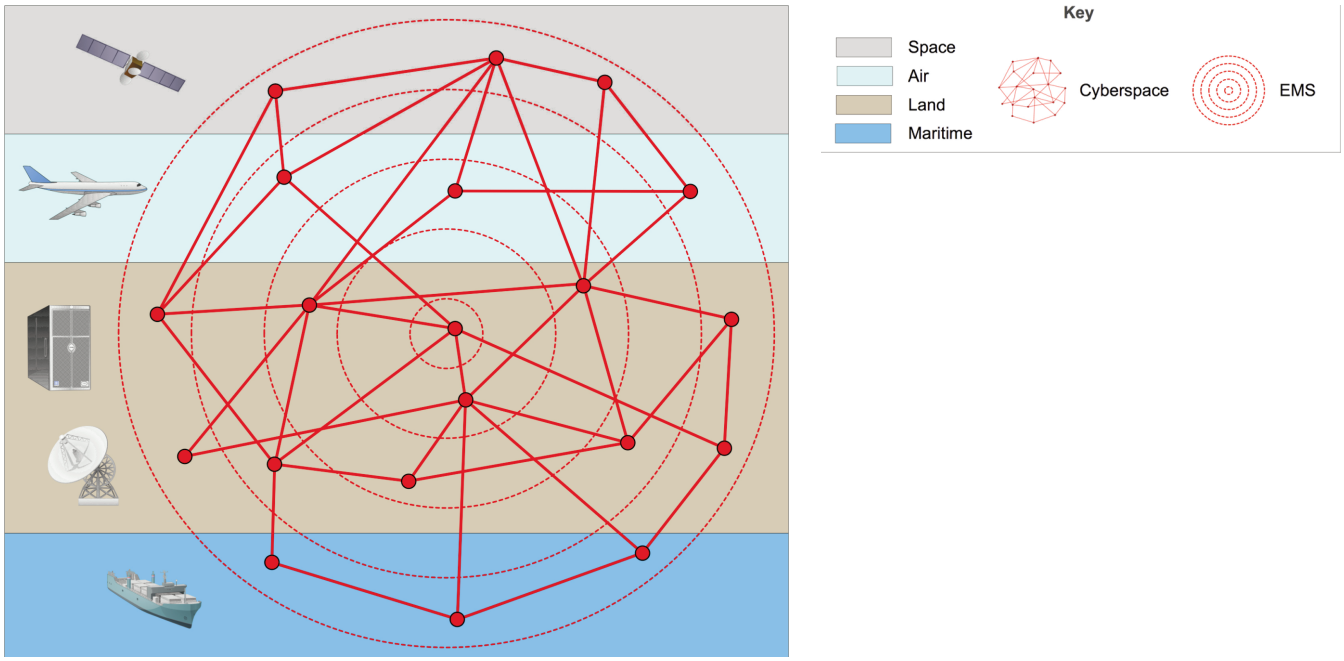

Because of the ubiquitous nature of cyberspace and the EMS, the six domains in **Figure 2.1** are tightly interconnected such that an event in one domain can cause an event (or a series of cascading events) in one or more of the other domains. Cyberspace and the EMS reinforce the fact that an operational framework is not confined to a physical place. Traditional battlefields are confined to physical space. While the outcome of a battlefield can create social and political effects around the world, the actual physical impact is limited to the physical battlefield. The inclusion of cyberspace and the EMS greatly expands and complicates the situation, transforming a limited physical battlefield to a global battlefield. For example, a computer virus executed in cyberspace may strike its intended target and also indiscriminately strike other systems around the world (US Army (2014)).

At a more abstract level, cyberspace is an environment created and maintained for the purpose of facilitating information exploitation, human interaction and general communications. It can be better understood, especially from an intelligence perspective, if it is viewed as comprising three dimensions (US Army (2014)):

Physical dimension

This comprises the core technical infrastructure: networked hardware and software across land, sea, air and space that exploits the EMS to enable the flow of information between producers, consumers, audiences and systems.

Informational dimension

This is the content (generally referred to as information but can also include data and knowledge) that is at rest or in transit within cyberspace, including machine-readable content, numbers, text, audio, pictures and video.

It is also here where cyber persona reside: digital representations of individuals or other entities that use cyberspace and have one or more identities that can be identified, attributed and acted upon. These identities may include email addresses, social network

names, web forum names, IP addresses and telephone numbers. For intelligence analysts cyber personas are key for attributing responsibility and targeting the source of a cyber threat.

The characteristics of cyberspace and the EMS provide the threat actors behind cyber personas with considerable measures of anonymity. Individuals, politically motivated groups and criminals can have a larger cyber persona than some nation states. Cyber personas can be complex, with elements in many virtual locations not linked to a single physical location or form. Therefore significant intelligence collection and analysis capabilities may be required to resolve them and this is why attack attribution is so difficult.

Cognitive dimension

This comprises the knowledge, values, beliefs, concepts, intentions and perceptions of individuals and groups transmitting and receiving information. These actors are the creators and users of the content that moves through the physical layer. This dimension provides the societal, cultural, religious and historical contexts that influence the perceptions of those producing the content and those consuming it. Governments, criminals, activists and hackers all think, perceive, visualise, understand and decide within this dimension.

2.2.3 Cyber threat/advanced persistent threat

As mentioned in Section 1.2.1, a new generation of smart adversaries are launching targeted cyber attacks against organisations. Such cyber threats are also termed '*advanced persistent threats*' (*APTs*) (Daly (2009)). Originally used to refer to state-sponsored groups that conducted cyber espionage against specific targets for political or commercial advantage, the term APT has now become more generic and is used to refer to a variety of threat actors and their attacks.

Like '*intelligence*', 'APT' can be subject to inconsistent and self-serving definition. Terms like '*advanced*' and '*persistent*' are relative. Many organisations consider an attack to be '*advanced*' simply because it bypassed their traditional defences. The reality is that many of these attacks are not particularly advanced; they are simply designed to bypass traditional signature-based mechanisms. Similarly, while many APTs take the form of highly persistent and protracted campaigns, others can be executed very quickly.

The use of cyberspace as a tool or target for attack is nothing new; what differentiates this new breed of cyber attack are the following (CREST (2013); Dark Reading (2013); Kapuria (2011); Ragan (2014); Schneier (2011); Schneier (2014); Techopedia (2014)):

Professional threat actors

The threat actor community, once dominated by amateur hackers, has expanded to include a broad range of professional threat actors, all of whom are strongly motivated, organised and funded. They include:

- state-sponsored organisations stealing military, government and commercial intellectual property;
- organised criminal gangs committing theft, fraud and money laundering which they perceive as low risk and high return;
- non-profit hacktivists and for-profit mercenary organisations attempting to disrupt or destroy their own or their client's perceived enemies.

Targeted attacks

Unlike conventional attacks — for example massive amounts of malware randomly infecting any suitably vulnerable computer — APTs target specific organisations or people within them. One class of targeted attack is Computer Network Exploitation (CNE) where the goal is to steal (or exfiltrate) confidential information from the target. This is effectively espionage in cyberspace or, in information security terms, compromising confidentiality. The other class of targeted attack is Computer Network Attack (CNA) where the goal is to disrupt or destroy the target's operational capability. This is effectively sabotage in cyberspace or, in information security terms, compromising integrity and availability.

Sophisticated tactics, techniques and procedures

Behind APTs are threat actors with long-term strategic goals. The tactics, techniques and procedures (TTPs) they use are more sophisticated and demonstrate a high degree of skill, patience and persistence, often taking months or even years to execute. They begin with a period of intensive reconnaissance on the people and systems within the target. Information sources include hacker forums, social networking websites and job hunting websites. Armed with this intelligence they then gain access to an

office-based or mobile endpoint computing device via exploitation, deception or force. Once inside they then pivot laterally from this foothold into other parts of the system where they can implement their attack.

2.3 Cyber kill chain

Prior to the arrival of APTs the '*defence-in-depth*' model prevailed. This is based on perimeter defences recognising pre-defined threat signatures. The model focuses on a single tactical behaviour on the part of the adversary with no other contextual information.

The cyber kill chain model, devised by Lockheed Martin in 2011, shifts the focus from trying to keep all adversaries out to assuming that an adversary will get in at some point. The model is based on military experience of real-world attacks and has been modified to reflect the characteristics of cyberspace. The model enumerates the different stages of a cyber attack, beginning with reconnaissance and ending with the action the attacker is undertaking such as exfiltration (Hutchins, Cloppert and Amin (2011)).

The kill chain consists of seven stages:

- **reconnaissance**: research, identify and select targets;
- **weaponisation**: bind the intruder code to a delivery mechanism (eg PDF, Word document or email message) that has been crafted to deceive the target into accepting it (spear phishing);
- **delivery**: transmit the weapon to the intended target;
- **exploitation**: run the intruder code on the target's machine and take ownership of it;
- **installation**: download and install more software to the target's machine that allows the intruder to maintain a presence inside the target's network;
- **command and control**: establish a command channel back through the Internet to an intruder-controlled server;
- **actions on objectives**: exfiltrate confidential data or disrupt or destroy the target's operational capability, moving laterally inside the network to compromise more machines.

In principle the adversary has to be successful at every link in the kill chain in order for the attack to be successful. Defenders only have to break one link in the chain in order to thwart the attack and force the adversary to start again. In this way the persistence of an APT is turned into a liability, decreasing the adversary's likelihood of success with each intrusion attempt.

By setting out the basic TTPs of a generic cyber attack, this model enables organisations to devise defensive courses of action that target and engage an adversary. It also highlights any gaps in defence capability and serves as a framework for measuring the effectiveness of defensive actions.

The kill chain underlies what Lockheed Martin refers to as '*intelligence-driven computer network defence' which is described as 'a risk management strategy that addresses the threat component of risk, incorporating analysis of adversaries, their capabilities, objectives, doctrine and limitations*' (Hutchins, Cloppert and Amin (2011)).

Variations on Lockheed Martin's original kill chain exist within the security industry. For example:

- reconnaissance, incursion, discovery, capture and exfiltration (Kapuria (2011));
- motivation and decision to act, determine objective, select avenue of approach, acquire capability, develop access, implement actions, assess and restrike (INSA (2013));
- intelligence gathering, initial exploitation, command and control, privilege escalation and data exfiltration (CREST (2013));
- staging of attack components, reconnaissance against target, execution of the attack and exploitation of attack's successes (Jellenc (2013)).

Kill chain models tend to depict attacks in terms of a neat, linear progression of activity. In reality this may not always be the case. Furthermore, kill chains do not take the actor's motives into account, nor the involvement of collusive insiders, nor fragmented attacks by more than one actor. This is why understanding the sponsor behind the attack, similar to real-world espionage, is so important.

Nevertheless, the kill chain approach enables defenders to develop defensive strategies around every link in the chain. In order to do that, however, requires good quality intelligence on what the adversary does during each link in the chain.

2.4 Intelligence

2.4.1 Information

Apart from the more general use of term to describe the ability to acquire knowledge and skills, intelligence is more specifically used in military, police or political environments to describe information, usually used or collected covertly, about an adversary or hostile activities. Its use to date in the business world has been largely in non-threat domains such as customer intelligence.

Intelligence is a particular kind of information. Intelligence and information are often used interchangeably as are information and data. To properly understand information (and therefore intelligence) it is necessary to put it in context and a useful model is the data information knowledge pyramid shown in **Figure 2.2**.

Figure 2.2 Information and intelligence in context

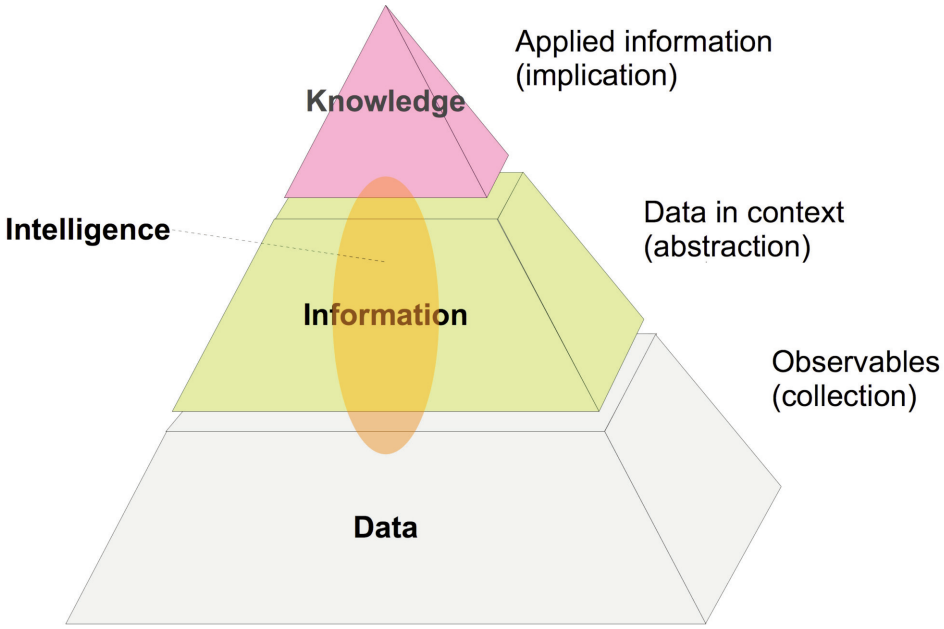

The three levels of this model are summarised below.

Data

Data equates to elementary facts and observables. For example, *name, age, postal address, telephone number, bank balance*, etc.

When describing the indicators that describe a cyber attack, the Lockheed Martin kill chain refers to elementary '*atomic indicators*' that retain their meaning in the context of an intrusion, examples being IP addresses, email addresses and vulnerability identifiers (Hutchins, Cloppert and Amin (2011)). These equate to data. On its own, data does not provide any intrinsic value.

Information

Information is data in context, or a higher-level abstraction or viewpoint made on the basis of one or more data items.
A general definition of information, drawn from classical information theory, is '*that which reduces uncertainty*' (Shannon (1948)). An example from the banking domain might be the abstraction '*account is dormant*' on the basis that the balance on a credit card account has been nil for the past nine months.

The Lockheed Martin kill chain refers to '*computed indicators*' which are derived from data involved in an intrusion, examples being hash values and regular expressions (Hutchins, Cloppert and Amin (2011)). These equate to information.

Data and information are often used interchangeably despite being different things. One potential source of confusion is that information can itself be subject to further abstraction and manipulation, in other words, one person's information can be another person's data.

Knowledge

The layer above information is knowledge, or the interpretation and exploitation of relevant information in order to solve a problem or make a decision. This is usually undertaken by humans but can also be done by machines. Very often knowledge is expressed in the form of an '*if-then rule*' (also known as a '*heuristic*', '*implication*', or, more commonly, a '*business rule*'). For example, to continue the previous banking example, a suitable heuristic might be '*If an account has been dormant, and this month's spending is very high, then it may have been taken over by a fraudster*' (where '*has been dormant*' and '*very high*' are information-level data abstractions of data).

The Lockheed Martin kill chain refers to '*behavioural indicators*' which are collections of both computed and stand-alone indicators, often subject to qualification by quantity and possibly combinatorial logic. An example might be '*The intruder would initially used a backdoor which generated network traffic matching [regular expression] at the rate of [some frequency] to [some IP address], and then replace it with one matching the MD5 hash [value] once access was established*' (Hutchins, Cloppert and Amin (2011)). This equates to knowledge.

As with information, knowledge can be subject to further abstraction and manipulation, resulting in higher-level constructs such as wisdom, intuition and so on. However, these serve to complicate the picture and are all variations on the core theme of a higher-level knowledge layer where information is structured and applied.

2.4.2 Intelligence

Intelligence is therefore a particular kind of information. Formal definitions of intelligence vary. Some might say that intelligence can be very simply defined as anything that is classified. For many of the vendors who are rebadging their existing information security products, intelligence is a marketing term that can mean whatever they want it to mean.

The original Hoover Commission definition from the Cold War is '*Intelligence deals with all the things which should be known in advance of initiating a course of action*' (Clark (1955)).

More recently, the US Army defined intelligence as '*…the product resulting from the collection, processing, integration, evaluation, analysis, and interpretation of available information concerning foreign nations, hostile or potentially hostile forces or elements, or areas of actual or potential operations. The term is also applied to the activity that results in the product and to the organizations engaged in such activity*' (United States Army (2010)).

There is no single agreed definition of intelligence although definitions seem to be converging and sharing some common terminology. Definitions generally vary with regard to the word count (and therefore clarity, many of them taking up multiple sentences) and whether they focus on intelligence as a product or a process.

At the Bank of England Cyber Working Group held on 9 January 2014 the following working definition of threat intelligence was drawn up: '*Threat Intelligence is the contextualised output of a strategically-driven process of collection and analysis of information pertaining to the identities, goals, motivations, tools and tactics of malicious entities intending to harm or undermine a targeted organisation's operations, ICT systems or the information flowing through them*'.

This was followed up at the CBEST threat intelligence workshop held on 13 March 2014, with three further definitions:

- evidence-based interpretation of data, collected on or against the identities, goals, motives, TTPs and targets of malicious actors;
- information that provides relevant and sufficient understanding for mitigating a harmful event;
- the process of analysing data creating contextualised knowledge to mitigate a threat. A threat is an event with potential to cause harm.

For the purposes of this report the second definition above has been adopted and refined:

Intelligence
Information about threats and threat actors that provides sufficient understanding for mitigating a harmful event.

The '*sufficient understanding*' wording in the above definition chimes with a quote from a NSA representative at the RSA 2014 security conference that '*Information doesn't become intelligence until it is useful to someone*' (Bianco (2014a)).

Given the above generic definition, cyber threat intelligence is simply information about threats and threat actors that provides sufficient understanding for mitigating a harmful event in the *cyber domain*.

Most recently, industry analysts have identified three levels of cyber threat intelligence (Gartner (2015)):

- **tactical**: technical intelligence such as using threat indicators to proactively hunt for and defend against adversaries;
- **operational**: intelligence focused on the motivations intent and capabilities (including TTPs) of adversaries;
- **strategic**: intelligence about the risks and implications associated with threats used to inform business decisions and direct cyber security investment.

With regard to existing information security products, intelligence must provide a level of analysis that the customer does not already have which enables them to make new defence-related decisions. Traditional systems tell a customer, for example, that they have been infected by a virus. Intelligence systems, on the other hand, tell them who was behind the attack, why they attacked, their TTPs and key indicators of the attack. This enables customers to respond (rather than react) to the attack and tune their prevention and detection systems (451 Research (2014)).

In the remainder of this document the terms intelligence and threat intelligence are used interchangeably and relate to cyber threats as a default.

3 Process

3.1 Introduction

This section presents an overview of the process underpinning a best practice threat intelligence capability. In other words, the *why*, *what*, *when*, *where* and *how* of threat intelligence. The process is an amalgamation of best practice models and is described in generic terms to make it applicable to in-house intelligence units or external service providers.

3.2 Strategy and scope

Before describing the threat intelligence process it is worth setting out the initial considerations relating to strategy and scope that are typically undertaken in best practice scenarios (Dorrington (2004); KPMG (2013); Verisign (2013)). This ensures that the purpose of the threat intelligence function is both understood and supported by all stakeholders, in particular its customers who issue directives and consume the final products.

As a first step, the broad scope and high-level objectives of the intelligence function need to be set out. Essentially, the job of an intelligence function is to deliver useful and relevant threat intelligence products that will help decision makers protect the organisation's assets against cyber attack and, ultimately, protect shareholder value. However, because intelligence often involves dealing with incomplete information describing ambiguous situations, all stakeholders should share a common understanding of the goals, responsibilities and limitations of the intelligence effort.

A number of constraints on the intelligence function should therefore be clarified. For example:

- what role the intelligence function should take, ranging from informing decision making through to making decisions itself;
- whether the intelligence function will be unable to meet its customer's proposed needs as a result of legal, ethical or practical constraints;
- the extent to which the intelligence function will need to ensure the legal admissibility and evidential weight of all information collected, analysed and disseminated electronically (BSI (2008));
- the accountability and ownership of the intelligence function.

Commercial organisations, like law enforcement agencies, cannot dedicate resources to counter every threat they face. Therefore the allocation of resources to implement a threat intelligence capability should be informed by a prioritised understanding of assets, threats and vulnerabilities.

This will then lead to a strategy for achieving the goals of the intelligence function. Once the strategy has been set out, key measures are identified that will allow the organisation to monitor and measure the performance of, and therefore manage, the strategy. Maturity models (discussed further in Section 5) are useful reference aids for this purpose, enabling the intelligence function to be assessed on a scale of '*informal*' to '*highly repeatable and efficient*', where the highest grade indicates well-documented processes and communications, a high degree of automation and the ability to identify and address insufficiencies in a quantitative manner.

A plan can then be drawn up describing how the strategy will be implemented in detail and who will be responsible for the achievement of specific goals. Having identified what needs to be done, who is going to do it and how it will be measured, the organisation can then identify what organisational structure is required for the threat intelligence function and what infrastructure is needed to support that structure.

3.3 Life cycle overview

Threat intelligence in the commercial arena is still a relatively immature industry. In order to make it more mature the process needs to move away from an informal, *ad hoc* approach towards one that is more rigorous and methodological. As well as making the process more transparent it will also make it more consistent, testable and repeatable across industry sectors. Another way to view this is moving the process from an art to more of a craft. That said, a degree of art will always be required

given the need for human intuition, curiosity and imagination in the process, not to mention the need to configure the process according to the demands of each organisation.

At the highest level, the threat intelligence life cycle follows the classic computing process model of input-process-output. Most threat intelligence functions have adopted this basic model and, while the terms for the three phases vary, the most commonly used ones are collection-analysis-dissemination. They also precede the input/collection phase with an initial direction-setting activity and join everything up into a four-phase cycle (Holland (2013); McMillan and Pratap (2014); Verisign (2013)).

Some threat intelligence functions choose to conclude the cycle with a review activity that leads to an adjustment in future direction and planning. This recognises the concept of continuous improvement as embodied in the plan-do-check-act (or adjust) approach advocated by Deming (1986), a simple method for testing information before making a major decision. In this way the intelligence function can match the pace of change in both the intelligence field and the threat environment. Acting on completion of a task is, in any case, the only way of ensuring the customer derives maximum value from the intelligence they receive.

Interacting with this five-phase cycle are two control functions responsible for managing the efficiency of the intelligence generation process and managing the quality of the intelligence product delivered to the consumer.

The final consolidated life cycle model is presented in **Figure 3.1**.

Figure 3.1 Threat intelligence life cycle model

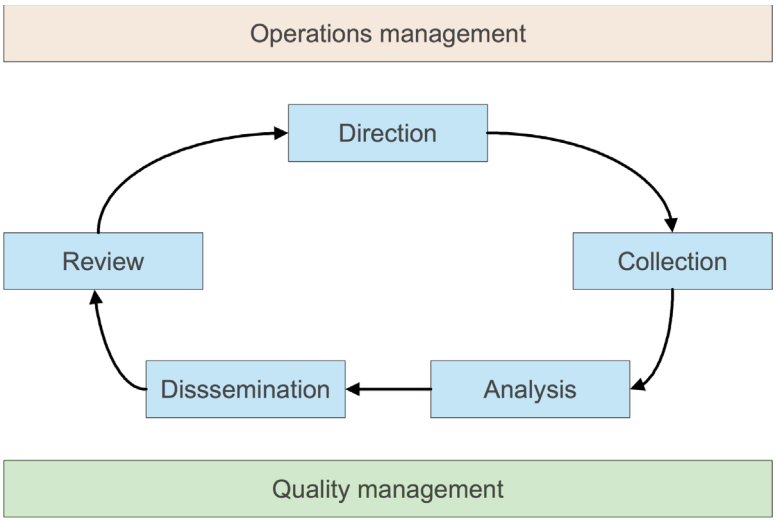

The life cycle model has been kept simple for the purposes of clarity and therefore two caveats should be noted. First, such models can give an over-idealised impression of a neat, linear progression of activity. In operational reality there may well be iteration between certain phases, for example:

- difficulty experienced in collecting certain kinds of data may require some iteration (negotiation) between collection and direction;
- initial analysis may indicate that further data needs to be collected, necessitating some iteration (refinement) between analysis and collection.

This mirrors the general trend in the technology industry towards a more rapid, iterative and incremental style of product development (known as '*agile*') where requirements and solutions evolve through close collaboration with the customer and a rapid and flexible response to changing requirements (Beck *et al* (2001).

The second caveat is that life cycle models can be misread as plans, giving the impression that all activities are of equal complexity and duration, which is not necessarily the case. For example, collection of data, and the subsequent preparation of that data for the purposes of analysis, can be a considerable undertaking.

3.4 Direction

The intelligence cycle begins with direction from the customer, ie an appropriate representative of the body of people consuming the intelligence product. The intelligence manager liaises with the customer, using standard requirements gathering techniques, to help the customer understand and define what they want. In addition to the content of the intelligence product, how it should be presented and disseminated are also covered.

The customer's intelligence requirements typically divide into:

- **long-term directives** that set the broad scope which usually persist for 1–2 years;
- **medium-term directives** oriented towards a particular topic which are usually handled within weeks or months;
- **short-term directives** that are tactical and narrower in scope and are usually handled within days.

The intelligence manager evaluates the customer's requirements to ensure that they are within the remit and capability of the intelligence team. They are then translated into specific intelligence collection, analysis and dissemination tasks. Very often they may start on a small scale and then become incrementally more complex as the intelligence cycle proceeds. Care must be taken to ensure requirements are not misunderstood otherwise this will lead to collection of the wrong data.

Because threat intelligence is a cyclic, evolutionary process, the direction phase will also, where appropriate, take into consideration the results (successes and failures) of the previous loop around the cycle when planning the next round of intelligence tasks. Where there are significant variations between actual and planned results, the root causes are identified and corrective actions put in place to improve the intelligence product. Where there is no need to improve the scope is refined so that it either moves into further detail or focuses on a new area.

3.5 Collection

3.5.1 Overview

On the basis of the customer's intelligence requirements, data is then collected and turned into a format suitable for analysis. Since data is the bedrock for analysis, a due level of care needs to be taken since errors made during this phase can lead to downstream errors during the analysis phase. Collection usually consumes the greatest amount of budget because of the time, effort and cost involved in collecting data from diverse sources.

3.5.2 Intelligence sources

Intelligence data can come from a variety of sources, the most significant ones being:

- **Human Intelligence (HUMINT)**: intelligence derived overtly or covertly from human sources based on a relationship between an intelligence agent and the agent's handler;
- **Covert Human Intelligence Sources (CHIS)**: intelligence obtained by a person who establishes a relationship with another person for the covert purpose of using it to obtain or provide access to any information (eg relating to geopolitical matters). This may be conducted on-line or face-to-face. It also includes intelligence derived from sources on the Deep Web (eg gaining access to a discussion forum using a false identity) that cannot be classified as OSINT since it is not public;
- **Open Source Intelligence (OSINT)**: intelligence derived overtly from publicly available sources.
- **Signals Intelligence (SIGINT)**: intelligence derived overtly or covertly from the interception of signals, whether communications between people (Communications Intelligence or COMINT) or from electronic signals not directly used in communication (Electronic Intelligence or ELINT);
- **Technical Intelligence (TECHINT)**: although this is a variation of SIGINT it should not be confused with intelligence obtained '*by technical means*' in that it does not involve any form of covert activity. One example of TECHINT is the signals generated routinely by hardware devices or software applications connected to an organisation's computer networks (eg log data). Another is the result of examining the inner workings of malicious software code or other technology-based attack methods.

Of these intelligence sources, OSINT, TECHINT and CHIS are the primary area of focus for CBEST because they are the most accessible intelligence for commercial organisations and avoid the risks and legal pitfalls of collecting HUMINT and SIGINT.

A final potential source, particularly relevant to financial services and law enforcement, is financial intelligence (FININT) or gathering information about the financial affairs of entities of interest.

3.5.3 OSINT

According to the US Army Field Manual 2.0, OSINT is '*the discipline that pertains to intelligence produced from publicly available information*'. This information is provided *'without the expectation of privacy' and can be 'published or broadcast for general public consumption; available on request to a member of the general public; lawfully seen or heard by any casual observer; or made available at a meeting open to the general public*' (US Army (2010)).

OSINT sources are many and varied. They include the following (Holland (2013); 451 Research (2014); Chuvakin (2014a)):

- observations and suspicions submitted by employees trained in security awareness;
- general on-line research, such as looking for discussions of particular organisations or topics, monitoring mentions of a company brand, finding personal information about an individual;
- public comments, indicating malicious intent, made by potential threat actors in forums, chat rooms and social media (the cyber personas discussed in Section 2.2.2);
- automated spidering systems that crawl through websites, databases and other data stores;
- capturing and archiving (scraping) web data such as text, graphics, video and audio;
- harvesting or deriving other stores of data not related to network traffic or host-based activity, such as topographical maps, geolocation data, census data, enrolment lists, driver records, transaction histories, metadata, etc;
- automated systems that collect data about attacks such as client and server honeypots, spam traps, phishing traps, botnet traffic emulators, live botnet connections and sinkholes;
- capturing, observing and reverse-engineering malware;
- Border Gateway Protocol (BGP) monitoring;
- Tor usage monitoring;
- log data produced by Security Information and Event Management (SIEM) tools;
- output from detection systems such as anti-virus, anti-malware, anti-spam, network monitoring, network behaviour analysis, intrusion detection and vulnerability scanners;
- output from forensic incident investigations;
- no-cost public threat data feeds (containing, for example, malicious IP addresses, domains and URLs) provided by specialist information security organisations and vendors;
- government intelligence sharing schemes involving law enforcement agencies, government security bodies and Community Emergency Readiness Teams (CERTs);
- industry intelligence from business partners, formal intelligence-sharing bodies (eg Internet Safety Advisory Committee (ISAC), Financial Services Information Sharing and Analysis Center (FS-ISAC), European Network and Information Security Agency (ENISA)) and informal one-on-one links with trusted industry contacts;
- commercial providers of fee-based threat intelligence services, either aggregating and enhancing existing public threat data feeds or providing proprietary threat intelligence products and consultancy based on their own OSINT (and, in some cases, HUMINT) collection and analysis methods.

Both custom and open source tools are available for collecting OSINT. For example, the Collective Intelligence Framework from the Research and Education Networking Information Sharing and Analysis Center (REN-ISAC) aggregates public threat intelligence feeds for ad hoc querying (REN-ISAC (2014)) and software tools are used by a variety of public and private sector organisations to collect and aggregate data, including creating graphical images of relationships between entities of interest (Paterva (2014)).

3.5.4 Collection strategy

Constant monitoring of every possible piece of required data to a high degree of detail is not technically tractable. In addition, too high a volume of data risks overloading downstream analysis. An intelligence function will therefore need to shape its collection strategy according to the following considerations (Verisign (2013)).

Firstly, breadth vs. depth. In most cases finding the data is not the issue; the harder part is filtering out that data that is relevant to the analysis task. A careful balance will therefore be made between '*broad but shallow*' and '*detailed but narrow*'.

Secondly, monitoring frequency. Monitoring is a cyclic process driven by a pulse. The pulse should be sufficiently short such that the monitored entity does not deteriorate beyond correction between pulses. On the other hand it should be sufficiently long such that it does not incur unnecessary computational expense or cause undue delay (Hickman *et al* (1989)). Monitoring frequency will therefore be one or a combination of the following:

- **periodic monitoring**: monitoring the environment at a regular frequency (or pulse rate) which may range from minutes to months or more;
- **analysis-driven monitoring**: monitoring the environment in an *ad-hoc* manner which is driven by the current state of the analysis (eg the current hypothesis under consideration);
- **event-driven monitoring**: monitoring the environment in an *ad-hoc* manner that is driven by specific events occurring, or anticipated to occur, within the threat landscape (ie data-driven).

Each of these has its advantages and disadvantages. The aim is to ensure maximum efficiency and effectiveness of the analysis given the customer's requirements and the characteristics of the cyber threat.

The above monitoring methods are examples of pull-type monitoring. There is also the option of implementing push-type monitoring whereby a near real-time source of data flows in from an organisation's technical data feeds (eg log data) and social media sites. The result of this constant monitoring is then stored in a historical database ready to add valuable context when an intelligence request comes in and the formal intelligence cycle begins.

3.5.5 Processing

Before raw data can be analysed it needs to be processed to render it amenable to downstream analysis. This will involve the use of automated tools that can perform useful data processing functions such as parsing, correlating, filtering, de-duplicating and aggregating. This is a critical and often-overlooked step in the threat intelligence cycle. Although it is generally described as being part of the collection phase it could just as easily be said to form a bridge between collection and analysis.

3.6 Analysis

3.6.1 Overview

As a generic knowledge-based task, analysis can be decomposed into a taxonomy of sub-types that include classification, diagnosis, assessment, monitoring and prediction. During the analysis phase raw data is transformed into information in the form of patterns, trends, clusters, sequences and so on. This is achieved via a series of primitive inferences such as selection, classification, abstraction, decomposition, specification, comparison, matching, instantiation, correlation and transformation (Hickman *et al* (1989)). If the information generated by analysis provides sufficient understanding for mitigating a harmful event then it can be termed intelligence.

The above definition is more or less echoed by the US Army's Field Manual 2.0 which describes analysis as: '*The process by which collected information is evaluated and integrated with existing information to produce intelligence that describes the current — and attempts to predict the future — impact of the threat…on operations*' (US Army (2010)).

3.6.2 Analytical strategies

Analysis is carried out using a mix of machines and human analysts. Machines typically perform simpler, high volume tasks that reduce a huge amount of input data down to a more manageable subset with a lower signal-to-noise ratio. Human analysts then apply a critical level of judgement to this filtered data to ensure the final intelligence product contains minimal false positives.

Analytical strategies can be data-driven or hypothesis-driven depending on the intelligence requirement. Machines can play a leading number-crunching role during data-driven analysis and human analysts can apply their intuition, curiosity and imagination during hypothesis-driven analysis, the most effective approach being a combination of innovative human and systematic machine. An example of these two strategic approaches applied to the generic analytical task of assessment is shown in **Figure 3.2** (Hickman *et al* (1989)).

Figure 3.2 Data-driven and hypothesis-driven assessment strategies

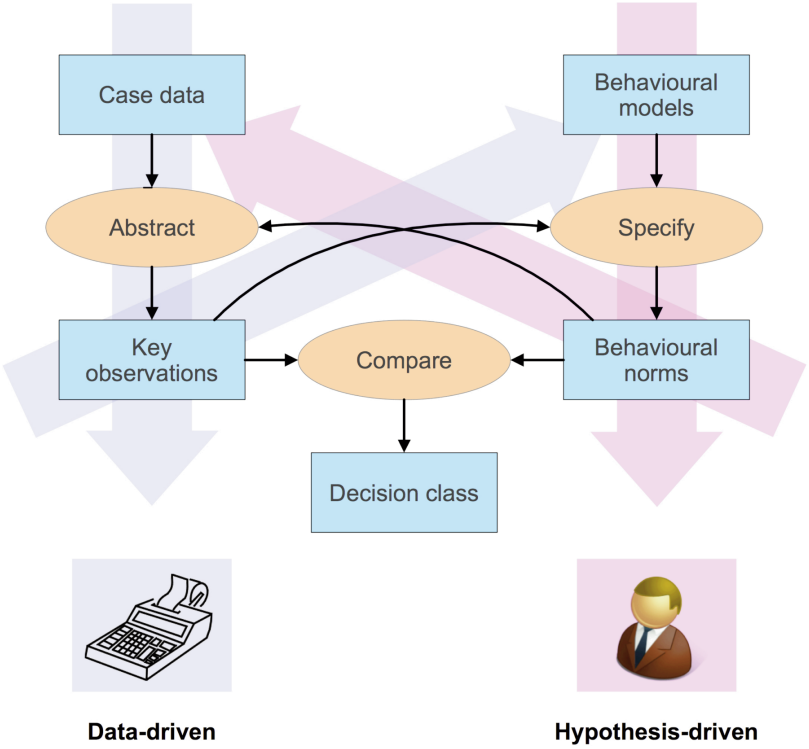

3.6.3 Machine-based analytical techniques

Using machines to undertake or support intelligence analysis is a mature discipline based on artificial intelligence research conducted since the 1960s. A summary of the techniques, categorised according to the kind of threat being analysed, is given below.

Known knowns

These are threats previously encountered and recognised by means of identifying similar characteristics. This is based on an analyst's knowledge and expertise that in turn can be expressed in the form of a production rule or other form of machine-executable algorithm.

Unknown knowns

These are threats that are known about but have never been seen, or have been previously seen but are not recognised now because of altered characteristics. These can be identified using matching techniques such as:

- **hard matching**: where a threat is identified by matching against a repeated identifier;
- **fuzzy matching**: where repeat identifiers are resolved through a fuzzier form of matching that returns a list of results based on likely relevance even though the exact words and spellings may not match exactly;
- **geo-matching**: using geolocation data for identifying clusters of significant activity or hotspots;
- **social network analysis**: identifying networks of new, unknown threats on the basis of their association with other, known, threats.

Unknown unknowns

These are threats that have not been previously encountered. They can be identified using two categories of techniques:

- **supervised learning**: where threat characteristics are induced by the machine (eg in the form of rules) on the basis of historical examples where the outcome (threat or no threat) is known;
- **unsupervised learning**: where the machine applies an automated learning technique (eg neural networks such as multi-layer perceptrons and self-organising maps) on a large quantity of data in order to determine threat characteristics for itself.

These techniques are summarised in **Figure 3.3**.

Figure 3.3 Machine-based analytical techniques

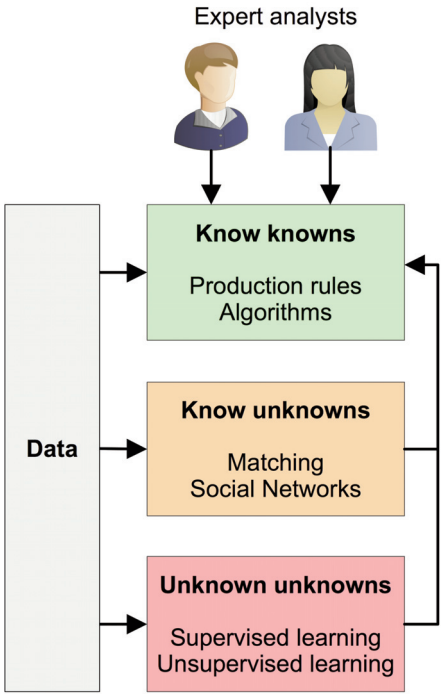

3.7 Dissemination

3.7.1 Overview

Once the data has been collected and analysed, the intelligence function then disseminates the intelligence product to its consumers. Recipients of intelligence reports will be located at strategic, operational and tactical levels inside business functions (primarily risk and security) and technology functions (primarily operations). Suitably summarised intelligence reports also go to senior executives.

Dissemination is not a trivial undertaking. If the consumer is to accept and benefit from an intelligence product then three essential criteria need to be met:

- **the right content**: good quality intelligence must provide sufficient understanding to allow consumers to mitigate a harmful event;
- **the right presentation**: intelligence must be concise, understandable (ie jargon-free and matching the language of the recipient) and strike the right balance between narrative, tables, numbers, graphics and multimedia;
- **the right time**: intelligence must be disseminated within a time frame that enables its consumers to make effective, proactive decisions.

The above criteria are highly interconnected. The best intelligence in the world will be useless if it cannot be understood or arrives too late. Similarly a '*style over substance*' situation can potentially arise when glossy intelligence reports mask low-grade content. In the CBEST scenario it will therefore be essential that the right language, using the same set of base data, be chosen to make threat intelligence understandable and usable for penetration testing providers and financial institutions.

Further details on the content of a threat intelligence model can be found in the sister CBEST report *An Introduction to Cyber Threat Modelling* (CBEST (2016a)).

3.7.2 Forms of delivery

Intelligence products are typically delivered to their consumers in the following ways:

- **simple alerts** sent out by phone, text or email;
- **detailed reports** comprising narrative enriched with tables, numbers, graphics and multimedia;

- **machine-readable data feeds** based on a proprietary or open standard structured threat intelligence notation (described in more detail below), for Security Information and Event Management (SIEM), anti-virus software, firewalls, intrusion prevention systems (IPS), intrusion detection systems (IDS) and forensic tools;
- **custom-designed output** for in-house systems;
- **Application Programming Interfaces** (APIs) enabling direct system connection for the purposes of intelligence query or retrieval;
- **secure online portals** providing on-demand access to an always-current threat intelligence database and a range of analytical functions that could be as basic as from simple queries to more complex data mining.

Each of these has its advantages and disadvantages. Customised threat intelligence reports in a narrative format are much easier for business people to understand. On the other hand, it can be difficult to translate these reports quickly and easily into a format that machines can read and factor into the tasks they execute. This is covered in more detail below.

3.7.3 Machine-readable threat intelligence

CBEST threat intelligence will be disseminated from threat intelligence suppliers to security testers in a form closer to the narrative (human readable) end of the spectrum than the machine-readable end. Nonetheless machine-readable threat intelligence (MRTI) is an important and growing topic. Given the rapidly changing nature of the threat, intelligence must be acted on quickly to receive its full value. In many cases its value can reduce to zero in days or even hours (Farnham (2013)).

Thus the term '*machine-readable threat intelligence*' has been coined by industry analysts, defining it as '*a capability that allows SIEM and other security controls to make operational security decisions based on information about the prevailing threat landscape*' (Gartner (2014a)).

Other analysts have commented: '*By year-end 2016, 'threat intelligence broker' offerings will emerge, providing machine-readable threat intelligence from multiple sources to an array of technical security controls, independent of vendor*' (Lawson and McMillan (2014)). Another analyst describes the goal '*to make intelligence sharing with context occur at wire speed*' (Holland (2013)).

MRTI, as described above, takes the form of a direct data feed for automated prevention and detection systems. This contrasts with intelligence that takes the default form of verbal or written narrative (or what might be termed '*human-readable threat intelligence*').

MRTI is based on a structured threat intelligence notation of which there are a large number currently under development. These emerging standard notations are both vendor and community sponsored and many of them overlap. Most of them are based on XML (Extensible Markup Language), a well-recognised standard for encoding a document in a format that is both human-readable and machine-readable. Standards written in XML are also highly extensible.

The intelligence these standards describe is mainly tactical and technical in nature (ie the TECHINT described in Section 3.5.2), the main focus being indicators of compromise (IOCs). These are forensic remnants of an intrusion residing in operating system and network devices, eg IP addresses, domain names, uniform resource locators, file hashes, registry key values, service start-ups and HTTP requests. An analyst can piece together these breadcrumbs from endpoints and networks to understand the anatomy of an attack.

Despite the focus being on IOC TECHINT, there are signs that softer, human-oriented indicators, such as motivation and geopolitical markers, are beginning to appear in some of the standards. A number of organisations are now using specialist tools to export and manage data that has been fused from both machine-oriented and human-oriented threat intelligence (Paterva (2014)).

Some examples of the more prominent vendor and community-sponsored standards are given below.

OpenIOC (Mandiant)

Open Indicators of Compromise (OpenIOC), launched by Mandiant in 2011, is an extensible XML schema for defining and sharing threat indicators (OpenIOC (2014)). Although this is a vendor-sponsored emerging standard, it has also been released as an open standard and endorses a community look and feel. The schema provides a comprehensive set of attributes (around 500) for

defining technical IOCs in considerable detail. As well as being extensible, OpenIOC can be converted or parsed to other formats that might contain information that could feed into or benefit from the threat information contained in an IOC.

OpenIOC is primarily used in Mandiant products but others are making use of it. The OIC Bucket website provides a community resource to submit and share OpenIOC files (IOC Bucket (2014)). McAfee has released OpenIOC files for operation Troy and lists several McAfee products that can consume OpenIOC files (Walter (2013)). An open source project, pyioc, is also available which provides a set of tools to handle OpenIOC files (Bryner (2013)).

CybOX (Mitre)

Cyber Observable eXpression (CybOX), launched by Mitre in 2010, is an extensible XML schema for defining and sharing IOC details known as observables (Mitre (2014a)). CybOX provides over 70 pre-defined objects that can be used to define observable technical security events or stateful properties.

STIX (Mitre)

Structured Threat Information eXpression (STIX), launched by Mitre in 2012, is an extensible XML schema for defining and sharing IOCs and TTPs (Mitre (2014b)). These consist, in part, of observables defined using CybOX (see above). Relationships between constructs can also be defined, eg a TTP can be related to a specific threat actor. Extensions have been defined to inter-operate with other standards such as TLP, OpenIOC, Snort and YARA Editor. STIX is sponsored by the US Department of Homeland Security (DHS) and maintained by Mitre.

CRITs (Mitre)

Collaborative Research Into Threats (CRITs), under development by Mitre, is a threat intelligence gathering and analysis platform that makes use of STIX and related formats (Mitre (2014c)). It is intended for use by a select community of threat-sharing partners inside a '*walled garden*'. There is little public information available and the platform remains at a prototype stage.

TAXII (Mitre)

Trusted Automated eXchange of Indicator Information (TAXII), launched by Mitre in 2012, defines a set of services and message exchanges for sharing threat intelligence (Mitre (2014d)). TAXII is the preferred method of exchanging STIX (see above). It uses XML and HTTP for message content and transport. It supports a range of sharing models, such as hub-and-spoke or peer-to-peer, operating in push or pull modes. Users can categorise and share intelligence with the partners they choose.

TAXII is not a specific information sharing initiative and does not attempt to define trust agreements, governance, or other non-technical aspects of intelligence sharing. TAXII has been adopted as a planned standard by Microsoft as part of its Microsoft Active Protections Program (MAPP) (Bryant (2013)). TAXII is also used by the Financial Services Information Sharing Analysis Center (FS-ISAC) (Connolly (2013)).

MAEC (Mitre)

Malware Attribute Enumeration and Characterization (MAEC), under development by Mitre, is a standardised language for defining and sharing intelligence about malware based upon attributes such as behaviours, artefacts and attack patterns (Mitre (2014e)). This compares with the single metadata entity commonly employed in signature-based malware detection.

CVE (Mitre)

Common Vulnerabilities and Exposures (CVE), under development by Mitre, is a dictionary of common names (CVE Identifiers) for publicly known information security vulnerabilities (direct criminal access) and exposures (indirect access or '*stepping stones*') (Mitre (2014f)).

IODEF (Internet Engineering Task Force)

Incident Object Description and Exchange Format (IODEF), launched in 2007 by the Internet Engineering Task Force, is an extensible XML schema for sharing information about computer security incidents between computer security incident response teams (CSIRTs) (Danyliw, Meijer and Demchenko (2007)). IODEF is also referred to as IETF Request For Comments (RFC) 5070. It provides a data model (over 30 classes and sub-classes) to accommodate most commonly exchanged data elements and associated context for indicators and incidents. It also offers an approach (albeit a limited one) for documenting the workflow necessary for triaging and communicating this data.

Being community-driven, IODEF is highly extensible and benefits from refinement and improvement by an open source community independent of a vendor sponsor. The Anti-Phishing Working Group has extended the IODEF standard to support the reporting of phishing and other email incidents. It is also being used as a storage format in the Collective Intelligence Framework (see below) and features in products from DFLabs, Arcsite and Foundstone (Moriarty (2013)).

IODEF-SCI (Managed Incident Lightweight Exchange)

IODEF for Structured Cyber security Information (IODEF-SCI) is an extension to the IODEF standard proposed by the Managed Incident Lightweight Exchange (MILE) working group (Takahashi (2013)). The extension adds support for additional information that comprises attack pattern, platform, vulnerability, scoring, weakness, event report, verification and remediation.

RID (Internet Engineering Task Force)

Real time Inter-network Defense (RID), under development by the Internet Engineering Task Force, is a standard for sharing incident-handling data (Moriarty (2012)). RID is also known as IETF Request For Comments (RFC) 6545. The RID XML schema is based on IODEF with extensions.

VERIS (Verizon)

Vocabulary for Event Recording and Incident Sharing, launched in 2010 by Verizon, is a schema for collecting and sharing security incident intelligence (VERIS (2014)). The schema covers victim demographics, incident description, discovery and response, impact assessment and, in a limited manner, IOCs. VERIS is intended for strategic usage rather than tactical. It underpins the annual Verizon Data Breach Investigation Report.

OTX (AlienVault)

Open Threat eXchange (OTX), launched in 2012 by AlienVault, is a publicly available threat intelligence sharing service (AlienVault (2013)). OTX interoperates with the company's Open Source SIEM (OSSIM) system. OSSIM users can configure their SIEM systems to upload threat intelligence to OTX. Collected intelligence is validated by AlienVault and then delivered to all OSSIM users that subscribe to OTX. OTX threat intelligence can also be accessed by Collective Intelligence Framework users (see below). Unlike closed, invitation-only intelligence sharing networks, OTX intelligence is available to anyone who chooses to participate. It therefore claims to be the world's most authoritative crowd-sourced threat intelligence exchange.

CIF (REN-ISAC)

The Collective Intelligence Framework (CIF), developed out of the Research and Education Network Information Sharing and Analysis Center (REN-ISAC) in 2009, is a system for managing threat intelligence (CIF Project (2009)). CIF is based on the IODEF standard. Users can combine known malicious threat intelligence from multiple sources and deploy it accordingly. The most common types of threat intelligence warehoused in CIF are IP addresses, domains and URLs related to malicious activity. CIF also includes information on the type of threat, severity of an attack and the confidence in the data. CIF provides the ability to control access and place restriction levels on intelligence. The system is being used by REN-ISAC members.

3.7.4 Threat intelligence platforms

Recent years have also seen the growth of threat intelligence platforms (TIPs) that enable organisations to manage incoming structured and unstructured threat intelligence that is typically delivered by MRTI data feeds. While each platform varies in functionality, most offer some or all of the following generic capabilities (Gartner (2014c)):

- collection and normalisation of MRTI (typically STIX-based) from multiple sources;
- correlation, pivoting and enrichment of data in order to add context;
- categorisation into indicators of compromise, threat actor type, geography, etc;
- integration of derived information into downstream security prevention and detection tools;
- co-ordination of the workflow of multiple users during incident response;
- sharing derived intelligence with other organisations at wire speed.

Each TIP tends to focus on a particular area. Some are SOC-oriented and focus on deploying MRTI in operations. Others are analyst-oriented, focusing on near-term data analysis and longer-term risk analysis. Others focus on integration between external threat intelligence feeds and internal SIEM systems. But all act as a single funnel for channelling and analysing the growing fire hose of MRTI (particularly newly-emerging exploits and vulnerabilities) that emanate from multiple threat intelligence services and open-source organisations (Dark Reading (2015)).

For many analysts this represents a '*third wave*' of MRTI integration. This began with organisations employing single MRTI data feeds, followed by multiple MRTI feeds with limited analysis, and now enhanced analysis, sharing and integration (Gartner (2014c)).

3.7.5 External intelligence sharing

As well as exploiting intelligence for their own benefit, both producers and consumers of threat intelligence can share intelligence with external industry and government organisations. This helps resolve the asymmetry that exists whereby cyber criminals gain value from collaborating and sharing while their targets behave like individual islands. Intelligence sharing enables organisations to maintain awareness of a highly dynamic and diverse spectrum of threats and arguably the financial sector represents one of the best examples of this (Payments Council (2014)). When intelligence is shared then the adversary has to make a mistake just once, ie be detected, and all the defenders will know about it (451 Research (2014)).

Intelligence sharing, where the number of targets outweighs the number of attackers and the end result is great than the sum of its parts, is therefore seen as the next wave of cyber security. That said, intelligence sharing needs to be performed judiciously, ie understanding what is appropriate to share with a wider audience versus a single trusted partner. It may be that too wide a dissemination could tip off an attacker.

While threat intelligence professionals find value in sharing threat information through informal, personal networks, the results are inconsistent and unscalable. It is acknowledged in the cyber security industry that better frameworks are needed for communicating threat intelligence. Such frameworks should include:

- standardised reporting terminology and processes;
- indemnification against liability for information sharing or directed action for cyber security purposes;
- the ability for users to create circles of trusted peers and specify which elements of their own threat intelligence they want to share as well as how, when, where and to whom;
- a technical infrastructure to share and analyse threat intelligence at machine speed.

Structuring principles are either hub and spoke, hierarchical, network or a combination of these. Each model has implications for the rate of dissemination, the transmitter's control over what is sent, flexibility of follow-on inquiry and anonymity (Jellenc (2013)).

In absence of an industry-standard framework, current sharing mechanisms include:

- private or restricted face-to-face meetings and phone calls;
- emails, forums and message boards;
- web portals with wiki-type capabilities;
- web portals acting as document management systems, usually holding PDF or Microsoft Word™ files;
- web portals (some with APIs) allowing downloads of structured (or semi-structured) data;
- web portals offering social networking facilities with secure access and sharing controls.

The aim is to build a community of members who are incentivised to trust one another. Because trust usually happens between individuals rather than organisations, it is worth noting that any formalised and sophisticated intelligence-sharing scheme will always be underpinned by informal and elementary trust-based mechanisms; these never go away. Members may start out trusting one another and creating their own trusted circles but if conflicts arise or certain members are replaced, trust can diminish and with it the flow of useful intelligence (451 Research (2014)).

Threat intelligence sharing initiatives are many and varied. In the United Kingdom most intelligence sharing takes the form of government-to-industry intelligence hubs. Intelligence sharing in the other direction is not so well established, the most visible example being the Cyber Security Information Sharing Partnership (CISP) that was used by participants in the Waking Shark II exercise (held in November 2013) to share real-time threat information as the scenario unrolled. In the United States intelligence sharing is more established. A lineage of successful intelligence sharing programs (spearheaded by the financial services sector and the Defense Industrial Base) has become a core pillar of the country's cyber security infrastructure (Jellenc (2013)).

Intelligence sharing initiatives include those listed in **Table 3.1**.

Table 3.1 Threat intelligence sharing initiatives

Acronym	Name	Further information
FS-ISAC	Financial Services Information Sharing and Analysis Center	https://www.fsisac.com
DIB CS/IA	Defence Industrial Base Cyber Security/Information Assurance Program	http://dibnet.dod.mil/
DSIE	Defense Security Information Exchange	www.whitehouse.gov/files/documents/cyber/Defense%20Security%20Information%20Exchange%20-%20DSIE%20summary%20-%20William%20Ennis.pdf
NCFTA	National Cyber-Forensics & Training Alliance	www.ncfta.net
REN-ISAC	Research and Education Networking Information Sharing and Analysis Center	www.ren-isac.net
Ops-T	Operations Security Trust	https://portal.ops-trust.net
n/a	IID ActiveTrust	http://internetidentity.com
ENISA	European Network and Information Security Agency	https://www.enisa.europa.eu
CERT-EU	Computer Emergency Response Team for the EU institutions, bodies and agencies	https://cert.europa.eu/cert/filteredition/en/CERT-LatestNews.html
FSIE	Financial Services Information Exchange	http://itlaw.wikia.com/wiki/Financial_Services_Information_Exchange
NCA	National Crime Agency	www.nationalcrimeagency.gov.uk
NCCU	National Cyber Crime Unit	www.nationalcrimeagency.gov.uk/about-us/what-we-do/national-cyber-crime-unit
UKCERT	UK Computer Emergency Response Team	www.ukcert.org.uk
GovCertUK	Computer Emergency Response Team for UK Government	www.cesg.gov.uk/policyguidance/GovCertUK/Pages/index.aspx
CISP	Cyber Security Information Sharing Partnership	https://www.cisp.org.uk
IBSIG	Investment Banking Special Interest Group	n/a (hosted meetings)
CSIG	Cybersecurity Special Interest Group	n/a (hosted meetings)
FCAS	British Bankers' Association Financial Crime Alerts Service (under development)	https://www.bba.org.uk/news/bba-voice/uniting-to-tackle-financial-crime/#.VfvJH7R77v0

With regard to TIPs discussed above, many of them are being used to share and receive MRTI among a large number of user-created '*circles of trust*'. Only members who have been manually verified as being clients of the threat intelligence vendors' data feeds are able to access them. Many TIPs are also employing trust methods such as the Traffic Light Protocol described in the following section.

3.7.6 Establishing trust

Although attractive in principle, large-scale intelligence sharing across government and private sector organisations faces the challenge of establishing trust to incentivise organisations to share commercially sensitive information. There is a direct conflict between maintaining the value of what is essentially secret knowledge about an adversary and sharing (or selling) it as widely as possible (451 Research (2014)). In addition to the barriers that can exist between government and the private sector, the defence contractors and ICT companies targeted by APTs are, at different times and in different contexts, sometimes partners, sometimes rivals and sometimes vendors and consumers of each other's goods and services (Jellenc (2013)).

Fear of legal action and reputational damage appears to be one of the biggest impediments to sharing threat intelligence. Participants must find ways credibly to commit to only using shared intelligence for the expressed purposes that warranted the

sharing in the first place. There are also concerns about the negative reputational consequences that could result from an organisation disclosing the fact that is has been successfully infiltrated by a hostile agent (Jellenc (2013)).

Assuring the quality of the membership identification and induction process is one element of the trust solution. The aim is to ensure that members of an intelligence-sharing scheme are sufficiently:

- competent in their ability to have access to, and share, good quality threat intelligence;
- trustworthy and discreet enough to not engender reticence in any member that would hinder sharing.

To this end, membership identification and induction options include:

- open applications evaluated individually;
- criteria and threshold-based applications;
- preconfigured (via a trusted third party, eg UKCERT) organisational memberships with designated individuals;
- vouchsafing by existing members only.

With regard to ensuring that participants share and use intelligence only in ways that the group as a whole deems acceptable, there are three mechanisms:

- security classification;
- legal-regulatory;
- trust-building via acts of reciprocity.

In practice, one or a combination of these mechanisms can be used. Focusing on security classification, a range of classification schemes are available for controlling the way in which intelligence is shared.

For example, the UK Government Security Classifications Policy defines:

- **Top secret**: information whose release could cause considerable loss of life, international diplomatic incidents, or severely impact on-going intelligence operations;
- **Secret**: information whose release could cause serious harm such as threats to life, compromising major crime investigations or harming international relations;
- **Official-Sensitive**: information whose release could have more damaging consequences (for individuals, an organisation or government generally) if it were lost, stolen or published in the media;
- **Official**: all routine public sector business, operations and services.

The above scheme, applicable from April 2014, replaces the older '*Top Secret-Secret-Confidential-Restricted-Protect Government*' Protective Marking Scheme.

Similarly, the US Department of Defense classifies intelligence as:

- TS-SAP (Top Secret — Special Access Program);
- TS-SCI (Top Secret — Sensitive Compartmented Information);
- TS (Top Secret);
- Secret;
- Confidential.

Of relevance to threat intelligence sharing programmes is the newer US Department of Defense classification of '*Controlled Unclassified Information*' (CUI). This is information that is technically unclassified but of sufficient sensitivity to warrant special handling. This is relevant because it overcomes the following impasse:

- if too much intelligence becomes classified then it cannot be used effectively to mitigate a harmful event;
- however, if intelligence is not given any kind of protected designation then participants will be suspicious about how intelligence might be used once it has begun circulating among the group.

The new CUI designation is reinforced by legal-regulatory mechanisms, binding contracts (viz Non-Disclosure Agreements) and informal trust-based commitments. Current CUI labels are as follows:

- Dissemination and Extraction of Information Controlled by Originator (ORCON);
- For Official Use Only (FOUO);
- Caution — Proprietary Information Involved (PROPIN);
- Not Releasable to Foreign Nationals (NOFORN);
- Law Enforcement Sensitive (LES);
- Limited Distribution (LIMDIS);
- Authorised for Release to [name of country(ies)/international organization] (REL TO [NAME]).

In comparison, the intelligence sharing rules defined by the US Defense Security Information Exchange are relatively simple. All members and their companies are subject to a Non-Disclosure Agreement that permits sharing on three levels:

- Non-Attributional;
- For DSIE Eyes Only;
- Public Domain information.

All intelligence shared within the DSIE is by default non-attributional and cannot be shared outside the constraints of the DSIE without permission from the intelligence owner.

The US Multi-State Information Sharing and Analysis Center defines an equally simple Traffic Light Protocol to ensure that sensitive information is shared with the correct audience. Based on the schema originally developed by the UK National Infrastructure Security Co-ordination Centre (NISCC), it employs four colours to indicate different degrees of sensitivity and the corresponding sharing considerations to be applied by the recipient(s) (MS-ISAC (2014)). The protocol is similarly used by the US Computer Emergency Readiness Team (US-CERT) (US-CERT (2015)). The protocol is summarised in **Table 3.2**.

Table 3.2 MS-ISAC/US-CERT Traffic Light Protocol

Classification	Summary	Example
Red	Personal for named recipients only	In the context of a meeting, Red information is limited to those present at the meeting. In most circumstances Red information will be passed verbally or in person.
Amber	Limited distribution	The recipient may share Amber information with others within their organisation but only on a 'need-to-know' basis. The originator may be expected to specify the intended limits of that sharing.
Green	Community-wide	Green information can be circulated widely within a particular community. However, the information may not be published or posted publicly on the Internet nor released outside of that community.
White	Unlimited	Subject to standard copyright rules, White information may be distributed freely, without restriction.

With regard to the Green classification above, the current definition is somewhat contradictory and is likely to be refined in order to make it logically consistent.

Finally, the UK National Intelligence Model '*5x5x5*' includes series of codes that specify how a piece of intelligence should be handled (NCIS (2000)):

- 1: Open source — no restrictions;
- 2: Restricted to clients only;
- 3: Restricted to specific clients;
- 4: Restricted to specific clients with conditions;
- 5: No dissemination without authority.

3.8 Review

3.8.1 Overview

Once the intelligence product has been received and used by its business and technology consumers, it is then time to review the success or failure of the operation and gauge overall customer satisfaction. The key question therefore is whether the intelligence product provided sufficient understanding for mitigating a harmful event.

The review phase is the last of the five-phase intelligence cycle and, depending on its outcome, will either terminate the cycle or begin another loop around the cycle in order to adapt and enhance the intelligence product. Most customers are rarely satisfied with a single delivery since receiving an intelligence product usually acts as a catalyst for additional requests.

The intelligence function must therefore manage its customer's expectations, educating them about any novel issues that have arisen, and ensure any subsequent loop around the cycle is scoped carefully. Review is also important because intelligence has a shelf life and must therefore be reviewed periodically to ensure it remains relevant.

The question of whether the intelligence was sufficient is particularly relevant for organisations seeking to take a more proactive stance. During a cyber attack, there is a natural tension between monitoring the attack to gain further intelligence about the adversary versus disrupting the attack and minimising harm.

Review corresponds to the check phase of the plan-do-check-act (or adjust) continuous improvement approach (Deming (1986)). This means:

- comparing actual results against expected results to ascertain any differences then understanding and documenting why this occurred;
- looking for any deviations in the implementation of the original plan;
- reviewing previous cycles around the loop in order to spot any trends.

The results of the review will then feed into the next round of intelligence planning, ie the direction phase.

3.8.2 Tuning detection

The purpose of intelligence is to help mitigate a harmful event. Given is it not possible to completely deter or prevent all attempts at cyber attack then such attacks need to be detected as quickly as possible. The two overriding considerations are:

- find all attempts at attack that is reasonably possible so that action can be taken to prevent further harm;
- do this as quickly as possible so as to limit the harm.

Proactive detection can be measured (and optimised) using the following criteria:

- time taken to detect;
- true positive vs. false positive ratio;
- true positive vs. false negative ratio.

Detection is founded on two things: knowing what to look for and looking for it. Intelligence tuning therefore involves the following:

- identifying and refining detection typologies;
- modifying the trigger conditions of these typologies to optimise the actual detection rate against the false negative rate;
- establishing the frequency at which these typologies should be applied (the monitoring pulse mentioned earlier), thereby balancing performance with cost of use.

Intelligence tuning is not necessarily about improving the actual detection rate. This might come at the cost of increased investment resources, both capital and human, leading to diminishing returns.

In addition, focusing too closely on one small, highly predictable group might give a very good false positive rate, but at the expense of failing to identify other cases that a more generalised, but less accurate, approach may have detected (ie false positive rate improves at the cost of a degrading false negative rate).

Finally, as damage may accrue as time passes, it will be important to focus on the time-to-detect rate. This means accuracy (false positives) might be sacrificed in order to detect serious or complex cases earlier.

3.8.3 Confirmation bias

Another important factor to consider during review of intelligence is confirmation bias. This is a phenomenon whereby analysts actively seek out and assign more weight to evidence that confirms their hypothesis and ignore or assign less weight to evidence that could refute their hypothesis.

Analysts display this cognitive bias when they gather or remember information selectively, or when they interpret it in a biased way. They also tend to interpret ambiguous evidence as supporting their existing position. This results in statistical errors that can in turn lead to flawed intelligence.

A similar phenomenon, groupthink, occurs when groups share analytical judgements at face value without thorough questioning (Digital Shadows (2015)).

3.9 Operations and quality management

Interacting with the five-phase intelligence cycle are two key management functions. The first of these, operations management, maintains the efficiency of the intelligence generation process. This single point of control simplifies interactions and eliminates duplication of effort. Key tasks include:

- assessing intelligence requests from customers in terms of their scope, requirement and intended outcome;
- assessing the feasibility of fulfilling a request given time, resource and capability constraints;
- prioritising requests and then allocating and managing resources through the intelligence cycle.

An important responsibility is to regulate the flow of requests to avoid the intelligence function being overwhelmed by too much data or too many tasks. If an intelligence function is stressed it may fixate on the influx of new data rather than dealing with the data it has in hand. This can be managed by developing a '*battle rhythm*', a military management doctrine that uses Standard Operating Procedures (SOPs) to maintain control over personnel and assets in extremely stressful situations. This also allows operations management to assess the capacity at which the intelligence function is running (Verisign (2013)).

Quality management, the second management function, maintains the quality of the intelligence product delivered to the consumer. Quality management consists of quality assurance and quality control. Quality assurance is a set of activities for ensuring quality in the processes by which intelligence products are developed. Usually conducted by an external review team, it is process-oriented and focuses on defect prevention or '*doing things right*'. Quality control is a set of activities for ensuring quality in threat intelligence products. Usually conducted by an internal review team, eg during an end-of-cycle review, it is product-oriented and focuses on defect identification or '*doing the right things*'.

4 Organisation

4.1 Introduction

This section presents an overview of the organisation, roles and skills required for running a best practice threat intelligence capability. In other words, having described the processes (*what*, *when*, *where* and *how*), we now consider the *who*.

Because threat intelligence remains in an early stage of definition it is not possible to define an exhaustive list of the organisational elements needed to produce it well. Nonetheless there is enough information available to produce an initial specification. This subject is developed further in the sister report *CBEST Services Assessment Guide* (CBEST (2016b)).

4.2 Structure

While an intelligence function organisation can be structured a number of different ways, the most commonly adopted are hierarchical and flat models.

The advantage of the classic hierarchical management model is the clear chain of command that enables it to respond to the customer quickly. It also closely aligns with government and military functions. A hierarchical structure provides support to team members who are new to intelligence. However, the layered structure can cause intelligence to be degraded or lost between command layers. This can create delays and slow the analysis of new information. This delay is informally known as '*blink potential*' (Verisign (2013)).

In the flat structure team members can easily pass data between peers. This fluid, '*unblinking eye*' mode of operation reduces the risk of loss of intelligence. The freedom of a flat management structure also suits more experienced analysts. However, the lack of management control and oversight can have negative consequences, particularly with regard to preventing scope creep and co-ordinating operations (Verisign (2013)). In practice a hybrid approach is often adopted, such as federated or hub-and-spoke models where there is a mix of centralised control and distributed execution.

Another question is where an intelligence function should be located within an organisation. It could, for example, reside within IT, operations, finance, risk or (information/physical) security. In each case the function will reap the benefits that arise from being integrated with that part of the business, ie technical efficiency, business insight, return on investment, independent oversight and security capability respectively. On the other hand the function can become niche, biased and blinkered as a result of residing in a business silo.

Many organisations will choose to subsume their intelligence function within their existing (information/physical) security function or, alternatively, transform their security function into an intelligence function. Still others will establish an intelligence function so that it is completely separate from the existing security function. However, given tight budgets and the constantly evolving threat landscape, maintaining two operations may not be sustainable. Intelligence and security are increasingly being viewed as partners since the intelligence product forms an input to the security cycle (Verisign (2013)).

Regardless of where the intelligence function fits in the organisation, it should be conducted in partnership with other organisational functions that may be consumers of intelligence or valuable sources of input for intelligence products. This includes risk, finance, marketing, legal, human resources, internal audit, IT development and IT operations.

Organisational structure aside, with regard to organisational culture, the best environment is one that encourages self-awareness, peer review and questioning of existing procedures.

4.3 Roles

An intelligence function typically consists of managers, analysts and operations specialists. They may be organised according to the phases of the intelligence life cycle (described in Section 3.3) or according to the categories of threat actors or TTPs. Some of the specialist roles include the following (Verisign (2013); Jellenc (2013)):

- **intelligence co-ordinator**: sets scope and direction, liaises with customers and is accountable for the team's performance and operational security;
- **operations team**: handles administrative and co-ordination functions including refining customer requirements, allocating resources and managing them through the intelligence cycle;
- **quality team**: manages the quality of the intelligence product delivered to the customer by undertaking quality assurance and quality control activities;
- **intelligence analysts**: perform core intelligence tasks focusing on social, cultural and geopolitical analysis, with generalists synthesising disparate research and specialists applying a deep understanding of particular domains;
- **malware engineers**: reverse-engineer malicious code samples and monitor their behaviour inside a controlled environment;
- **technical security engineers**: develop collection and analysis tools, analyse network traffic and undertake vulnerability research.

It should be noted that one individual might perform more than one role depending on their skillset as well as the scope, size and maturity of the intelligence function.

4.4 Skills

An intelligence function requires an additional set of skills over and above the standard technical skills found in an information security function such as intrusion detection, penetration testing, programming and incident handling.

Overall, intelligence is a blend of art and science. Intelligence personnel tend to deploy a variety of skills in disciplines such as psychology, sociology, linguistics, languages and geopolitics in comparison with information security personnel whose skills lie primarily in computer technology. With intelligence the approach to execution is experimental, exploratory and creative compared with a largely linear execution approach in information security (Jellenc (2013)).

Threat intelligence in commercial environments is a relatively immature discipline compared to its traditional role in the public sector. There are therefore fewer individuals who can demonstrate the standard 10,000-hour expertise standard, 2,500 hours being a more realistic figure. Furthermore, apart from textbooks that examine the nature of intelligence analysis, such as Heuer (1999), there is no industry standard intelligence analyst skills profile.

That said, there are some general skill characteristics that many government intelligence analysis functions have used for years and are commonly agreed within the industry. These include the ability to:

- apply classical research methods to collect, organise and evaluate data;
- apply analytical methods such as hypothesis generate-and-test, discourse analysis and behavioural profiling;
- deal with incomplete information describing ambiguous situations and apply a high level of intuition, curiosity and imagination in order to join the dots and make a judgement;
- apply critical thinking, assess a situation dispassionately and argue both sides of a position, alternate between each one as new information arises;
- maintain awareness of the influence of cognitive bias and other cognitive traps that can compromise analysis;
- speak other languages and understand the social, cultural and geopolitical characteristics of a particular country or region;
- manage human relationships effectively when undertaking HUMINT;
- communicate effectively with customers and non-experts in the field;
- remain calm under pressure;
- demonstrate the personal characteristics required for obtaining a high level of security clearance.

Team personnel may have a background in information security or may come from the private security, police, military and intelligence communities. As with all teams success ultimately lies in a balance of diverse skills and personality types.

Guidance can also be taken from existing business intelligence functions. These focus on assessing market conditions and business competitors, ie the voice of the customer. This contrasts with threat intelligence that focuses on malicious actors who can damage the organisation's ability to provide services, ie the voice of the enemy (Verisign (2013)). Nevertheless, business intelligence is a relatively mature discipline and the organisational lessons it has learned are ones that threat intelligence teams might find useful.

4.5 Professional standards

An intelligence function needs to operate at a high level of professionalism with all team members acting diligently, securely and with integrity. Three particularly important areas are summarised below.

Legal and ethical data collection

While OSINT is easily the most accessible form of information it is so abundant that it can be difficult to distinguish signal from noise. It can also be difficult to judge its authenticity. Although OSINT collection appears to be a risk-free practice it can present both legal and ethical issues that vary from country to country.

It is therefore imperative that an intelligence function demonstrates that the means by which it acquires information is both legal and ethical. Two relevant initiatives in this area are the codes of ethics promoted by:

- **OSIRA (Open Source Intelligence and Research Association)**: an international body dedicated to enhancing the knowledge and expertise of OSINT practitioners in both the public and private sector (OSIRA (2014));
- **SCIP (Strategic and Competitive Intelligence Professionals (SCIP)**: a global non-profit membership organisation for professionals involved in competitive intelligence (SCIP (2014)).

Also relevant here are the provisions of the Computer Misuse Act for securing computer material against unauthorised access or modification (Crown (2016)).

Legal admissibility and evidential weight

Organisations that intend to supply intelligence to law enforcement need to ensure the legal admissibility and evidential weight of all information collected, analysed and disseminated electronically.

BS 10008 ('*Evidential Weight and Legal Admissibility of Electronic Information*') is the British Standard that outlines best practice for the management and storage of electronic data. This includes transferring electronic data between systems and migrating paper records to digital files. It also gives guidelines for managing the availability and accessibility of any records that could be required as legal evidence (BSI (2008)).

Security and data protection

Threat intelligence functions must satisfy higher-than-normal security and data protection requirements, over and above standard requirements regarding physical security, information security and business continuity. This is particularly important for CBEST since the goal is to fuse government intelligence with commercial intelligence. Threat intelligence for the highest-level CBEST tests will require access to extremely sensitive material and personnel will therefore need to satisfy, where appropriate, government security clearance requirements.

5 Maturity

5.1 Introduction

This section discusses maturity models relating to the production and consumption of threat intelligence. An important caveat is that all models, maturity or otherwise, are approximations of reality and should be treated as such. Or, to quote the noted statistician George Box, '*Essentially, all models are wrong, but some are useful*' (Box and Draper (1987)). The extent to which an organisation needs to ascend a maturity model will depend on its particular circumstances and requirements.

5.2 Intelligence production

Maturity models for threat intelligence producers are relatively common and straightforward. They typically comprise three levels, namely:

- **initial**: developing core capability;
- **industrialised**: expanding scope and refining operations;
- **optimised**: becoming fully proactive.

It is at the industrialised level where significant changes occur, in particular:

- beginning the shift from an inward-looking reactive stance to an outward-facing proactive one;
- developing analytical specialities within the team;
- increasing the level of automation.

At the optimised level the team is proactively identifying emerging intelligence, some of which may relevant to the customer before the customer has even inquired about it.

Figure 5.1 presents a typical threat intelligence producer maturity model where each level builds on the capabilities of the one beneath (Verisign (2013)).

Figure 5.1 Threat intelligence producer maturity model

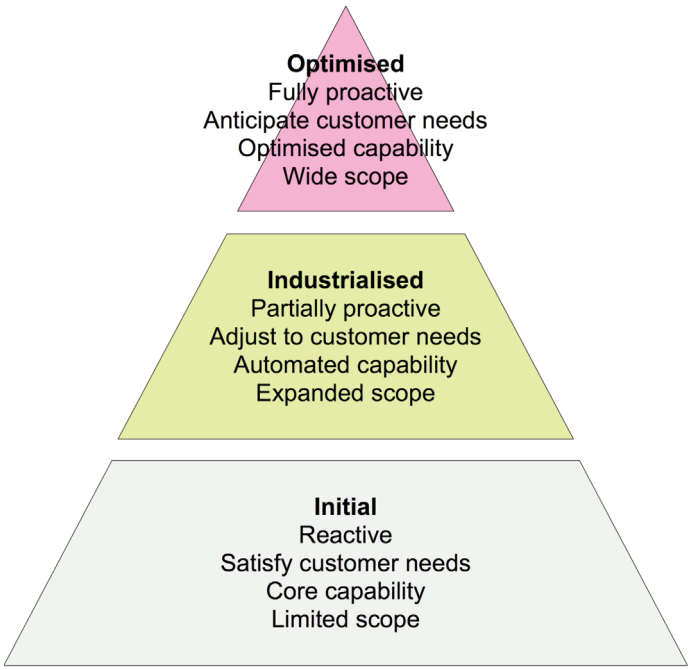

5.3 Intelligence consumption

5.3.1 Introduction

The second type of maturity model measures how well a consumer of threat intelligence exploits the intelligence products they receive. While maturity models for threat intelligence producers are relatively simple and well defined, consumer models are more complex and remain at an earlier stage of development. An initial model is therefore presented here for consideration and future refinement.

5.3.2 Business intelligence input

When devising an intelligence consumption maturity model, two useful inputs come from suppliers and analysts of general business intelligence products used by finance, marketing, risk and operations functions. This market has matured significantly over the past two decades.

The first useful input is the now familiar three-tiered classification of knowledge originally advocated by marketing professionals and subsequently adopted by their security counterparts. This distinguishes between:

- **what we know**: eg a particular threat actor is targeting our organisation;
- **what we know we don't know**: eg we think we might be susceptible to a particular class of threat but need to find out more (using threat intelligence);
- **what we don't know we don't know**: eg we are at risk in a particular area but will be completely unaware of this until it happens (or, with threat intelligence, we become aware of it).

This is summarised in **Figure 5.2**.

Figure 5.2 Three-tiered classification of knowledge

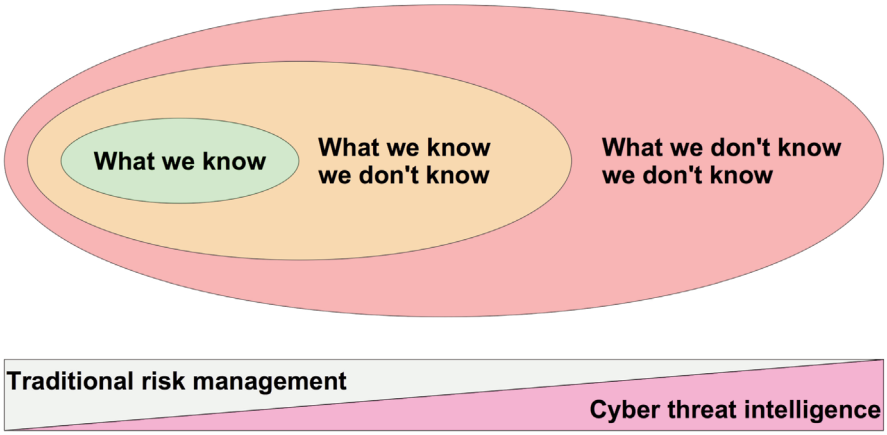

Traditional risk management works better with '*what we know*' and '*what we know we don't know*'. Cyber threat intelligence works across all three classes and particularly excels at '*what we don't know we don't know*'. This equates to the classic '*anomalous state of knowledge*' (ASK) where users of search systems are unable to precisely formulate what they need (Belkin (1980)).

A second useful input is the five-stage business intelligence maturity model (Brobst and Rarey (2001); Gartner (2013)) that distinguishes between:

- **reporting**: strategic intelligence describing what has happened, in response to pre-determined user queries;
- **diagnosis**: strategic intelligence explaining why something happened, derived by users exploiting query, drill-down and slice-and-dice functions in an *ad hoc* manner;
- **prediction**: strategic intelligence predicting what will happen, derived from predictive models created by data mining tools analysing historical data;
- **operationalisation**: tactical intelligence that indicates what is happening now, derived from a continual data feed, that is put to immediate use to support tactical decision-making in the field;
- **activation**: the increased automation of tactical and strategic decisions (intelligence events instantly triggering a machine-based decision) creates a tighter coupling between intelligence and operations, or '*what do I want to happen?*'.

In addition to the above insights from general business intelligence, another useful input into a threat intelligence consumer maturity model is the distinction between six different classes of intelligence that describe a generic cyber attack, namely *who*, *why*, *when*, *where*, *how* and *what*. These can be mapped onto the standard model of a cyber attack that comprises:

- the threat actor;
- the threat actor's TTPs (or modus operandi);
- the intended target and the incident that will arise from attacking that target.

Figure 5.3 summarises the mapping. This broadly aligns with the many cyber attack models currently under development in the security industry such as STIX (Structured Threat Information Exchange) (Barnum (2014)). Because of the increasing take-up of STIX by the threat intelligence market **Figure 5.3** shows, in italics, the equivalent STIX terminology.

Figure 5.3 Mapping intelligence classes onto a cyber attack model (STIX terminology shown in italics)

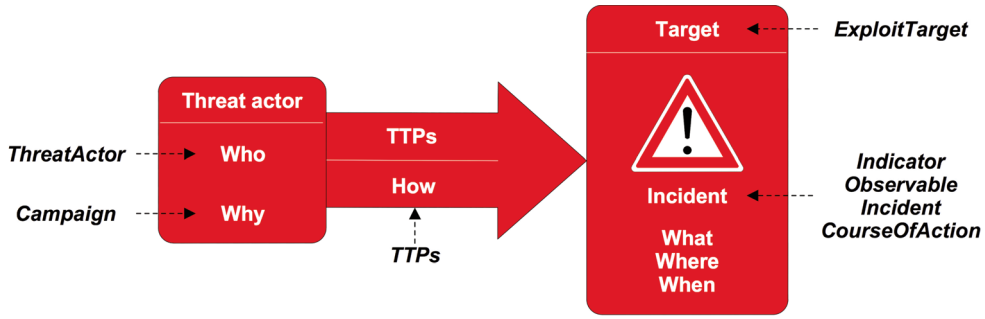

5.3.3 Consolidated model

Given the above inputs, **Figure 5.4** presents a consolidated maturity model where each level builds on the capabilities of the previous one.

Figure 5.4 Threat intelligence consumer maturity model

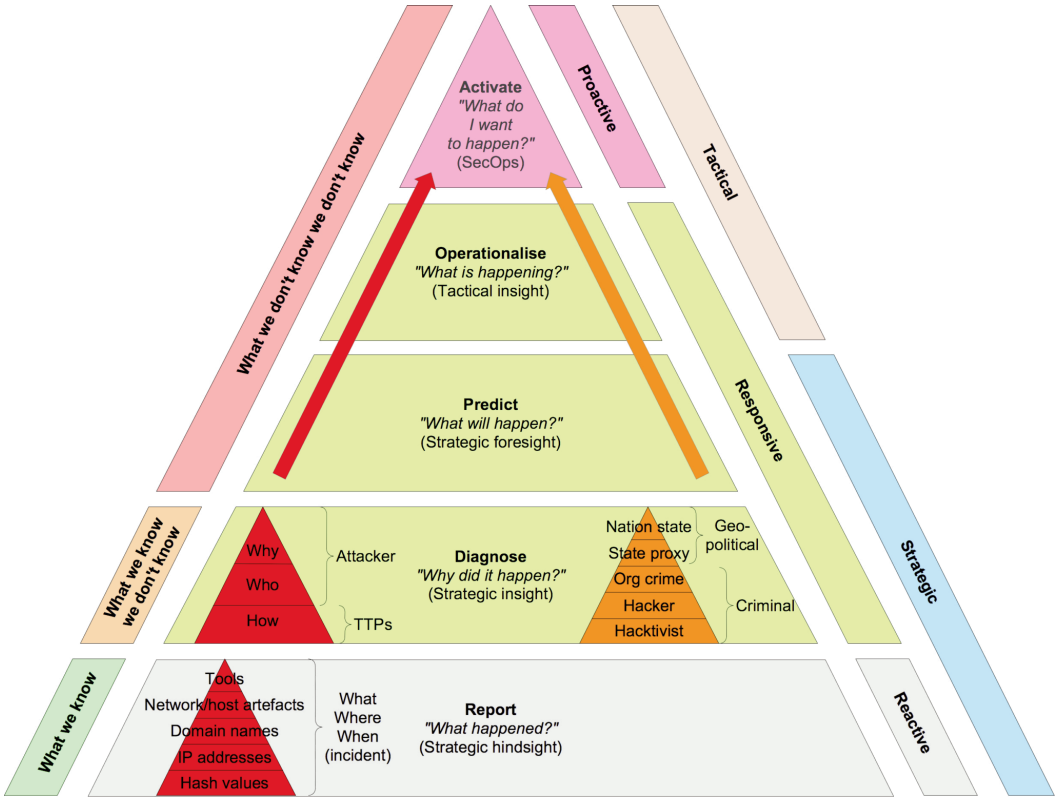

An explanation of each level of the maturity model is provided in the remaining sub-sections.

Report

At the lowest level, Report, the consumer makes use of intelligence that describes what happened, ie the details surrounding a successful attack. This comprises:

- **evidence that an attack has taken place**: ie indicators of compromise providing technical details of the incident that reside within a sub-maturity model that indicates the relative difficulty of collecting each indicator type (namely: hash values, IP addresses, domain names, network/host artefacts and tools) (Bianco (2014b));
- **evidence of how the attack has compromised the organisation**: ie details relating to the disclosure, theft, compromise or destruction of valuable organisational assets.

On the basis of this intelligence an analyst can piece together the evidence to understand the anatomy of an attack: the *what*, *where* and *when*. This is the lowest level of intelligence consumption maturity since the customer simply consults the data to determine whether they have been attacked. What is termed intelligence at this level only describes symptoms of the attack rather than information about the adversary behind the attack.

This is strategic hindsight based largely on '*what we know*' (or '*known knowns*'). It is also a reactive stance that can lead to defensive, kneejerk reactions that lack considered decision-making. In the worst case this can result in unpredictable outcomes and loss of control.

Diagnose

At the next level up, Diagnose, the consumer makes use of intelligence that diagnoses why an attack happened. Within the Diagnose level we can identify increasing intelligence maturity levels that describe:

- **TTPs**: the *how* of an attack;
- **the attacker**: *who* they are and *why* they are attacking.

The Diagnose level is a significant advance on the Report level since, rather than focusing on the symptoms or effects of an attack, the consumer derives a much greater depth of understanding about the adversary behind the attack.

Intelligence at this level provides valuable context for interpreting detected IOCs. For example, by knowing the IP addresses of all the machines running a particular data exfiltration attack (run by a botnet) an analyst can learn about the botnet infrastructure, its location, type of vulnerability exploited and style of traffic emanating from that IP. Even though the analyst may not know the *who* or *why*, they will learn more about the *how*, *where* and *when*.

The analyst can also discover how long the enemy has been inside the perimeter fence, what other IOCs may be detectable given the attacker's TTPs and what the best response should be given the nature of the adversary. In this way the consumer exhibits a more intelligence-led mindset.

Regarding the attacker, we can also identify increasing intelligence maturity levels that describe the type of attacker and the intelligence needed to defend against them, ie:

- hacktivists, hackers and organised criminals: criminal insights;
- state proxies and nation states: geopolitical insights.

The level of capability/maturity required for a threat intelligence consumer will depend on the threat profile for that consumer's organisation. Intelligence relating to *who* and *why* is becoming increasingly important as adversaries become more sophisticated. However, while many consumers will be competent at understanding the *who* and *why* of a hacktivist attacker they may be less able to understand the equivalent characteristics of a state proxy or nation state attacker.

Diagnose therefore involves strategic insight, or '*what we know we don't know*' (or '*known unknowns*'). It is also a responsive rather than reactive stance in that an element of decision-making is involved when reacting to an attack event.

Predict

At the third level up, Predict, the consumer makes use of more advanced forms of intelligence that attempt to predict what will happen. This can help focus monitoring and mitigation activities. This is a responsive stance employing strategic foresight to ascertain '*what we don't know we don't know*' ('*unknown unknowns*' or ASKs).

Operationalise

At the fourth level up, Operationalise, the consumer has the appropriate people, processes and technology in place to enable them to take intelligence indicating what is happening now and use it immediately to support tactical decision-making in the field. This is a responsive stance employing tactical insight based on what is happening now.

Activate

At the fifth and final level up, Activate, there is such a tight coupling between intelligence and operations that an intelligence event instantly triggers an automated operational response.

In the same way that the field of DevOps attempts to integrate IT development and IT operations, at this level of maturity an equivalent emerges in the form of SecOps. This sees a tighter integration between security (Intelligence, SIEM, Information Security Operations Centre, Network Operations Centre, Physical Security Operations Centre, Incident Response) and IT Operations.

A common example of this is feeding intelligence directly into automated detection systems such as SIEM, anti-virus, anti-malware, anti-spam, network monitoring, network behaviour analysis, intrusion detection and vulnerability scanners. The consumer will also feed intelligence it receives into its own data mining tools to discover new patterns, clusters and trends.

This is a highly proactive stance where the consumer's mindset moves towards '*what do I want to happen?*'

Note that offensive security is not within the scope of this maturity level. There is a myriad of both ethical and legal questions associated with any offensive philosophy. It is generally recognised that those with mature security programmes can consider counterintelligence operations (eg using use decoys to induce adversaries to reveal additional intelligence) but should leave 'hacking back' to governments and militaries (Holland (2013)).

5.3.4 Further refinement of this model

As already stated, threat intelligence consumer maturity models, including the one presented above, remain at an early stage of development and need refinement. For example:

- although there are five layers to the model, it could be sliced or structured in different ways to emphasise the categories on either side of the pyramid;
- the different maturity levels within the Diagnose level require more structure and elaboration;
- while many will agree that obtaining intelligence on what is happening right now is harder than making a prediction about the future, others will disagree and say the levels should be reversed.

Furthermore, other maturity models exist that combine producer and consumer models into a single model comprising (Gordon (2015)):

- ad hoc/monitor (lowest level of maturity);
- spectator/reactive;
- consumer/proactive;
- producer/adaptive;
- mission partner/transformative;
- mission integrated/influential (highest level of maturity).

5.3.5 Intelligence-led cyber resilience

By operating on the upper levels of the threat intelligence consumer maturity model described above, and making appropriate enhancements to their organisation, processes and technology, organisations can reap the benefits of intelligence-led cyber resilience.

In order to appreciate what this means, it is important to understand how cyber resilience differs to traditional information security practice. The latter:

- focuses on the perimeter and looks inward;
- is founded on computer technology skills;
- investigates, remedies and complies;
- adopts a probabilistic red-amber-green approach to risk;
- reports on the past.

In contrast, cyber resilience:

- is outward-focused;
- is founded on skills such as psychology, sociology, linguistics, languages and geopolitics in addition to technical skills;
- experiments, explores and creates;
- adopts a goal-driven, uncertainty reduction approach to risk;
- reports, diagnoses, predicts, executes and influences.

If we consider all the activities involved in containing cyber threats, information security focuses on a subset of these, namely prevention, detection, response and recovery. Even within that focus the coverage is not total: disruption of the cyber kill chain (Hutchins, Cloppert and Amin (2011)) before the command and control stage is uncommon and post-incident response and recovery are more reactive than responsive.

Intelligence-led cyber resilience, on the other hand, is more than rebranded information security. By exploiting threat intelligence (who is out to get them and why, and how, when and where they are likely to do it) in every step of the threat containment cycle, and continuously learning and adapting defences in response to the nature of the threat, it enables an organisation to become proactive in managing the vast array of new and evolving advanced cyber threats. Critically, the attacker's kill chain can be disrupted at a much earlier stage, ie delivery, and the intelligence about the disrupted attacker and their TTPs used, in a predictive manner, to improve the capability of preventative and detective countermeasures (Chuvakin (2014b)).

Figure 5.5 summarise the above discussion. As well as showing the scope of information security and threat intelligence, it also overlays the CESG's '*10 Steps*' cyber security measures (CESG (2012)) and the '*Cyber Essentials*' scheme backed by the Department for Business, Innovation and Skills (BIS (2014)) to provide further context.

Figure 5.5 Intelligence-led cyber resilience

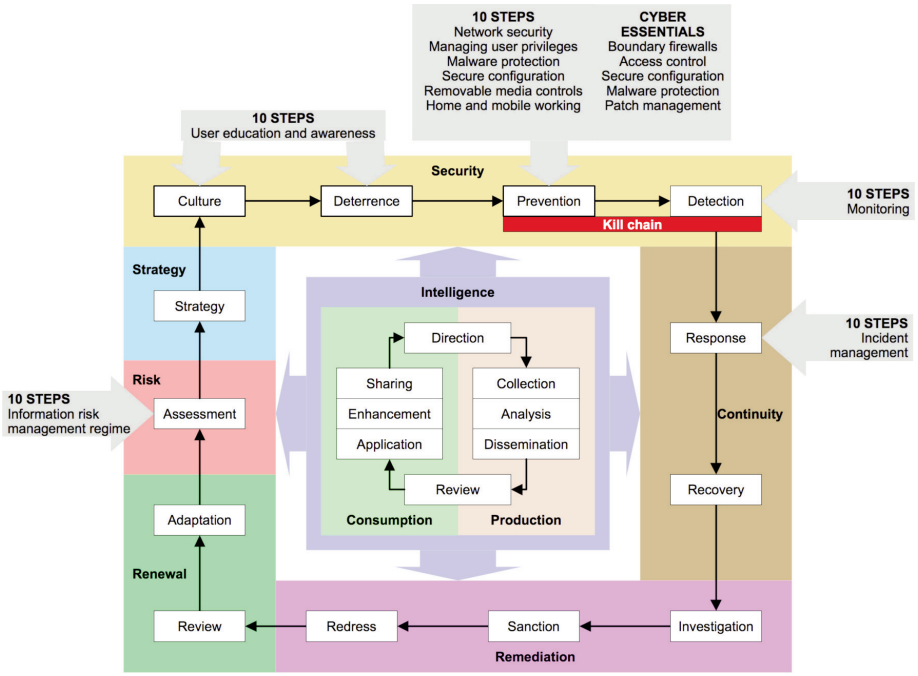

5.3.6 Implications for risk management and business strategy

Operating on the upper levels of the threat intelligence consumer maturity model means ensuring that cyber risk is factored into the organisation's risk appetite and risk management framework. Advising policy teams about potential future threats means changes can be made proactively to controls across the '*three lines of defence*', namely business operations, oversight functions and independent assurance (FSA (2010)).

Intelligence can go even further by challenging entrenched methods of defining and managing risk. Traditional risk management involves the attempt to enumerate, up-front, all the things that could go wrong and then assigning probabilities and other numerical values. As well as confusing probability with severity, and generating a huge number of often irrelevant or ambiguous red-amber-green risk register entries, this approach struggles to cope with today's turbulent, unpredictable and interconnected operating environment where threat actors reside.

The intelligence-led approach sheds more light on threat actors, their motivations and their TTPs. By reducing uncertainty about threats it also reduces risk. This is because risk, at its core, is all about uncertainty, or the degree to which the chance of an organisation achieving its goals depends on things it cannot control, predict or understand (Slater (2013)). Risk management can exploit the insights from threat intelligence to understand where the risks, ie the greatest uncertainties, really lie. Armed with this insight it can then determine what countermeasures are really needed and which ones are not worth the financial outlay.

All too often what is called risk management is simply compliance. Many who claim to be risk averse are in reality blame averse. By making it possible, through improved situational awareness, to better understand, predict and ultimately control uncertainty, threat intelligence can encourage a move away from compliant behaviour and acting out of ignorance to actively managing the right risks.

More broadly still, intelligence-led decision-making and continuous learning can be something that corporate boards can adopt. Having a clear view on cyber threats will help ensure they make better-informed strategic business decisions. A good example from the public sector is the London 2012 National Olympics Co-ordination Centre (NOCC) which, as a mature intelligence capability, was integrated into the governance and decision making structures of UK policing (ACPO (2011)).

In this way, traditional information security can be transformed from an isolated expert function to an influencer of organisational culture (Hult and Sivanesan (2013)). This in turn will enable the organisation to embrace digital business while protecting itself from targeted cyber attack.

6 Conclusions

6.1 Introduction

This section presents a summary of the key points and discusses future developments in cyber threat intelligence.

6.2 Summary

Industry and analyst commentators all agree that the time has come to invest resources into understanding and countering professional cyber attackers. An approach based on threat intelligence complements the existing preoccupation with vulnerability- and asset-centric security (Chuvakin (2014c)).

In comparison to its counterpart in the government and law enforcement sector, cyber threat intelligence in the commercial environment remains a relatively immature discipline. For example, there is still no formally agreed definition of what constitutes cyber threat intelligence. It is also the subject of much vendor hype. That said, in the financial sector there is a high level of information sharing (eg FS-ISAC and CISP) and immaturity remains at the strategic level rather than the operational or tactical level.

Despite its relative immaturity, commercial intelligence is in a good position to exploit several decades of government intelligence best practice. It is a positive sign, for example, that intelligence suppliers have broadly settled on a common intelligence life cycle model.

Within the intelligence life cycle it is the input and output phases, ie collection and dissemination, where the greatest issues lie such as ethical collection, trustworthy sharing and the role and relative merits of machine-readable threat intelligence. Intelligence sharing remains at a relatively low level of maturity, much of it taking place at the tactical or operational level rather than at the strategic level. A range of relevant MRTI standards have emerged with different degrees of adoption, functionality and effectiveness. Of these, STIX appears to be gaining the widest support in the industry (Bryant (2013); Gartner (2014c); Tripwire (2015)). That said, a great deal of threat intelligence is also shared using simpler formats such as plain text, PDF, CSV, XML, PCAP, YARA Editor, SQL and JSON.

Defining the requisite roles and skills in a threat intelligence team also remains at a relatively low level of maturity. Organisations need to appreciate that cyber threat intelligence analysis differs from traditional information security, in particular the way it requires skills in the humanities rather than a sole focus on computer science.

While maturity models for threat intelligence producers are relatively simple and well defined, consumer models are more complex and remain at an earlier stage of development. However, when properly structured and communicated, such models can be compelling means of conveying the benefits of intelligence-led cyber resilience to systemically important financial institutions. Not only will an organisation proactively manage an array of new and evolving advanced cyber threats but there is also the potential for improving risk management and high-level business strategy.

As for traditional information security, adopting an intelligence-led approach has the potential to transform it from a largely reactive function that investigates, remedies, complies and reports to a more responsive function that also diagnoses, predicts, executes and influences.

6.3 Future developments

In terms of the classic hype cycle defined by Gartner (2014b), commercial threat intelligence is in the '*trough of disillusionment*'. Although recent surveys show that organisations are steadily increasing their use of threat intelligence to better detect and respond to cyber attack, many are critical of the reliability of this intelligence as well as its ability to be actionable (Ponemon (2015)). Having built a threat intelligence capability, or procured a threat intelligence platform from a third-party service provider, innovators and early adopters are now asking what they can usefully do with them.

While TIP vendors and security analysts are keen to promote the virtues of MRTI, its main purpose is to make the data import process easier. Although some TIPs can perform elementary analysis of MRTI, it remains a data feed that supports a higher-level geopolitical threat intelligence narrative but does not replace it. Furthermore, MRTI focuses on the sharing of tactical indicators of compromise (or '*known knowns*'). Although these are valuable they occupy just one part of the intelligence spectrum and their use alone is not indicative of a high level of intelligence maturity.

Many analysts based in Security Operations Centres (SOCs) are therefore frustrated with the limitations of generic threat intelligence data feeds based on MRTI. Others who originally welcomed an increase in information via threat intelligence platforms now want the fire hose to be turned down and given only what they really need to know. Not only do they want more tailored intelligence but they also want it more closely integrated with their security infrastructure so that responses can take place inside that interface rather than adjacent to it.

To this end, the virtues of a hybrid human-machine analysis are highlighted in a recent industry analyst report that states: '*A benefit of carbon-based analysis is that people understand context and can provide a deeper understanding of critical issues. In this scenario, automation is used to enable the analyst since human analysts cannot scale*' (Gartner (2015)).

As commercial threat intelligence moves out of the trough of disillusionment and onto a more enlightened and productive growth path, vendor offerings will consolidate and evolve and an early majority of organisations will take up the technology. It is at this point when providers and users of threat intelligence services will see increased benefits from a more robust, holistic and tailored approach to generating threat intelligence which will play a far more strategic role. This is because generic threat intelligence can only achieve so much. While it is helping to improve defensive capability by providing a better understanding of threats and threat actors, organisations need to develop a stronger sense of situational awareness.

Situational awareness is classically defined as '*the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and a projection of their status in the near future*' (Endsley (1995)). Cyber situational awareness brings together all the information that an organisation possesses about itself, its partners, its attack surface and its threat environment in a form that is tailored to that specific organisation. This will help it to understand what it needs to do right now and also how to adapt its security countermeasures in the longer term.

A number of threat intelligence providers are playing into this evolving market, the key points of differentiation being the extent of OSINT coverage (threat actors and targets across the visible and dark Web), efficiency of search algorithms, involvement of expert analysts and curators, reduction of false positives, tailoring of intelligence, latency of alerts, ease of deployment and quality of innovation. Their aim is to provide comprehensive intelligence at the tactical, operational and strategic levels that have been identified by industry analysts (Gartner (2015)).

The above discussion is summarised in **Figure 6.1**.

Figure 6.1 Current position and future evolution of commercial threat intelligence

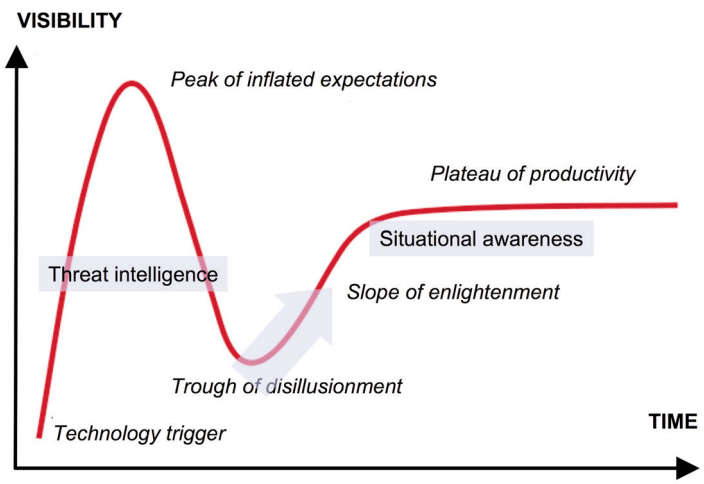

Threat intelligence is a moving target and this report will only ever be a snapshot of the current state of the art. As CBEST continues to evolve the following areas, relating to the concept of operations for cyber threat intelligence, should therefore be explored in further detail:

- further development of a data collection code of ethics;
- an industry standard for machine-readable threat intelligence;
- an industry standard trust model for intelligence sharing.

References

451 Research (2014), 'Threat intelligence', 451 Research, LLC.

ACPO (2011), 'National coordination — the National Olympic Coordination Centre (NOCC)', available at www.acpo.police.uk/ACPOBusinessAreas/OLYMPICS/Commandcontrolandcoordination.aspx. Association of Chief Police Officers.

AlienVault (2013), 'Welcome to the AlienVault Open Threat Exchange', available at www.alienvault.com/open-threat-exchange. AlienVault, Inc.

Barnum, S (2014), 'Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX™)', Version 1.1, Revision 1. The MITRE Corporation.

Beck, K *et al* (2001), 'Manifesto for Agile software development', available at http://agilemanifesto.org. Agile Alliance.

Belkin, N (1980), 'Anomalous states of knowledge as a basis for information retrieval', *The Canadian Journal of Information and Library Science*, Vol. 5, pages 133–43. University of Toronto Press.

Bianco, D (2014a), 'Use of the term 'Intelligence' in the RSA 2014 Expo', available at http://detect-respond.blogspot.co.uk/2014/03/use-of-term-intelligence-at-rsa.html.

Bianco, D (2014b), 'The pyramid of pain', available at http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html.

BIS (2014), 'Cyber essentials scheme', Department for Business, Innovation and Skills.

Box, G and Draper, N (1987), *Empirical model-building and response surfaces*, John Wiley & Sons, Inc.

BrightPoint (2015), 'Sentinel product overview', available at www.brightpointsecurity.com/products/products-overview/. BrightPoint Security.

Bryant, J (2013), 'New MAPP initiatives', available at http://blogs.technet.com/b/bluehat/archive/2013/07/29/new-mapp-initiatives.aspx. Microsoft Corporation.

Brobst, S and Rarey, J (2001), 'The five stages of an active data warehouse evolution'. Teradata Corporation.

Bryner, J (2013), 'Python tools for ioc (indicator of compromise) handling', available at https://github.com/jeffbryner/pyioc. GitHub, Inc.

BSI (2008), 'BS 10008:2008 evidential weight and legal admissibility of electronic information'. British Standards Institution.

Cabinet Office (2011), 'The UK cyber security strategy: protecting and promoting the UK in a digital world'. Crown Copyright.

CBEST (2016a), 'An Introduction to Cyber Threat Modelling', Bank of England.

CBEST (2016b), 'CBEST Services Assessment Guide', Bank of England.

CESG (2012), '10 steps to cyber security'. Crown Copyright.

Chuvakin, A (2014a), 'On threat intelligence sources', available at http://blogs.gartner.com/anton-chuvakin/2014/02/26/on-threat-intelligence-sources/#comments. Gartner, Inc.

Chuvakin, A (2014b), 'Delving into threat actor profiles', available at http://blogs.gartner.com/anton-chuvakin/2014/03/14/delving-into-threat-actor-profiles/?utm_medium=twitter. Gartner, Inc.

Chuvakin, A (2014c), 'How to collect, refine, utilize and create threat intelligence', available at www.gartner.com/doc/2738618/collect-refine-utilize-create-threat. Gartner, Inc.

CIF Project (2009), 'Collective Intelligence Framework', available at https://code.google.com/p/collective-intelligence-framework/. REN-ISAC, Indiana University, Internet2 and The National Science Foundation.

Clark, M (1955), 'Intelligence activities'. Interim technical report to Congress. Commission on Organization of the Executive Branch of the Government [the Hoover Commission].

Connolly, J (2013), 'The Trusted Automated eXchange of Indicator Information (TAXII)'. The MITRE Corporation.

CREST (2013), 'Cyber security incident response guide'. Version 1. CREST (GB).

Crown (2016), 'Computer Misuse Act 1990', available at www.legislation.gov.uk/ukpga/1990/18/contents. Crown copyright.

Daly, M (2009), 'The Advanced Persistent Threat (or informationized force operations)', LISA 2009 Conference. Available at www.usenix.org/legacy/event/lisa09/tech/slides/daly.pdf. USENIX.

Danyliw, R, Meijer, J and Demchenko, Y (2007), 'The Incident Object Description Exchange Format', available at www.ietf.org/rfc/rfc5070.txt. Internet Society.

Dark Reading (2013), 'Advanced Persistent Threats: the new reality', available at www.darkreading.com/vulnerabilities---threats/advanced-persistent-threats-the-new-reality/d/d-id/1139716?. UBM Tech.

Dark Reading (2015), 'Threat Intelligence Platforms: the next 'must-have' for harried security operations teams', available at www.darkreading.com/threat-intelligence-platforms-the-next-must-have-for-harried-security-operations-teams/d/d-id/1320671. UBM Tech.

Deming, W (1986), *Out of the crisis*. MIT Press.

Digital Shadows (2015), 'The dangers of groupthink', available at www.digitalshadows.com/blog-and-research/the-dangers-of-groupthink/. Digital Shadows, Inc.

Dorrington, P (2004), 'A strategic approach to reducing losses'. SAS Institute, Inc.

Endsley, M (1995), 'Towards a theory of situation awareness in dynamic systems', *Human Factors*, Vol. 37(1), pages 32–64. Human Factors and Ergonomics Society.

Farnham, G (2013), 'Tools and standards for cyber threat intelligence projects'. The SANS Institute.

FSA (2010), 'Enhancing frameworks in the standardised approach to operational risk — guidance consultation'.

Gartner (2013), 'Extend your portfolio of analytics capabilities'. Gartner, Inc.

Gartner (2014a), 'Technology overview for machine-readable threat intelligence'. Gartner, Inc.

Gartner (2014b), 'Research methodologies: hype cycle', available at www.gartner.com/technology/research/methodologies/hype-cycle.jsp. Gartner, Inc.

Gartner (2014c), 'Technology overview for threat intelligence platforms'. Gartner, Inc.

Gartner (2015), 'Vendor landscape: S&R pros turn to cyberthreat intelligence providers for help'. Gartner, Inc.

Gordon, M (2015), 'The evolution of cyber security'. Lockheed Martin Corporation.

Heuer, R (1999), 'Psychology of intelligence analysis'. Center for the Study of Intelligence, CIA.

Hickman, F, Killin, J, Land, L, Mulhall, T, Porter, D and Taylor, R (1989), 'Analysis for knowledge-based systems'. Ellis Horwood Publishing.

Holland, R (2013), 'Five steps to building an effective threat intelligence capability'. Forrester Research, Inc.

Hult, F and Sivanesan, G (2013), 'What good cyber resilience looks like', *Journal of Business Continuity and Emergency Planning*, Vol. 7, No. 2. Henry Stewart Publications.

Hutchins, E, Cloppert, M and Amin, R (2011), 'Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains', in proceedings of 6th Annual International Conference on Information Warfare and Security. Lockheed Martin Corporation.

IOC Bucket (2014), 'About iocbucket.com', available at http://iocbucket.com/about.

INSA (2013), 'Operational levels of cyber intelligence'. Intelligence and National Security Alliance.

Jellenc, E (2013), Unpublished research materials. VeriSign-iDefense, Inc.

Kapuria, S (2011), 'IT threat intelligence to anticipate, stop and counteract targeted attacks'. Symantec Corporation.
KPMG (2013), 'Cyber threat intelligence and the lessons from law enforcement'. KPMG International Cooperative.

Lawson, C and McMillan, R (2014), 'Technology overview for machine-readable threat intelligence'. Gartner, Inc.

McMillan, R and Pratap, K (2014), 'Market guide for security threat intelligence services'. Gartner, Inc.

Mitre (2014a), 'Cyber Observable eXpression: a structured language for cyber observables', available at http://cybox.mitre.org. The MITRE Corporation.

Mitre (2014b), 'Structured Threat Information eXpression: a structured language for cyber threat intelligence information', available at http://stix.mitre.org. The MITRE Corporation.

Mitre (2014c), 'Threat-based defense', available at www.mitre.org/capabilities/cybersecurity/threat-based-defense. The MITRE Corporation.

Mitre (2014d), 'Trusted Automated eXchange of Indicator Information: enabling cyber threat information exchange', available at http://taxii.mitre.org. The MITRE Corporation.

Mitre (2014e), 'MAEC Language — Version 4.1', available at http://maec.mitre.org/index.html. The MITRE Corporation.

Mitre (2014f), 'Common vulnerabilities and exposures: the standard for information security vulnerability names', available at http://cve.mitre.org/index.html. The MITRE Corporation.

Moriarty, K (2012), 'Real-time Inter-network Defense', available at http://datatracker.ietf.org/doc/rfc6545/. Internet Society.

Moriarty, K (2013), 'Implementations on incident object description exchange format', available at http://siis.realmv6.org/implementations/. Workshop on Security Incident Information Sharing.

MS-ISAC (2015), 'Traffic Light Protocol (TLP) matrix and frequently asked questions', available at https://msisac.cisecurity.org/resources/tlp/. Center for Internet Security.

NCIS (2000), 'The National Intelligence Model'. National Criminal Intelligence Service.

OpenIOC (2014), 'OpenIOC: an open framework for sharing threat intelligence', available at www.openioc.org/.

OSIRA (2014), 'About OSIRA', available at http://osira.net/about-osira/. Open Source Intelligence and Research Association.

Paterva (2014), 'Maltego', available at https://www.paterva.com/web6/products/maltego.php. Paterva (Pty) Ltd.

Payments Council (2014), 'Cyber threat intelligence: an analysis of an intelligence-led, threat-centric, approach to cyber security strategy within the UK banking and payment services sector'. Payments Council.

Ponemon (2015), 'The Importance of cyber threat intelligence to a strong security posture'. Ponemon Institute, LLC.

Ragan, S (2014), 'CSO's guide to advanced persistent threats', available at www.csoonline.com/article/747158/cso-s-guide-to-advanced-persistent-threats?source=CSONLE_nlt_secresmetrics_2014-02-03. CXO Media, Inc.

Recorded Future (2015), 'Cyber threat intelligence', available at www.recordedfuture.com/cyber-threat-intelligence/.
Recorded Future, Inc.

REN-ISAC (2014), 'Collective intelligence framework', available at https://code.google.com/p/collective-intelligence-framework/. Research and Education Networking Information Sharing and Analysis Center.

SCIP (2014), 'About the Strategic and Competitive Intelligence Professionals', available at www.scip.org/AboutSCIP.php. Strategic and Competitive Intelligence Professionals.

Schneier, B (2011), 'Advanced persistent threat', *Crypto-Gram Newsletter*, 15 November, available at www.schneier.com/crypto-gram-1111.html. Schneier on Security.

Schneier, B (2014), 'Computer network exploitation vs. computer network attack', available at www.schneier.com/blog/archives/2014/03/computer_networ.html. Schneier on Security.

Shannon, C (1948), 'A mathematical theory of communication', *Bell System Technical Journal*, Volumes 27/28, Issues 3/4, pages 379–423/623–56. Alcatel-Lucent.

Slater, D (2013), 'Risk: the country of the blind? Managing uncertainty, concerns and constraints in the real world'. Cambrensis Ltd.

Takahashi, T (2013), 'IODEF-extension for structured cybersecurity information', available at http://tools.ietf.org/html/draft-ietf-mile-sci-08. Internet Society.

Techopedia (2014), 'Computer Network Exploitation (CNE)', available at www.techopedia.com/definition/27909/computer-network-exploitation-cne. Janalta Interactive, Inc.

Tripwire (2015), 'Why we should care about STIX and TAXII', available at www.tripwire.com/state-of-security/security-data-protection/cyber-security/why-we-should-care-about-stix-taxii/. Tripwire, Inc.

US Army (2010), 'Field Manual 2–0: intelligence'. Department of the Army.

US Army (2014), 'Field Manual 3–38: cyber electromagnetic activities'. Department of the Army.

US-CERT (2015), 'Traffic Light Protocol (TLP) matrix and frequently asked questions', available at www.us-cert.gov/tlp. Department of Homeland Security.

VERIS (2014), 'VERIS Community: a resource for learning about the VERIS framework', available at www.veriscommunity.net/doku.php. Verizon.

Verisign (2013), 'Establishing a formal cyber intelligence capability'. VeriSign-iDefense, Inc.

Walter, J (2013), 'Operation Troy: OpenIOC release', available at http://blogs.mcafee.com/executive-perspectives/operation-troy-openioc-release. McAfee, Inc.