Cyber Threat Intelligence in Government: A Guide for Decision Makers & Analysts

Version 2.0
March 2019

Digital, Data & Technology

CYBER SECURITY PROGRAMME

Executive Summary

The UK government, led by the National Cyber Security Centre (NCSC), is currently maturing its cyber security capability. A significant number of government departments have delivered or are in the process of delivering cyber security improvement programmes to mature their cyber security posture.

In line with the current opinions on best practice, an increasing number of departments are taking a threat led approach to cyber security, seeking to understand who in the threat landscape means to do them harm, and using that information to improve their ability to prevent or mitigate attacks. To achieve this, departments are seeking to create and deliver Cyber Threat Intelligence (CTI) capabilities.

CTI can be essential for investigation and incident response processes. Understanding the tactics, techniques and procedures used by threat actors enables their presence to be anticipated on a network. Also, having an understanding of the intent and capability of threat actors allows CTI analysts and leads to react appropriately in the event of a breach and mitigate its impact as much as possible. CTI is crucial for allowing CSOCs and network administrators to understand which attacks could be the most likely or have most impact, and therefore allow them to prepare accordingly.

Many departments that are looking to augment their newly developed and existing cyber security capabilities with CTI are struggling to define what the capability should look like and how best to drive value from it. This paper provides an overview for UK government departments and organisations (collectively termed as, “departments”) on how to deliver a CTI capability. This covers how to set a CTI strategy, what a CTI function should deliver, how that content should be delivered and how to effectively resource a capability. These conclusions have been driven by collaborating with large government departments, the NCSC and commercial partners. Whilst most departments consulted as part of this exercise had some CTI capability, the level of maturity and coverage varies significantly, and this guide supports both newcomers to threat intelligence and existing departments who may have nascent or incomplete capabilities.

The key takeaways from this guide are as follows:

1. CTI is a supporting capability for cyber security defences; it does not replace a dedicated protective monitoring capability. Prior to investment in CTI, departments should uplift existing capability to the Minimum Cyber Security Standard (1).

2. CTI is a broad field, with both immature customers and immature product and service vendors. To avoid wasteful spending, departments should create a CTI strategy, and pilot open source tools to better inform requirements.

3. Resourcing a CTI function is particularly challenging. Individuals with desirable skills in this field are difficult to find, particularly within civil service pay bands. Departments should attempt to resource from cyber, intelligence, and non-technical backgrounds to build a diverse team.

Collaboration is necessary to create an effective CTI community. The threats faced by different departments are often extremely similar, and there is value in sharing not just intelligence but also lessons learnt and best practice. This can likewise be applied across industry, with threat sharing groups allowing competitors to share threat intelligence information for the benefit of all.

Table of Contents

List of Figures

List of Tables

1 Introduction

1.1 Background

The UK Government is currently investing heavily in Cyber Security, with departments now being mandated to deliver a minimum standard of cyber security (1). Departmental maturity is varied, with departments aware of and looking to deliver to this standard and above.[1]

Departments are exploring Cyber Threat Intelligence (CTI), and current best practice advocates a "threat-led" approach. This involves understanding who a departments threat actors may be, their motivation and capability, and the subsequently distributing actionable intelligence on them.

The broad coverage of CTI has driven the creation of a wealth of products and services from vendors. However, departments may possess limited CTI knowledge and experience, hindering their ability to purchase appropriate tooling and preventing the creation of an environment that aligns with best practice.

Departments with a CTI capability possess a range of scopes, budgets and features and their ability to deliver in line with best practice varies. To tackle this problem, this project has engaged with the National Cyber Security Centre (NCSC), government and commercial partners to create this guide, with a view to providing a government focussed view of CTI capability development.

1.2 Target Audience

This guide is aimed at individuals who oversee or deliver threat intelligence capability to a department. This document provides a roadmap to delivering a CTI capability and an overview of the activities, deliverables and technologies required. Where necessary technical detail is included. An overview and roadmap for delivering a CTI function containing the key conclusions from the paper can be found in section 2. In depth discussion of the considerations to take when delivering a CTI capability, either maturing specific areas or delivering from scratch can be found in the remaining sections.

1.2.1 Getting the Most Out of This Paper

This paper provides a detailed description of an end-to-end CTI capability, at a basic but competent level. This paper does not provide guidance on developing an already mature capability further, or on developing cutting edge capability.

The authors recommend that individuals responsible for delivering the CTI capability, read sections 3, 4 and 10 to support their initial decision-making.

If you are accountable for an existing CTI function, and you wish to understand how this can be matured to deliver greater value, consider reading sections 4 and 9.

Whether you are a decision maker, team lead or analyst and you wish to explore each area of CTI in detail, this is divided according to the threat intelligence lifecycle in the following sections:

[1] For simplicity, the terms "government" and "organisation" will collectively be referred to as "department" for the remainder of this document.

- Section 4: Direction
- Section 5: Collection
- Section 6: Processing
- Section 7: Analysis
- Section 8: Dissemination

With additional sections on continuous improvement and organisation (including resourcing):

- Section 9: Continuous Improvement
- Section 10: Organisation

Whilst each section stands on its own, specific concepts and technologies are introduced the first time they are referenced.

1.3 NCSP Funded Publications

This guide has been authored by the Home Office Cyber Security Programme. The authors of this guide are grateful to the Cabinet Office for providing funding for this project from the National Cyber Security Programme.

This guide is one of three documents being published as part of NCSP funded projects, each of which are mutually complementary. They are as follows:

- Cyber Threat Intelligence – A Guide for Decision Makers and Analysts (this paper).
- Detecting the Unknown – A Guide to Threat Hunting
- Controlling Your Footprint – A Guide to Digital Risk and Intelligence

1.3.1 Cyber Threat Intelligence

Cyber Threat Intelligence is the process of collecting, processing and analysing information regarding adversaries in cyberspace, in order to disseminate actionable threat intelligence, by understanding adversaries' motivations, capability, and modus operandi, to inform cyber security mitigation measures.

This paper provides an overview for UK government departments and organisations on how to deliver a CTI capability. This covers how to set a CTI strategy, what a CTI function should deliver, how that content should be delivered and how to effectively resource a capability.

1.3.2 Threat Hunting

Threat Hunting is the proactive, iterative and human-centric identification of cyber threats that are internal to an IT network and have evaded existing security controls.

This guide, produced via a literary review and engagements with public and private sector organisations, provides recommendations for Security Operations Centres (SOCs), government departments, and across HM Government, to detect unknown malicious activity through development of Threat Hunting as both capability and a profession.

1.3.3 Digital Risk and Intelligence

Digital Risk and Intelligence (DR&I) is the process of monitoring, detecting and remediating threats within the public domain, through the control of an organisation's digital footprint.

This paper will provide recommendations as to how government departments can better understand and control their digital footprint through developing and maturing Digital Risk and Intelligence capabilities. The recommendations in this paper have been broken down into three levels: Threat Intelligence team - the quick, more easily implemented, short term recommendations; Government Department - the medium-term recommendations that will bolster the capability of the threat intelligence teams; Cross-Government functions - the longer-term recommendations that will allow government departments to better protect their digital footprints for the future.

1.3.4 Full capability adoption

We recognise that each of these publications recommends a dedicated team and investment for each capability, and in an ideal world each would stand alone with discrete objectives. However, it is recognised that there are synergies between each which can be utilised to facilitate a more streamlined capability.

Each of the areas covered by these papers cover different elements of the MITRE Cyber Attack Lifecycle (2):

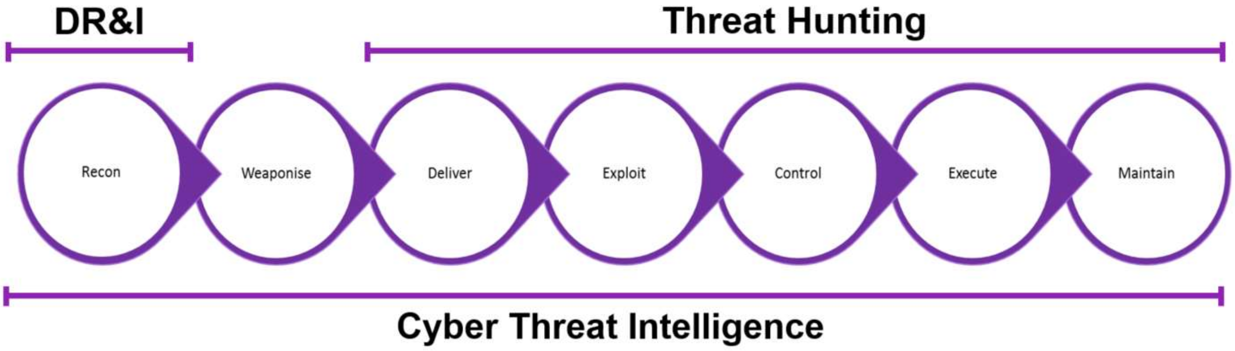

Figure 1 – Capability Scope Comparison

Clearly there are overlaps in the focus of the distinct functions, for example in the reconnaissance phase whilst CTI and DR&I have different objectives, there is a similarity in content and focus. Depending on business requirements there may be other areas where further integration can be of benefit, but fundamentally adoption of each capability needs to be on the basis of its value, based on cost vs business benefit.

If adopting all three capabilities, we recommend the following considerations be made:

- All three capabilities are subservient to each of the capabilities outlined in the minimum cyber security standard. If the minimum standard is not met, it is highly likely that investment in those areas will be more beneficial than these capabilities.
- Establishing a mature capability in all three areas represents a significant business investment. Particularly in the public sector, scrutiny of this investment will be high and we recommend that the business case for each ensures that there is genuine value for money in each area. Across all of the organisations we collaborated with none of them had made a commitment to maturing all three capabilities beyond a nascent state.

- Access to data and visibility of data is critical to all functions, both internally and externally. We would recommend that the specific pre-requisites for data access in your organisation are understood prior to investment – other organisations consulted have made significant investments, and subsequently failed to realise the benefit due to a lack of data access.

- A nascent CTI and Threat Hunting capability should grow together, if they have complementary requirements. A mature threat hunting capability that has no CTI capability to feed it intelligence will be limited, and likewise a CTI capability feeding information to a CSOC with no threat hunters is likewise limited in value.

For further details on each of these points, please refer to each of the papers specifically.

1.4 Contributors

The authors would like to thank all parties who have provided their time and insight in contribution to this guide. Information was gathered predominantly through workshops with a structured agenda; with additional context from cross government working groups, informal events and pre-existing papers. Whilst a full list of named collaborators is not included for privacy reasons, a list of organisations consulted can be found in the Appendix. If you require further information or wish to speak to anyone involved in the creation of this paper, please contact the authors directly.

If you wish to contact the authors with comments or questions, please email:

ctiplaybook@homeoffice.gov.uk

A full bibliography of referenced sources is available in Appendix 12.6.

2 Creating a CTI Function

2.1 Do I need Threat Intelligence?

Departments are not (currently) mandated to collect, integrate or analyse CTI related to their networks or data. The NCSC minimum cyber security standard presents a minimum set of measures (1), including:

> “As a minimum, Departments shall capture events that **could** be combined with common threat intelligence sources e.g. Cyber Security Information Sharing Partnership (CiSP) to detect known threats.”

This states that departments shall collect data (from their own infrastructure and devices) but that integration with CTI is not mandatory. However, the standard also says that departments should look to exceed the standard wherever possible, so whilst it is not currently mandated that organisations consume threat intelligence content it is recommended that they do so.

Irrespective of the minimum standard, threat intelligence is likely to be valuable to departments who already have a mature cyber security posture. Whilst the appropriate pre-requisites for a CTI capability to deliver value will vary from department to department, we would recommend as a catch-all statement that organisations only consider significant (i.e. more than 5% of security budget) investment in CTI after they have met or are on a realistic roadmap to achieving *all* of the 10 sections of the minimum cyber security standard. The budget allocated to the CTI investment should be based upon the department’s assets and risk appetite.

Assuming that has already been achieved, the final test is the following:

- Can your department make use of actionable threat intelligence provided to it?

If your department receives intelligence, can it use that intelligence in a way that enhances the defences of the department, whether technical or procedural? For example, if your CSOC is given information but cannot process it in sufficiently fast timescales for that information to be useful to it, there is no point in providing that information in the first place. Alternatively, if a CTI function provides information in the form of reports to system owners who are not empowered to act on the contents of the report to defend their systems, there is limited value in creating the report.

If all the above pre-requisites can be met, CTI can be an extremely valuable enabler to any department. The remainder of this section will cover some indicative use cases for CTI and the basic elements of creating and maturing a CTI function over the first year of activity. The remaining sections of the paper will discuss the detailed process of delivering a basic but competent CTI function, and the associated challenges which must be overcome to do so.

2.2 CTI Use Cases

CTI can be used in several ways. Before a department makes a relatively expensive and long-term decision on how to improve its CTI capability, it should understand the use cases for CTI. Seven key use cases are identified below (adapted from (3), (4), (5)):

Use Case	Objective	Intelligence Required
Validate Alarms/Events	Validate alarms/events and decide which to escalate to the incident response team for remediation.	Threat data: data connecting individual indicators, threat actors, techniques, etc.
Enhance Automated Response	Automate the triage process of investigations by helping Security Information and Event Management (SIEM) and analytics tools correctly prioritise alarms and events presented to the CTI lead/analyst.	Threat data: threat indicators and severity ratings, linked to attacks targeting specific industries, applications, etc.
Inform Departmental Risk Profession	Enhance the security assurance and risk management process with contextual content from intelligence gathering	Threat data: threat indicators and severity ratings, linked to attacks targeting specific industries, applications, etc.
Prioritise Vulnerabilities	Create a metric for evaluating vulnerabilities, by measuring the overlap between the problems which can be fixed and those with the most impact, given the time and resource available.	Vulnerability data: CVEs linked to attacks against specific industries, CVE's linked to specific threat actors, etc.
Support Threat Hunting	Proactively uncover hidden attacks on a department's network, related to current incidents, or threats targeting the department.	Threat data: indicators with links to context regarding campaigns, threat actors, techniques, history and targets.
Contain and Remediate Attacks	Disrupt attacker communications/ command and control, remove malware.	Threat data: intelligence knowledge base including data on techniques, history and targets of various threat actor groups.
Anti-Phishing	Enhance existing mail protection capabilities by enriching detection datasets with indicators.	Threat data: indicators with links to context regarding campaigns, threat actors, techniques, history and targets.

Table 1 – CTI Use Cases

These indicative use cases provide an overview of the types of activities a CTI function could undertake. If there is no appetite for the delivery of any of these use cases within your department, you should consider if CTI will meet your department's needs. Additional detail on these use cases is discussed later in the paper. The remainder of this section provides a roadmap to delivering these use cases and wider CTI capability to a department.

2.3 Roadmap

There are several steps which can be taken using existing or a minimum of new cyber security resources to develop a nascent CTI capability. We recommend tackling the steps in this section in the order they are listed. The steps:

- Represent the natural progression of the development of a CTI function
- Are prioritised in order of value
- Build on each other, requiring only the previous steps as pre-requisites
- Represent a progressive increase in expenditure, allowing for effective management of budget

2.3.1 Step 1 – Talk to Your Peers

This guide was written in collaboration with members of the cross-government CTI Working Group, the NCSC and commercial third parties. The CTI community is fundamentally collaborative, with the vast majority of departments facing the same threats as their peers or competitors. If your department is looking to stand up a CTI function from scratch, speak to the NCSC, your cluster lead or your regulator in addition to consuming the content of this guide.

2.3.2 Step 2 – Understand the Threat Landscape

Threat intelligence fundamentally involves gathering and distributing information about threats to a department. Those threats need to be identified and prioritised to ensure the CTI function provides the department with the most valuable intelligence. Understanding the threat landscape requires significant expertise so we recommend engaging with the NCSC and, if proportionate, engaging a commercial provider to provide subject matter expertise. It is not the purpose of this paper to conduct a market comparison; however there are several potential providers in the market who deliver this capability.

Performing a threat assessment enables departments to understand and prioritise key threats and threat actor groups. This analysis should be done in the context of the departments business, to understand who is likely to target the organisation and how they are likely to do it; together providing a prioritised view. This then informs the strategy. Additional detail on threat intelligence assessments is given in section 4.

2.3.3 Step 3 – Set a Strategy

As with any business function (finance, HR, security operations) CTI must demonstrate value to the department within which it sits and deliver in line with a pre-defined set of objectives.

The first exercise in standing up a CTI function is to define what the overall objectives of the function are going to be and identify what benefit that will provide the department. For most departments, this will involve a mix of technical and non-technical objectives.

The key objectives set within a department's CTI strategy should cover the following:

- What key outputs the CTI function will deliver
- What information the function will collect
- Which threat actor groups the function will focus its attention on
- Integration with Managed Service Providers (MSPs)

These objectives allow for the department to target its CTI function effectively; focussing on the key threats faced and allowing stakeholders to track its delivery and maturity. Some exemplar objectives could be created from the use cases given in section 2.2.

It is necessary to have experienced leadership drive the CTI function, including an effective CISO or other empowered senior staff member who can take who can take executive action if required. It is useful to have a CTI professional working on behalf of the department at this stage, to ensure that that the development process effectively captures requirements and delivers business benefit. It may be that existing members of the department (whether in the CSOC, physical security or other team) have the necessary skills and experience. If not, individuals may need to be recruited from outside the department.

A full overview of setting CTI strategy is given in Section 4.

2.3.4 Step 4 – Recruit and Stand Up the Capability

Once a strategy is in place, the next step is to recruit individuals to fill the roles identified. Some departments may assign the roles to existing staff in CSOCs, or network administrators. Alternatively, individuals may need to be recruited externally. However, this may prove to be difficult as the skills required are in high demand in the market. A breakdown of recruitment options is given in Section 10.

The number of resources required by a department will vary, depending on its threat profile (which does not necessarily correlate to a department's size). Your threat assessment will inform your threat profile and give an indicative view of how many individuals will be required to meet the CTI functions objectives. The CTI team can start with the following three roles (not necessarily three individuals), or they can be conducted as functions of other roles:

1. A CTI lead, who leads and manages the function, is responsible for delivering the strategy and delivers intelligence upwards to the board, senior management and system owners as required.

2. A CTI analyst who collects, analyses and processes information from non-automated sources (e.g. industry papers) and takes responsibility for profiling the activities of threat actors identified in the threat assessment.

3. A CTI analyst who collects, processes and analyses technical intelligence, and provides a dedicated point of contact to the CSOC. This analyst will support the profiling of threat actors, but their primary responsibility will be the maintenance and delivery of technical intelligence indicators to the CSOC and other defenders.

The CTI function may sit as a dedicated function under the head of security, or under the head of CSOC. CTI has several responsibilities which extend outside of a CSOC and into the wider department, and the CTI function needs the autonomy to prioritise those activities over those required by a CSOC, if necessary.

2.3.5 Step 5 – Mature Your Deliverables

The purpose of a CTI function is to provide intelligence relevant to its department. This does not require (particularly in a nascent capability) the CTI function to be at the cutting edge of intelligence investigation, analysing zero-day attacks from advanced persistent threats (APT). Start slowly – use Open Source Intelligence (OSINT) reports (and other relevant sources, e.g. reports shared by government and industry on CiSP) to refine the process of collecting, analysing and reporting intelligence relevant to your department. These reports will allow the function to begin to understand the Tactics, Techniques and Procedures (TTP) their threat actors are using, and create products based on that activity for use within the department.

As the cycle of intelligence collection, processing, analysis and dissemination improves, analysts can start to look further afield for information sources and to critically analyse content for reliability and relevance.

As the CTI function grows in maturity, it is likely that there will be a requirement to procure additional tools and feeds to provide a better service. This is likely to include:

- Additional threat intelligence feeds
- Intelligence analysis services

Threat intelligence feeds should be analysed not just on price, but upon relevance to the department. There are a number of feed providers in the market, and each focus on different threat actor groups and industry sectors. Pick the feed that provides the most relevant content to your department; this will help ensure the best value for money.

Intelligence analysis services (such as domain interrogation tools, malware analysers etc.) should be reviewed in line with the processes which generated the requirement for them. As these requirements begin to scale as the function analyses more information, there will be a requirement for automation.

2.3.6 Step 6 – Automation

In addition to reports and briefings, a key part of cyber threat intelligence is delivering technical intelligence to the CSOC. A key issue is that there is an overwhelming amount of information to process in this space, and a number of departments have been unable to effectively tackle the “firehose” of content.

Our recommendation is two-fold:

- Explore the usefulness of technical intelligence using open source technologies (e.g. the Malware Information Sharing Project, (MISP) and the existing capabilities of your SIEM tool. This allows your department to effectively define its requirements for technical intelligence processing, and the non-functional aspects of a system such as integration requirements and data volumes.
- Utilise centrally provided intelligence through platforms such as CiSP or feeds from trusted providers as much as possible. Existing threat feeds should be used to gain situational awareness of the evolving threat landscape.

The vast majority of departments who contributed to this paper purchased a commercial threat intelligence platform (TIP) to process their technical intelligence, and subsequently found that in some areas their purchase failed to meet their needs. We recommend an open source approach to ensure requirements are defined as comprehensively as possible before purchasing decisions are made, improving the value of investment.

2.3.7 Step 7 – Become a BAU Function

Once the CTI function is resourced, has begun creating deliverables and has a regular drum beat of activity, the next step involves integration with the other technology functions, such as new project on-boarding, joiners/movers/leavers processes, and integration with managed service providers, etc.

At this stage, the CTI function should be self-sustaining, with actionable intelligence provided across the department and being regularly judged based on tangible metrics of value. Once this stage has been reached, the CTI function will be at or above the level of maturity described in the remainder of this paper. Further development will depend on the department's priorities, but we would encourage departments to continue to collaborate to collectively better defend themselves and each other from threats.

2.4 Timescales

Delivery timescales for any business function will vary from department to department. Based upon the experience of collaborators of this paper, the timeline for the creation of a CTI function in an idealised environment is likely to take around 18 months, from initial engagement of threat assessment provider to the BAU operation of a CTI function. Some additional factors which will delay this process are:

- Creating a business case and securing funding
- Revision of strategy or scope adjustments
- Availability of appropriate resources suitable for recruitment into CTI roles
- Resource churn

It is therefore likely that most organisations will take 2-3 years to successfully stand up a basic but competent CTI function.

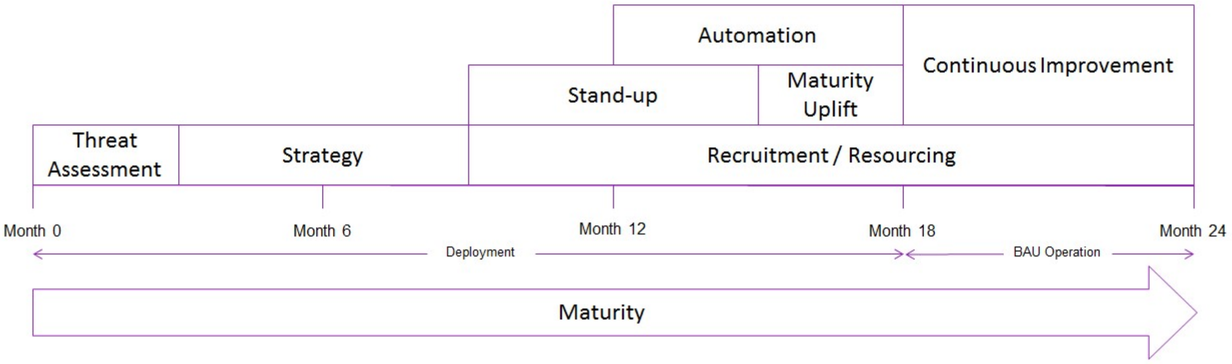

Figure 2 – CTI Development Roadmap

2.5 Cost

Cyber Threat Intelligence is a dedicated capability, which sits alongside existing security capability, both within cyber and traditional security roles. As stated above, CTI enables these functions to better protect a department, and is of limited value without them already being in a mature operational state. Consequently, it is not recommended that the budget assigned to these functions is repurposed to fund a CTI capability – CTI should be funded standalone.

This paper does not contain market analysis of the cost of CTI products, resources or services. However, in delivering the roadmap contained within section 2.3, departments should expect to spend at least £500,000 in delivering a basic capability over the initial 18 months, rising with third party licence costs and additional resources as the department matures. For more guidance on budgeting, reach out to the NCSC or your cluster lead.

3 Definitions, Scope and Structure

3.1 Cyber Threat Intelligence

Cyber Threat Intelligence has several definitions, but for the purposes of this guide, the following definition will be used:

> ***"Cyber Threat Intelligence is the process of collecting, processing and analysing information regarding adversaries in cyberspace, in order to disseminate actionable threat intelligence, by understanding adversaries' motivations, capability, and modus operandi, to inform cyber security mitigation measures."***

This is a broad definition; this guide concentrates on the delivery of a core CTI capability and does not define operational processes or other department specific procedures. Note that this definition deliberately implies that a CTI function provides information to defenders of an organisation to better enhance the security of that organisation – rather than making changes to an organisation's security posture itself. This guide remains vendor agnostic throughout.

3.1.1 Mission Statement

An exemplar mission statement for a Cyber Threat Intelligence function is proposed below:

> ***A Cyber Threat Intelligence function shall seek to collect, analyse and disseminate actionable intelligence to their organisation's defenders.***

A CTI function should be designed to support and improve the effectiveness of the defenders in a department. It is not itself a defensive capability and cannot deliver any benefit in isolation.

For a CTI function to provide value, mature capability should first exist in the areas of protective monitoring, risk management, change management, incident management and asset management. However, each department should implement CTI based on their specific requirements and threat profile.

3.2 The Threat Intelligence Lifecycle

The threat intelligence lifecycle is a model initially defined in the context of military intelligence, but the core principles remain relevant to civilian departments, for both traditional and cyber intelligence.

While a number of variations exist, this paper uses a lifecycle of 5 phases: Direction, Collection, Processing, Analysis and Dissemination.

- **Direction** refers to the strategy and objectives of a CTI function, and the requirements provided by their customers;

- **Collection** refers to the types, sources and mechanisms of gathering data;

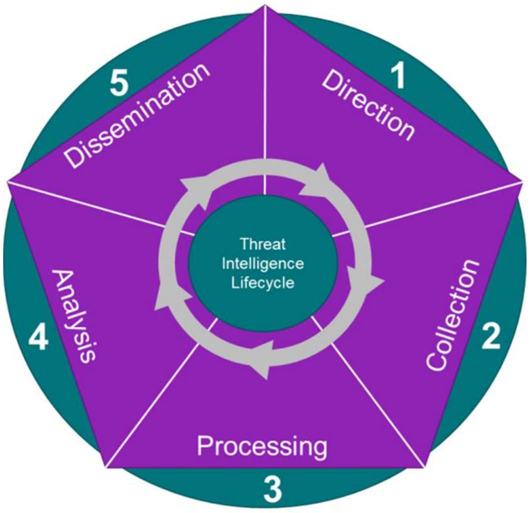

Figure 3 – Threat Intelligence Lifecycle

- **Processing** refers to the actions (automated or manual) that translate collected data into useful information for analysis;
- **Analysis** refers to creation of actionable intelligence from processed information;
- **Dissemination** refers to the distribution of intelligence products to the function's customers and partners.

3.3 Organisational Delivery Levels

Organisational/departmental activity is commonly split into three levels: strategic, operational and tactical. Within the context of CTI, a fourth level – technical – is commonly used, however this guide has included it within tactical, as that is where it is most relevant. Each level follows the threat intelligence lifecycle in a similar but distinct way, with each having different customers and therefore different requirements and outputs.

Strategic CTI is high-level and business-focussed, usually in the form of prose e.g. reports or presentations aimed at Senior Management Teams (SMT) within the department. The purpose of strategic CTI is to assist the SMT in making informed business decisions by providing them with an understanding of threats to the department, which will then feed into established strategic risk management and resource management processes. Common sources of strategic CTI include geo-political affairs, industry white papers and trusted networks.

Operational CTI, commonly in the form of tools, techniques and procedures, aims to understand threat actors and their likely attacks against the department. This allows security managers to allocate resources to take defensive action against the highest priority threats. Operational CTI also allows Incident Responders (IR) to respond more effectively during investigations by providing the required context to pivot from initial IOC (Indicator of Compromise) to potential attribution of the attack to a threat actor. Attribution allows a better understanding of the attackers' motivations, infrastructure, capabilities and target. Common sources of operational CTI include threat actor reports, incident reports, malware analysis, and occasionally social media and chat rooms.

Tactical CTI is more technical in nature and consists of the analysis of IOCs to allow the CSOC/network administrator to more effectively triage alerts and distinguish active attacks that require escalation from background noise. Tactical CTI is used to update technical controls based on the IOCs received from external and internal sources, e.g. adding block rules to firewalls and blacklist domains on the internet gateway, either automatically, or manually via the Network Operations Centre (NOC). Common sources of tactical CTI can stem from incident reports, audit/monitoring logs, threat hunting, and threat feeds.

CTI activities will not necessarily fit within one level, and these definitions should be used as a guide rather than a definitive structure.

4 Direction

1 Direction
2 Collection
3 Processing
4 Analysis
5 Dissemination
Threat Intelligence Lifecycle

This section provides an overview of the first step in the Threat Intelligence Lifecycle – Direction. Direction can be provided at the strategic, operational and tactical level from different customers within the department. Direction is the most important step to delivering a useful CTI function, as it drives all later phases. Effort should be spent to ensure requirements are gathered correctly from the start, to minimise the impact of changes at a later stage.

This section covers the steps required to set an effective CTI strategy and define requirements/objectives for the CTI team to meet. It is recognised that many departments may already have set a strategy, however we encourage all departments to conduct a review of this regularly, considering the following recommendations if applicable.

4.1 Cyber Threat Assessment

To understand the objectives of a CTI team, a department must first have a basic view of its threat profile. Whilst previous cyber risk management structures used within Government have supported this through a departmental threat assessment as prescribed in the Information Assurance Standard No. 1 (IS1) framework (6), these documents do not effectively articulate the threat landscape as it is currently evolving.

When considering a threat intelligence strategy, a mature cyber threat assessment (TA) should be the primary guidance for departments. The key differences between the previous methods of a threat assessment, and a mature cyber threat assessment, are shown in the table below:

Past Threat Assessment	Mature Cyber Threat Assessment
A wide range of threat actors are considered, regardless of functional relevance (e.g. radio enthusiasts).	Only threat actors with the capability and motivation to attack the department are assessed in detail.
Threat actors are grouped in a coarse fashion (e.g. “Foreign Intelligence Services (FIS)” or “Organised Crime Groups (OCG)”).	Specific threat groups are considered on a case by case basis, dependent on capability and motivation, and regardless of any formal label.
Departmental assets are considered coarsely in terms of business impact level.	There is an in depth understanding of the business, and critical business assets are considered individually with specific threat actors and attack scenarios considered.
Threats such as FIS are often discounted as being “out of scope of OFFICIAL”.	There is recognition that most threat actor groups are using commercially available, detectable attacks and that intelligence on their capability is of value.

Table 2 – Past TA vs. Mature Cyber TA

Use of a mature cyber TA allows for the threat to a department to be better understood, and therefore allows examination of how much of a department’s resources should be actively devoted to combatting that threat. Generally, departments should not spend more resources (capital or effort) on risk mitigation than the financial impact of the risk being realised, unless obliged to by legislation or regulations.

KEY POINT: Before starting a cyber threat intelligence programme, generate a mature cyber threat assessment.

If you do not have the expertise to create a mature cyber threat assessment for your department in-house, then it is strongly recommended that this activity is outsourced and may require engagement from the private sector. For large government departments, it is recommended that you consult the NCSC. For other departments, it is recommended that you engage with the threat intelligence contact for your cluster. A brief overview of the process is provided for information.

4.1.1 Threat Assessment Process

The first exercise is to ensure that the department in question understands its critical assets. For many large departments, these assets can include systems, networks, datasets and platforms which may reside across different locations and environments. Whilst from a technical perspective, important assets can (and should) be recorded in a Configuration Management Database (CMDB), this asset prioritisation exercise should also include intangible assets and intellectual property.

Once the view of a department's assets has been acquired, the assets can then be assessed as to their importance to the department, which will give a view of the cost of loss of those assets. A second view of the department's assets should also be acquired, which reflects the value of assets to an attacker. There may be some assets a department uses which may not be business critical, but in the wrong hands may prove extremely valuable. These two perspectives should be combined to create a view of what critical assets a department holds, and what the impact of compromise of those assets would be.

Once a department has a good understanding of their critical assets, the key threats to the department should be discussed. This problem can be broken down by using several sources of information. A useful list of input sources is given below (adapted from BAE Systems Applied Intelligence's Intelligence-Led Threat Mitigation paper (7)).

Input Source	Description
Incident Response Reports	Understanding attackers who have previously had success is key, as it is likely these adversaries will return. The tools, techniques and processes of previous attacks can inform threat intelligence analysis and further risk mitigation. IR reports can also include content from near misses or unconfirmed compromise.
Penetration Test Reports	Penetration tests (vulnerability assessments or red teaming activity) aim to detect functions or systems which are vulnerable to attack. This information can be added to the wider intelligence picture. Over time, the threat landscape and departmental systems change, and departments need to be wary of these assessments becoming inaccurate as each test is a point in time view of the vulnerabilities in a system.
Expert Advice	As discussed in the section above, CTI professionals can provide an impartial view of the threat landscape and distil from that the key threats to your department. SME's can leverage expertise from across a wide array of departments and contexts to track indicators and behavioural patterns.
Departmental Expertise	Those who are at the "sharp end" of a department's delivery will also have a view of the threats facing a department, even those outside specific security roles. This may include developers, administrators, business managers or members of legal. C-level will likewise have a view. Getting input from these members can prove valuable to understand threats, but also departmental maturity. Care should be taken to ensure that these are interpreted and prioritised appropriately.
Stakeholder Workshops	Collective discussion of threats between departmental members, suppliers and external parties (e.g. the NCSC) can provide useful input to the threat model, and better understand the scope of a department.

Table 3 – Input into Threat Assessment

The same paper provides the following detail on conducting a threat assessment:

> *"Both current threats as well as future ones should be considered. This may require input from strategy or development teams in order to gauge what the operating environment will look like in years to come. Geopolitical tensions and the rise of cyber capabilities in foreign militaries should not be underestimated as future threat sources.*
>
> *It is said that all models are wrong, but some are useful. It is important to bear this in mind and resist the urge to be too complete or precise. Threats which cannot be easily categorised initially, for example new attack sources using new capability, should be captured under a placeholder with a label such as 'emerging threats'. The model should be reviewed regularly and threats which have fallen into this last group can be given their own category as required."*

To complete the analysis mentioned above, the content from these various sources should be collated, analysed and fed into both a report and a model. Use of external partners in this process minimises the risk of confirmation bias[2], challenges entrenched departmental opinion and provides a much broader view of the threat landscape than can be gathered internally.

Each individual threat from the threat model can be further represented in widely-adopted models such as the Diamond Model of Intrusion Analysis (8), Lockheed Martin's Cyber Kill Chain (9) or other suitable methodology if appropriate.

Once both key threats and assets have been identified, the third step is to bring together both sets of information and prioritise which threat actors are likely to have the greatest impact on the department. This should take into account the attackers' motivation, capability and modus operandi. This can then be distilled into a departmental threat assessment report.

4.2 Setting a CTI Strategy

Once an organisation has identified its key threats, the strategy for the CTI function can be set. Whilst there is no fixed template for a CTI strategy, fundamentally it follows the same rules as setting a strategy for any other core IT or security function. This can include producing a vision or mission statement, setting the key objectives of the capability and defining the scope; as well as defining the products which the capability will produce. The CTI products will be both technical and non-technical, varying as they are presented across the strategic, operational and tactical levels.

KEY POINT: Set a strategy and objectives to ensure that CTI delivers value.

The strategy for the CTI function should be set with the budget owner, a representative of cyber risk owners across the business, the Head of Cyber Security and a relevant SME. It should avoid duplication with the objectives of existing teams. The strategy house below, documents an exemplar vision and objectives.

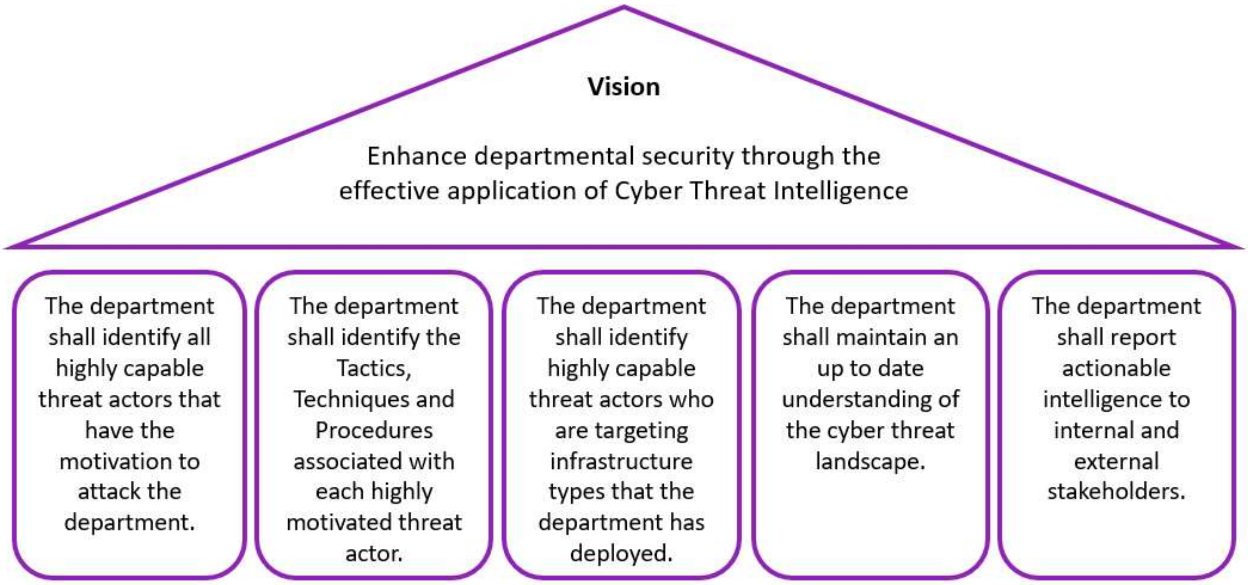
Vision
Enhance departmental security through the effective application of Cyber Threat Intelligence
The department shall identify all highly capable threat actors that have the motivation to attack the department.
The department shall identify the Tactics, Techniques and Procedures associated with each highly motivated threat actor.
The department shall identify highly capable threat actors who are targeting infrastructure types that the department has deployed.
The department shall maintain an up to date understanding of the cyber threat landscape.
The department shall report actionable intelligence to internal and external stakeholders.

Figure 4 – An Exemplar Strategy House

[2] Confirmation Bias is the tendency to search for, interpret, favour, and recall information in a way that confirms one's preexisting beliefs or hypotheses. (35)

4.3 Defining Operational Capability

Multiple methods of gathering requirements exist and the CTI function should follow their organisations standard business analysis approach to requirements elucidation e.g. workshops, stakeholder interviews, surveys, etc.

To begin defining requirements, the CTI function should ensure it engages with all relevant Business Units (BUs) on at least an annual basis to understand their plans and investments for the year, which will help to highlight potential threats. The intelligence obtained on a high priority threat actor can be compared to the target intelligence model to highlight areas where further effort is required.

Ad-hoc requests for CTI from customers will also drive requirements. To manage these requests effectively, an RFI process should be defined, with a template produced to capture all necessary information.

This could include fields such as:

- Requestor details;
- Date required by;
- Background to requirement;
- Requirement summary;
- Detailed requirements;
- Impact of delay;
- Intended dissemination.

4.3.1 Pre-requisites and Dependencies

The CTI function will have information requirements, which are essentially information dependencies needed to produce actionable intelligence. A useful example is the existence of a CMDB detailing hardware, software and operating systems used across the estate. These should be tracked to ensure they are met either by the function which provides them, or by the relevant dependency owners.

4.4 Financial Considerations

Departments may fund cyber security in a different manner; but, as CTI is an enabling service to defence, it is generally funded using the security or technical budget – usually whichever funds the operation of any existing CSOC or local security function. As stated in previous sections, CTI should not take from funding dedicated to other security enforcing functions – it requires standalone budget and management.

5 Collection

Threat Intelligence Lifecycle
1 Direction
2 Collection
3 Processing
4 Analysis
5 Dissemination

To fulfil the requirements set in the Direction phase of the Threat Intelligence Lifecycle, data needs to be collected. Collection is one of the most difficult steps of CTI to get right; ensuring that relevant data is collected in appropriate quantities from a multitude of different sources is a key challenge. This section will explore the different methods of data collection and the usefulness of each.

This section is split into two major areas – data that is commonly collected manually, and data that is commonly collected using automated means.

CTI content must be actionable, which in turn means it must be relevant, timely and accurate. A CTI function could ingest millions of IOCs, large volumes of data and information from feeds with relative ease, however unless it is actionable by the customers, then it is a burden on the function's resources. The content collected is therefore one of the most important areas to focus on when standing up a CTI function.

5.1 Manual Collection

There are various mechanisms for manually collecting data, but the key to gaining most value is to ensure that the collection processes are standardised. Key elements of manually collected data to be recorded are:

- Date and time
- Forecast timescales of relevance
- Nature of data gathered
- Specific technical records

These categories should integrate into the automated intelligence collection records detailed in Section 5.2.

5.1.1 CTI Sharing Networks

CTI sharing networks refers to formal or informal threat intelligence sharing between practitioners. This is common practice in government departments, for example, networks exist between those who had previously worked in the military or intelligence agencies, and are now working in the private sector, or in government departments. Some of these have developed into more formalised structures (such as the Cross Government Threat Intelligence Working Group). Regardless of origin, sharing intelligence between peers is critical to achieving success, and we encourage all departments be actively involved in contributing to cross government threat intelligence.

5.1.2 Open Source Intelligence

Open Source Intelligence (OSINT) refers to the collection of data from publicly available sources such as information on the Internet or in the media. This information is openly accessible for analysts to reach; however, the vast amount of content means that effective curation and management is required.

OSINT can include (but is not limited to) content from the following:

- Social media: social media can be monitored for threat information being published by researchers, commercial CTI providers or even the spokes-people for threat actor groups themselves. Intelligence gathered from social media must be scrutinised to ensure that the intelligence is accurate, and that is appropriate to gather it.

- WHOIS: a widely used internet record listing which identifies the owner of a website domain, and contact details. The CTI function can use WHOIS data to identify registered users of domain names, blocks of IP addresses or autonomous systems. For any registrant who hasn't opted to mask their information, the registrant's name, address, email and phone number can be searched. For example, threat actors may use registered domains for collecting ransomware payments using their own public email address, thereby unwittingly incriminating themselves, or associating themselves with more than one campaign.

- Domains and IP address analysis services: this information can be used to develop information about threat actors that control the infrastructure, including motivation, techniques, objectives, and more. This information can be gleaned from data sources that cross-index large volumes of information about domain registrants and IP address assignees. Various data points, such as domain registrants, IP address owners, DNS data, and more, can surface links between domains. This can help a department to obtain advanced warning of impending attacks where attackers are re-using IT infrastructure.

- CVE: Common Vulnerabilities and Exposures (CVE) is a catalogue of known vulnerabilities and the technologies which they relate to. The number of CVEs is growing, and it's relatively simple for a CTI function to enhance its capability by understanding CVEs relevant to its own infrastructure. This information can then be used (along with threat actor profiles identifying which threat actors commonly use which vulnerabilities) to prioritise vulnerability remediation. Anecdotally, it is common for departments to "accept the risk" associated with vulnerabilities which do not have high CVE scores. Enhancing knowledge of threat actor use of CVEs can provide context to that decision, which may affect the outcome of the decision to not remediate the vulnerability.

- Shortened URL Processing and Indexing: this is the translation of a long Uniform Resource Locator (URL) into an abbreviated alternative that redirects to the longer URL and understanding subsequent connections. Malicious threat actors use short URLs to conceal the actual URL, and plant malware and phishing links. The short URLs can bypass the security controls which block blacklisted domains, however they can be detected and analysed with appropriate tooling. Once analysed, shortened URLs can be either marked safe, or added to existing threat actor profiles if they redirect to known bad infrastructure.

Figure 5 shows four types of OSINT sources, ranked in order of speed to report, and in inverse order of reliability:

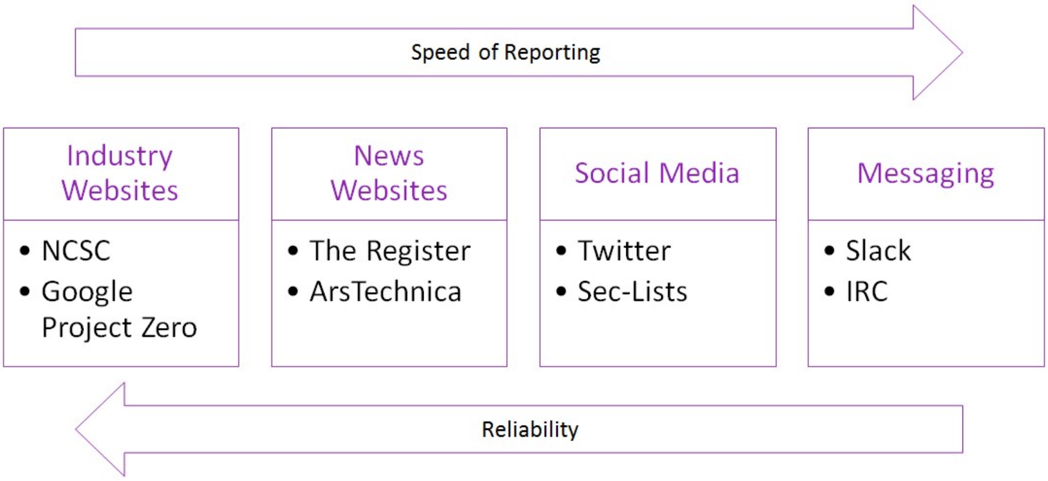

Figure 5 – OSINT Sources

In line with figure 5, whilst the speed of social media (in particular Twitter) makes it a widely used communications channel, analysts must proceed with caution to verify information being spread through viral sources. Information feeds such as Twitter have no filter and professional opinions are not clearly separated from personal opinions. Direct threats from criminal organisations via social media should be analysed to understand whether they have the means and motivation to commit an attack. Value from collecting OSINT requires critical analysis to differentiate reliable sources from unreliable. Effort should be taken to look past unconscious biases[3] as key information can often be found from sources that may not at first be considered.

KEY POINT: Multiple sources of OSINT must be used, while ensuring they represent value in proportion to the time invested for analysis.

5.1.3 Dark Web Content

Whilst there are use cases that require searching for content on the dark web, the clear majority of useful CTI content can be sourced from the open web or from commercial vendors. For a nascent CTI capability, we do not recommend exploring the dark web for content, as effort can be better spent elsewhere. If your organisation does wish to explore this area, speak to the NCSC or the National Crime Agency (NCA) for advice on reporting illegal content.

For a CTI capability, we recommend that a department's network should deny connectivity to any Tor entry and exit node, as well as track any attempted activity from departmental infrastructure and these nodes. CTI capabilities should also stay up-to-date on the ever-changing public list of Tor servers, attackers regularly use Tor as a bridge from their infrastructure to a target, and therefore traffic to and from Tor nodes should be monitored closely.

5.1.4 Limited Distribution Content

Limited distribution content is data that is either distributed only within a limited working group, or is commercially purchased, and provides a collection source for CTI.

There are many free, closed sources available to government departments; the most notable of these is the Cyber Security Information Sharing Partnership (CiSP). The NCSC manages CiSP and regularly shares content to it, along with other government departments and industry partners. It

[3] Unconscious Biases are learned stereotypes that are automatic, unintentional, deeply engrained, universal, and able to influence behavior. (36)

should be a priority for any government CTI function to maintain an active presence on the platform. However, CiSP is not a curated source, and in line with your department's threat assessment, content from it should be filtered based on relevance.

Commercially purchased content varies in usability and relevance. Different CTI vendors produce reports on different threat actors and campaigns, and the relevance should be a key consideration when looking to purchase access. It is not the purpose of this guide to compare vendor offerings; however, these sources of data should be considered when maturing a CTI capability. We recommend piloting collection process using OSINT or free sources before purchasing content.

When standing up a CTI function, care should be taken to ensure that the core sources of CTI are reliable, and we recommend leveraging content provided by the NCSC through CiSP as an information source to start with. However, it is recognised that in a mature CTI function the majority of content is likely to come from paid for sources which have been curated specifically for the department in question.

5.1.5 Use of Protectively Marked Content

Consumption of intelligence from protectively marked sources is encouraged; however the processes to do so are likewise protectively marked. As a consequence only a limited discussion is possible here.

Depending on partnerships, content can be received from a protectively marked report, memo or in an automated format. This content can be processed using similar processes to those used for other protectively marked intelligence but should accommodate for the restrictions placed on processing the information by its classification and or handling caveat.

Moving information from higher classifications to lower requires the use of an import-export process, or IMPEX. Content provided at a higher classification should not be re-used at a lower classification without express permission of the owner of that content, and that content should pass through the content owners IMPEX process prior to being used in a low classification setting. No specific systems or processes for IMPEX can be discussed here as the process varies on a system by system, stakeholder by stakeholder basis. We recommend speaking to NCSC or the provider of your high classification content if you have questions on this topic.

5.1.6 Legal Considerations

CTI functions must ensure that threat intelligence information is sourced using legal means, adhering to the Code of Practice for Covert Human Intelligence Sources if relevant. This code of practice provides guidance and rules on authorisations for the use or the conduct of covert human intelligence sources under Part 2 of the Regulation of Investigatory Powers Act 2000.

5.2 Automated Collection

Key to the effectiveness of CTI is the automated gathering of technical content that can be actively used to prevent attacks. Most defensive infrastructure capabilities (e.g. Intrusion Detection Systems (IDS)) across government are signature-based and require rules to be continuously updated to provide protection from emerging threats. Defensive infrastructure can ingest content such as the IP addresses of attacker's infrastructure, domains being used in phishing campaigns or hashes of known malware. This information is readily available to analysts, but the sheer volume makes manual processing and submission to defensive infrastructure impossible. This is widely recognised as one of the key challenges in the threat intelligence space.

5.2.1 Threat Intelligence Platforms

Threat Intelligence Platforms are a mechanism to collect, process, analyse and disseminate CTI. TIPs automatically ingest and reconcile data from various sources including OSINT, threat feeds and sharing partnerships such as the industry-centric Information Sharing and Analysis Centres (ISACs), or CiSP.

Once data is collected by a TIP, it is processed, which consists of consolidation across different sources and formats, removal of duplicates, and validation and scoring of IOCs. TIPs can integrate with various technical controls, and once the data is processed, they can automatically provide IOCs and effect rule and configuration changes to mitigate against known knowns, e.g. block known IPs used for command and control. This functions as an automatic method of dissemination for tactical CTI. If integrated with a SIEM solution, a CSOC analyst or network administrator can quickly leverage the TIP to provide context around observed IOCs, providing CTI at the operational level and assisting with decisions on actions or escalations. TIPs can also provide reports and metrics for senior and functional managers.

5.2.2 Malware Information Sharing Platform

The Malware Information Sharing Platform (10), originally called Cyber Defence Signatures, started out in 2011 as a project to overcome the frustration of IOCs being shared only in human-readable format, i.e. PDF or email. Over time, the North Atlantic Treaty Organisation (NATO) became involved and provided support, and now MISP is a project with a large community and is used by over 6,000 organisations.

We recommend utilising MISP to pilot automated threat intelligence capabilities. MISP has the benefit of having all of the core features of a TIP but is open source and free to deploy. MISP allows for the processing of threat feeds, threat actor profiling, SIEM integration and other capabilities, however it does not necessarily scale well to becoming an enterprise offering. Whilst MISP has no licence cost, for BAU operational capability the value of vendor support and the effort required to develop a mature platform should not be underestimated.

The remainder of section 5 and subsequent sections reference where MISP can be used to deliver automated intelligence capability.

5.2.3 Threat Feeds

Threat feeds are ongoing streams of data related to threats. They can range from relevant industry experts on Twitter, through to prose reports or IOCs being published.

Many SIEM vendors make third-party feeds available with their products, if relevant, these can be consumed. The NCSC are also in the process of creating a threat feed. When this feed becomes available, it is likely to be one of the most relevant feeds available for government. Further information can be sourced from the NCSC website.

KEY POINT: Begin threat feed analysis using a trusted feed.

Threat feeds most commonly use standardised methods for storing and distributing content. This involves the use of Structured Threat Information eXpression (STIX) and Trusted Automated eXchange of Indicator Information (TAXII). STIX and TAXII are industry-recognised specifications, designed to enable automated information sharing for cybersecurity situational awareness, real-time network defence, and threat analysis.

5.2.3.1 Structured Threat Information eXpression

STIX is a format for exchanging all aspects of CTI such as indicators, suspicion and attribution. It's not a sharing program or tool, but a component which supports programs or tools. Alongside STIX, a standardised methodology for its sharing was also developed –TAXII, which is discussed in Section 5.2.3.2.

The STIX 2.0 format defines twelve STIX Domain Objects (SDOs), with each representing a category of different attributes. STIX also defines two STIX Relationship Objects (SROs). These are Relationship, which links two SDOs, and Sighting*,* which represents the observance of related malware, IOCs, etc. While usually stored as JavaScript Object Notation (JSON), STIX can also be graphically represented, with SDOs as nodes and SROs as edges. Further details on SDOs and SROs can be found on the Organisation for the Advancement of Structured Information Standards (OASIS) Open GitHub. (11)

The SDOs and their relationships are shown in figure 5.

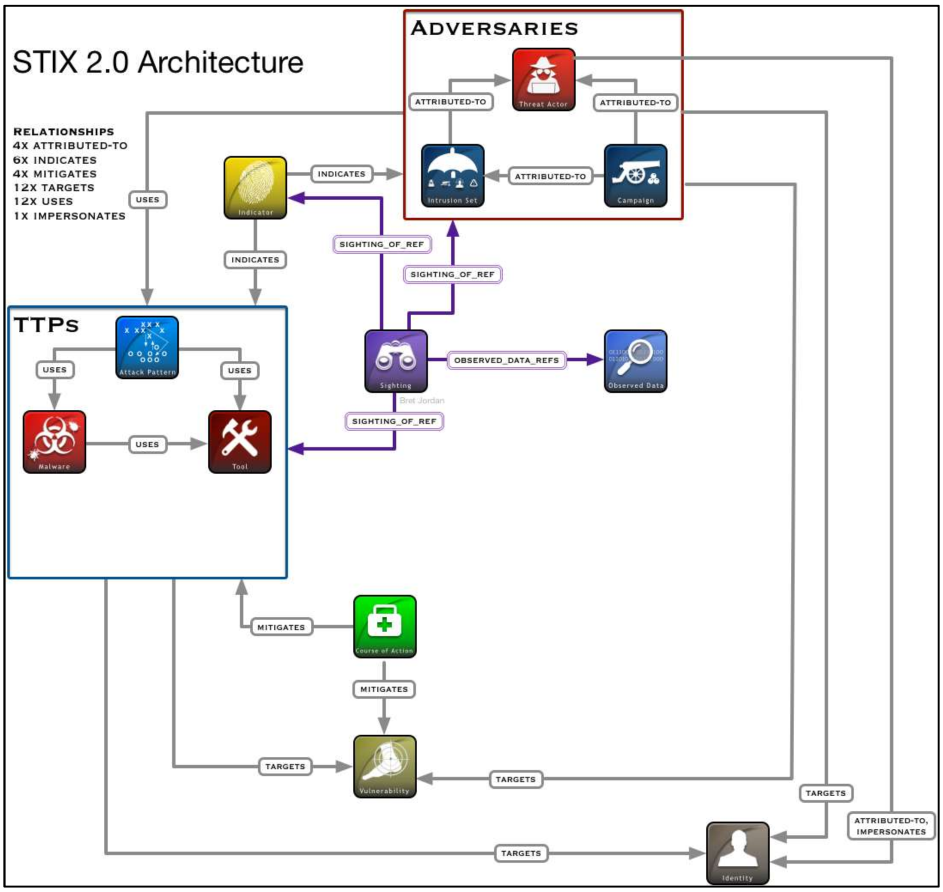

Figure 6 – STIX 2.0 Architecture

Further STIX Objects are Marking Definition and Bundles. Marking Definition defines the handling requirements of the data as defined by the Traffic Light Protocol (TLP, explained in Section 8.1.2), and Statement, which covers text-based handling requirements. Bundles are arbitrary groups of STIX Objects and Marking Definitions, primarily used for sharing.

A good real-world example of STIX in use is MITRE's JSON representation (12) of Mandiant's APT1 report. (13)

STIX's primary attribute is the Indicator, which provides insight into the infrastructure or TTPs that attackers are using. Examples of Indicators are shown in Figure 6.

Sharing of Indicators can reduce the likelihood of a successful attack by providing partners with a means of detection and analysis for their defensive infrastructure. Indicators can be gathered from numerous sources, such as malware analysis, threat hunting or automated detection.

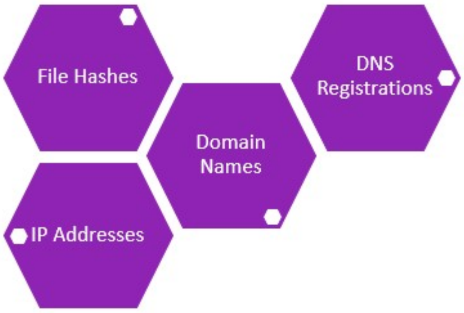

Figure 7 – Exemplar Indicators

5.2.3.2 Trusted Automated eXchange of Indicator Information

TAXII is a protocol designed to exchange STIX over Hypertext Transfer Protocol Secure (HTTPS); however, both STIX and TAXII are independent and one does not require the use of the other. TAXII is simple and scalable, offering two modes of sharing – Collection and Channel – that are defined by a set of requirements for Clients and Servers, and a Representational State Transfer (REST) Application Programming Interface (API).

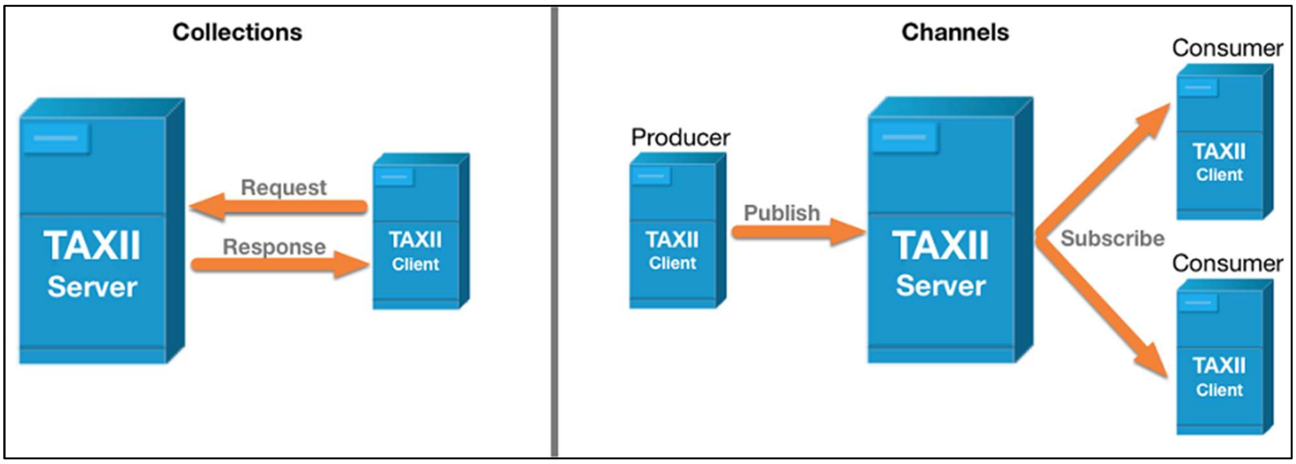

Figure 8 – TAXII Modes of Sharing

5.2.4.1 MISP Threat Feed Integration

MISP can be used to collect information distributed either using the MISP formats, STIX/TAXII, or others. Over fifty threat feeds are included in MISP, and these are a useful starting point to validating ingest processes and management. Integrating threat feeds can be done either through the User Interface (UI) or the python API. The simplest way to add feeds to MISP is to navigate to the list of feeds in the UI, check the ‘Enabled’ field, and then repeat for each feed to be integrated. For a third-party feed this can either be entered through the UI or the API, by passing the URL and any authorisation headers into the text box in the UI or the script for the API.[4]

[4] A full pilot of a MISP server and feeds was performed to inform this paper. Contact the authors if you wish to re-use the code used.

5.3 Natively Produced Content

Threat Intelligence is not just about content collected from third-parties, it relies on content being created and shared by departments. Effective creation comes primarily from detection events reported to a CSOC, network administrator, or from write-ups of first and third-party content from internal sources.

5.3.1 Creating Indicators

If an element of the defensive infrastructure creates an alert associated with IOCs that have not been previously seen, these IOCs should be reported. Care should be taken to ensure that these events are not false positives. Examples include beaconing that is detected by an IDS, or patterns associated with a brute force attack detected by a web application firewall. Processes for creating intelligence will be department-specific and should be defined as part of the CTI operational process.

CSOC analysts or network administrators may regularly come across examples of active malware which have been served to their department by malicious actors. It is useful to understand what the IOCs related to these samples are, so that subsequent attacks can be rapidly identified. To gain the most value from captured malware samples in a CTI context, the process in Figure 8 is recommended.

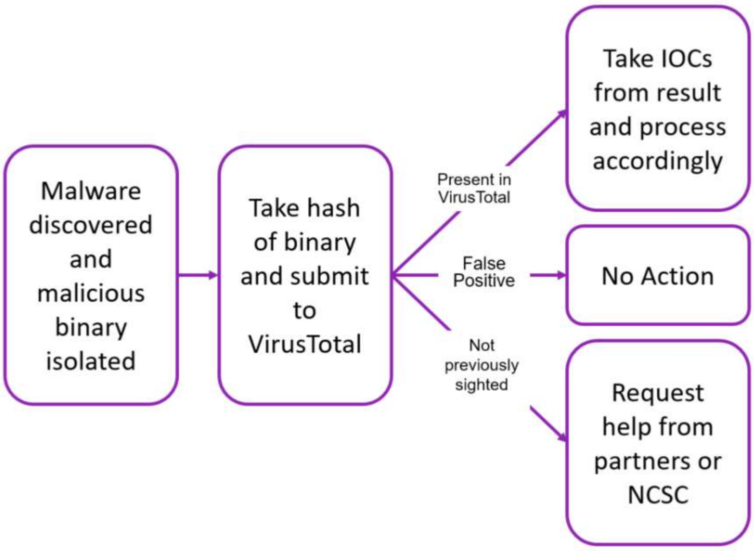
Malware discovered and malicious binary isolated
Take hash of binary and submit to VirusTotal
Present in VirusTotal
False Positive
Not previously sighted
Take IOCs from result and process accordingly
No Action
Request help from partners or NCSC

Figure 9 – Malware Analysis *(38)*

It is rare to identify malware which has not been previously identified. More detailed malware analysis than simply hashing is a specialist activity, and requires knowledge of reverse engineering, assembly and other areas. Consequently, individuals who have the skills required to be effective in this field are rare and command significant salaries, limiting the ability by departments to recruit and retain them.

As a consequence, we do not recommend analysing malware samples in-house; instead as part of incident response processes consider where third party assistance can be leveraged to meet this need.

5.3.2 Sightings of Indicators

If IOCs that have been collected are subsequently detected, a Sighting should be reported. This can be done via the SDO attribute within STIX, or the sighting flag within MISP. Other frameworks also include a method for acknowledging sightings.

5.4 Purchasing Tooling

Departments often purchase threat intelligence tools and services before understanding their requirements, so they do not realise the maximum benefit. To overcome this, define your automated intelligence gathering and processing requirements as described in Section 3. Also, it may be beneficial to collaborate with NCSC, consuming their CSOC tooling guidance. Requirements can then feed into an open source pilot for validation. Once the pilot is complete, requirements should be reviewed again prior to any commercial procurement.

KEY POINT: Consider an open source pilot to validate TIP requirements.

Once technical CTI capabilities have been piloted and understood comprehensively, the next step may be to purchase tooling. It is recognised that for most departments, some commercially purchased tools are likely to be needed to deliver the CTI function’s objectives.

An area to note is that ultimately, vendors are selling products/services, and as such marketing material should not be taken at face value. Departments should ensure the product meets their specific requirements. A demo or functional analysis may be of benefit, and we would recommend that you contact the NCSC or your cluster lead before engaging with the market.

6 Processing

All collected datasets are only as good as their content. Processing of collected data enriches the content to ensure that the most useful information is available to analysts. There are three main considerations when processing content:

- Reputation: how trustworthy is the source of this information?
- Relevance: is the content relevant to your department?
- Quality: is the content of sufficient quality to be useful?

This section describes how best to enrich content to better provide information for analysis, and to manage the volume of data available. This is a key challenge – a recent Ponemon (14) report showed seventy percent of respondents say CTI is often too voluminous and/or complex to provide actionable intelligence.

For each of the three areas, metrics can be created for automated and manually managed sources which allow intelligence to be processed and prioritised prior to analysis. We recommend that labelling of intelligence sources is done in a coarse fashion. Subdividing further than this may allow the best content to be highlighted more effectively, but is likely to create significant redundant sub-categories, and may lead to valuable content being de-prioritised.

As a CTI function matures, processing metrics can be refined to ensure that only useful, relevant content is made available to analysts with other content de-prioritised. The following sections give basic examples of metrics which could be applied to process data, but should be viewed as a starting point for further analysis rather than a mature solution.

6.1 Reputation

All data should be processed in the context of the reputation of its source. For example, if information is received from a trusted contact, it can be held in higher regard than that sourced from an unverified Twitter account, and this principle should be extended to threat feeds. The primary question to differentiate content is therefore “How trustworthy is the source of this information?”

To standardise this process, a metric for the reputation of content can be created and incorporated. Our recommendation is that sources from national or international authorities (e.g. Content from NCSC or NATO) is given a high reputation score as these organisations have a vested interest in providing good quality, actionable content. Feeds from commercial sources may also have a good reputation. When consuming publicly available content (such as voluntarily maintained or community threat feeds) consider giving these sources a lower reputation score unless assurances are available that the content provided is actionable and has a minimum of false positives.

Each item of CTI can be evaluated using the NATO System (also known as Admiralty Grading System (15)), which comprises a two-character notation assessing the reliability of the source, and the information’s assessed level of confidence. It can be tailored as necessary to refine its usefulness for CTI. We would recommend that whilst the NATO System is used as a guide, that fewer categories are used in each case, or that categories are merged as the difference between categories may be open to analyst interpretation.

Reliability:

The CTI source is assessed for reliability, based upon a technical assessment of its capability, or history. The notation for reliability uses one of six characters, from A to F:

A. Completely reliable: no doubt of authenticity, trustworthiness, or competency; has a history of complete reliability.
B. Usually reliable: minor doubt about authenticity, trustworthiness, or competency; has a history of valid information most of the time.
C. Fairly reliable: doubt of authenticity, trustworthiness, or competency but has provided valid information in the past.
D. Not usually reliable: significant doubt about authenticity, trustworthiness, or competency but has provided valid information in the past.
E. Unreliable: lacking in authenticity, trustworthiness, and competency; history of invalid information.
F. Reliability cannot be judged: no basis exists for evaluating the reliability of the source.

Credibility:

The CTI source is assessed for credibility, based upon the likelihood and levels of corroboration by other sources. The notation for credibility uses one of six numbers, from 1 to 6:

1. Confirmed by other sources: confirmed by other independent sources; logical in itself; Consistent with other information on the subject.
2. Probably True: not confirmed; logical in itself; consistent with other information on the subject.
3. Possibly True: not confirmed; reasonably logical in itself; agrees with some other information on the subject.
4. Doubtful: not confirmed; possible but not logical; no other information on the subject.
5. Improbable: not confirmed; not logical in itself; contradicted by other information on the subject.
6. Truth cannot be judged: no basis exists for evaluating the validity of the information.

6.2 Relevance

Collected data should be relevant to the CTI function and its requirements. Even for relevant sources (such as the NCSC threat feed), it may be that not all provided content is useful.

It is recommended that analysts use similar metrics to reputation to assess relevance for feed management, however there is no metric for relevance within the NATO System. Threat feeds often contain a wide variety of content - for example, many threat feeds include content related to ransomware that is targeted at individuals rather than enterprise. Feeds which prioritise this content can be de-prioritised, as existing defensive infrastructure such as web filters should mitigate the associated risks.

Additional down selection can be done at the event or document level. If a particular event has specific relevance to your department this should be prioritised; for example, indicators which are relevant to threat actor groups which were identified as significant threats to a department in the threat assessment should be marked as of interest. This principle can be applied across a number of contextual content types, e.g. threat actors involved, malware campaigns, technology or targeted departments.

6.3 Quality

The final consideration to be taken when processing CTI is the quality of the collected data. Key areas to examine include:

- Format of the content, e.g. IPv4 addresses should be XXX.XXX.XXX.XXX
- Associated context of the content
- Age of the content
- Accuracy of the content, i.e. “true” IOCs as opposed to false positives
- Frequency of updates in a feed

Consideration should be taken to ensure data is collected in appropriate formats. Sources that provide incorrectly formatted data should be seen as low quality compared to those that consistently provide content in compliant (and therefore automatically processable) formats. Indicators for example will always have specific formatting rules; IPv4 addresses should be in the format of XXX.XXX.XXX.XXX and domain names will always contain @.

Associated context provided by a feed should contribute towards its quality. If a feed regularly reports just IP address indicators, that should be viewed as lower quality than a feed that reports full contextual information, e.g. threat actors, campaign and domains associated with an IP indicator.

The age of content is also key, and should be processed in the context of IOC type, e.g. IPs have a short lifespan, hashes have a longer lifespan, etc. Threat feeds that regularly provide content past its useful lifetime should be de-prioritised. This is most apparent in the context of cloud, where attackers regularly spin-up and spin-down offensive infrastructure. This leads to the regular cycling of IP addresses and other re-usable signatures. Figure 9 demonstrates how difficult it is for attackers to modify indicator types, which in turn informs the lifespan of each indicator type.

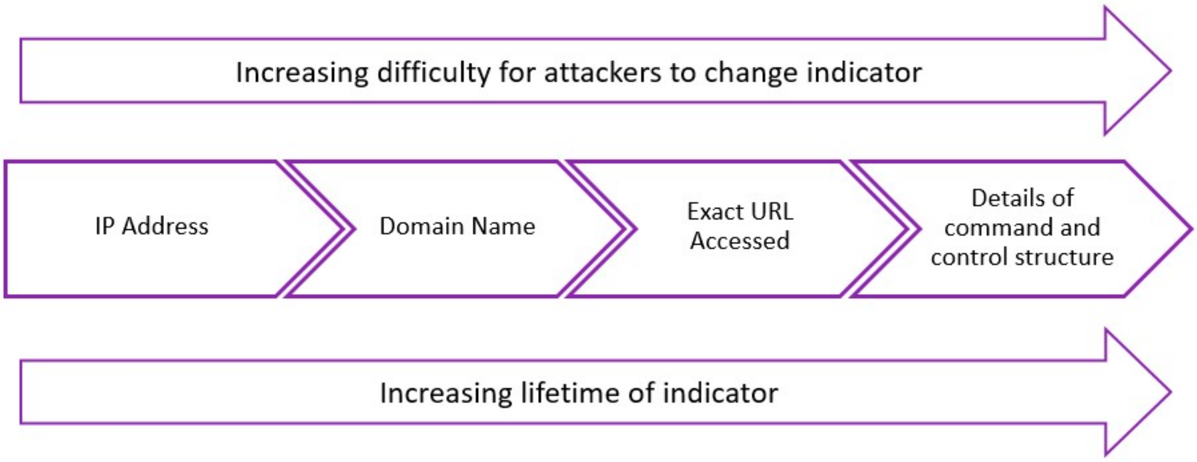
Increasing difficulty for attackers to change indicator
IP Address
Domain Name
Exact URL Accessed
Details of command and control structure
Increasing lifetime of indicator

Figure 10 – Lifespan of Indicators

The accuracy of the IOCs, (i.e. false positive rate) is also a key consideration. It should be noted that false positives may be delivered both accidentally and deliberately, and this may relate to the source’s reputation. Given the amount of effort required to check the validity of CTI indicators, threat feeds which regularly report false positives should again be de-prioritised.

The final consideration on data quality is whether the feed is regularly maintained – researchers have a variety of backgrounds and departmental priorities. In particular, open source threat feeds may be only sporadically updated compared to commercial threat feeds that are frequently updated.

6.4 Combining Metrics

Once metrics have been created for reputation, relevance and quality (regardless of whether the source of intelligence is automated or manually acquired) these can be combined to effectively prioritise content. The specific thresholds for "must analyse" content and content which can be ignored will depend on your department's risk appetite, team bandwidth and amount of content consumed. Metrics and thresholds should be regularly reviewed to ensure that the optimal amount of intelligence is made available for analysis.

6.5 Automation

Most Threat Intelligence Platforms provide some ability to automatically process and enrich content as it comes in. MISP has a feature called tagging, which can categorise incoming content based on a user-defined set of metrics. It is recommended to initially tag each event (MISP term for bundled associated IOCs) based on its reputation, content and quality. Tagging can be used to read from the event's JSON file and find the number of attributes, and their data types – a script is recommended for this. Automated tagging is best explored using the API.

6.5.1 Reputation

Tags can be applied to specific feeds, providing a first pass view as to the reputation of the incoming content.

6.5.2 Relevance

The content from each event can also be automatically tagged. It is recommended to start by using the API to search for your department's key threat actors and apply tags on these. Due to the variety of sources for the feeds, it is also recommended to create a "synonyms" JSON file, listing each Threat Actor and any alternate name they are known by (e.g. Sofacy is also known as Fancy Bear, APT28, Pawn Storm, etc.). This allows each threat actor to be identified by all its synonyms but tagged by a single common name. A limitation of MISP is that tags are not automatically grouped by synonym, so a script must be written to do this separately. These tags, combined with quality tags, can create an overall priority rating for analysts.

Further tags, such as incident classifications, malware varieties and stage in the Cyber Kill Chain can be added in the same way.

6.5.3 Quality

Two metrics that can classify the quality of received events are the count of the number of attributes, and the diversity of attribute types. A high count and diversity may correspond to a higher quality event, but the baseline figures for a high quality will be subjective to each department.

Additionally, several tools exist which allow the feed quality itself to be scrutinised. In relation to STIX, these include:

- STIX Validator: this validates that the STIX JSON content conforms to the 2.0 specification
- Pattern Validator: this validates that the patterning syntax conforms to the patterning expression, e.g. that IPv4 addresses comply with Classless Inter-Domain Routing (CIDR) notation

7 Analysis

1 Direction
2 Collection
3 Processing
4 Analysis
5 Dissemination
Threat Intelligence Lifecycle

Analysis of collected information is required to produce CTI products. Intelligence cannot provide absolute certainty; therefore, careful analysis is important to ensure that any conclusions drawn are sufficiently robust. The purpose of the analysis phase is to take processed content and convert it into actionable intelligence products for consumption by the CTI function's customers and partners.

7.1 Products of a CTI Function

CTI products must provide actionable intelligence to the department, namely intelligence that is accurate, relevant and timely, otherwise it may be of little value. A simple relevance test could be answering *yes* to questions such as "Does this threat/event/etc. require a change in our security posture now, or in the foreseeable future?" or "Does the Senior Management Team expect us to understand this threat/event/etc?" Aside from the content, CTI products must also be presented in the correct manner for each audience, i.e. jargon-free, concise and pitched at the right technical level.

There are a range of possible CTI products which can be produced. The table below, shows some typical products:

Strategic	Operational	Tactical
Annual threat assessments.	Project specific threat assessments.	Incident support.
Annual threat landscape reports.	Department relevant write-ups on key events (e.g. WannaCry, Spectre/Meltdown, etc.).	Enriched IOC feed to defenders.
Quarterly/six monthly briefings to the SMT.	Threat actor and campaign reports.	

Table 4 – CTI Products

Understanding the governance, incident and change processes is the key driver to deciding which products to target and at what frequency. The following products fall into the category of "situational" or "standing" products:

Situational Products:

- Threat Intelligence Alert: used in time-sensitive scenarios or to convey tactical threat intelligence
- Threat Intelligence Assessment: used to analyse the threat to a specific application or system (or departmental change programme) where new infrastructure or capability is being delivered

Standing Products:

- Standing Threat Intelligence Assessment: established CTI analytical position on a given threat and the residual risk it poses
- Security Forecast: six monthly view of the strategic threat landscape

Your department will need to reference its requirements to select the most appropriate products for it. Some examples are given in Table 3. A CTI product specifically excluded from the list is a monthly report to the business – something which was produced by a lot of organisations we collaborated with. The reason for its exclusion is that several the departments which provided insight to this paper stated that while this product was of interest, its actual value is limited. Unless a product can demonstrate tangible value (i.e. it makes a material difference to the defensive posture of a department) it should not be prioritised.

7.2 Threat Actor Reports

7.2.1 Purpose of Threat Actor Analysis

As the number of APTs increase, there is a need to proactively prevent attacks from such actors. Processed information can be further analysed to produce actionable intelligence. At a strategic and operational level, this tends to be in the format of prose reports, as detailed in Section 4.3. Analysis of the threat actors, their associated attributes, and relationships between them is pivotal to identifying themes, motives and TTPs. At the tactical level, this is likely to be in the format of indicators associated with campaigns run by a threat actor.

7.2.2 Naming Conventions

Naming conventions vary among commercial providers, and other departments. For example, CrowdStrike name actor groups after animals such as Panda and Bear, whereas Mandiant favour “APT” followed by a number. An up to date view of threat actor names (and the conversions between them) can be found at apt.threattracking.com.

7.2.3 Attribution

Ideally, all attacks would be attributable to a specific threat actor; however, this is unrealistic, and on the occasion that attribution is given, there is usually a caveated degree of uncertainty. It is highly likely that there has been, and will continue to be, cases where mis-attribution has taken place. As such, caution should be applied when consuming attributed content or attributing content directly. Reputation of the source of the information being analysed is key to making informed decisions about attribution.

7.2.4 Building a Threat Actor Profile

To gain an effective view of attackers’ activities and their modus operandi, it is useful to map intelligence to a threat actor profile. Models such as the Diamond Model, Lockheed Martin Cyber Kill Chain or MITRE ATT&CK framework can be of use here.

The profile should include details including:

- Motivations
- Infrastructure
- Targets
- Tactics, Techniques and Procedures
- Indicators
- Associated context, e.g. timing of attacks coincides with a specific time zones working day

Threat actor profiles can be further used to draw conclusions, either across a domain or for a specific attacker. For example, APT111 may be attributed to a spike in attacks on government owned cloud services and may have previously been associated with attacks on government. This allows the CTI function to gain confidence that APT111 is specifically targeting government departments.

Threat actor profiles can be generated in MISP in an automated fashion, using the processing mechanics in section 6.

8 Dissemination

5
Dissemination
1
Direction
2
Collection
3
Processing
4
Analysis
Threat Intelligence Lifecycle

Dissemination takes the products from the analysis stage and distributes them to the appropriate customers of the CTI function at the strategic, operational and tactical levels.

Separate from internal reporting, departments also need to consider the benefit of externally sharing relevant CTI (particularly operational and tactical CTI) to help prevent attacks within an industry or community. Shared intelligence of an attack on one department, such as attribution to a specific threat actor, their motivations, infrastructure, TTPs and IOCs, can enable other departments to improve their defensive posture and reduce the risk of compromise.

8.1 Distribution of CTI products

The method and format of distribution will vary from customer to customer to meet their tailored CTI requirements. For example, a high-level prose report on the trend of attacks on government can be provided to the SMT in an email; a low-level PDF prose report on a relevant threat actor's TTPs can be provided to Security Management and the Incident Response team via an Electronic Document and Records Management System (EDRMS, e.g. SharePoint); and validated IOCs can be provided to the CSOC and Network Operations Centre (NOC) via STIX/TAXII. Aggressive internal dissemination of CTI is key to optimising the benefit gained from the function. One discussion identified a department where approximately 25% of all staff received some form of CTI product via email (note however section 7's discussion of product value).

Across all departments surveyed, the distribution of CTI was noted to be critical to its usefulness. Where possible, CTI teams should maintain a view of who in their department is receiving intelligence and whether they are finding it useful.

8.1.1 Trust Relationships

The sharing of CTI has numerous benefits as previously described; however, it will likely require either the direct or indirect sharing of sensitive information between departments. This sharing requires a level of trust to be established between departments that the information shared will be appropriately handled and used only for expressly permitted purposes.

There is also the element of risk around sharing knowledge of threat actors, e.g. if a threat actor discovers that a TTP has been identified and mitigated against, then they will likely move onto a new and unknown TTP.

One method of establishing trust is via closed sharing forums. For example, CiSP is selective about which non-government organisations can become members, while MISP allows the establishment of arbitrary trust groups.

Whilst the above comments should be considered, fundamentally the success of CTI is in sharing content promptly to ensure all consumers can make use of it whilst it remains actionable. This may require creating relationships between companies which would not normally collaborate, such as direct competitors in the same sectors. One such example of this being successful is Target and Walmart in the US (16). All departments are encouraged to become active members of sharing communities such as the cross government CTI working group, part of CiSP or as part of an ISAC.

8.1.2 Document Marking

Traditional methods of sharing information in government are unsuitable for CTI where sharing is a priority (you will notice this document is deliberately not protectively marked). We recommend using the Traffic Light Protocol (TLP) to share CTI content due to its sector wide recognition, broad adoption and flexibility.

8.1.2.1 Traffic Light Protocol

The Traffic Light Protocol was setup by the UK's National Infrastructure Security Coordination Centre, the precursor to the Centre for Protection of National Infrastructure (CPNI), to encourage the sharing of sensitive information and help establish trust. The current standard is now defined by the Forum of Incident Response and Security Teams (FIRST) Standards Definitions and Usage Guidance (17).

TLP uses four colours to designate sharing boundaries and indicate when and how sensitive information can be shared. It is optimised for human-readable information however both STIX and MISP have incorporated TLP. Important principles of TLP are that the source is responsible for ensuring recipients understand and can comply with TLP markings, and that recipients seek explicit permission from the source prior to sharing more widely than the TLP markings indicate.

The four TLP designations are:

- TLP:RED = not for disclosure, restricted to participants only;
- TLP:AMBER = limited disclosure, restricted to participants' departments or organisations;
- TLP:GREEN = limited disclosure, restricted to a specific intelligence sharing community;
- TLP:WHITE = disclosure is not limited.

8.1.2.2 Criticism of TLP

Discussions conducted in preparation for this guide's publication generated suggestions that TLP is not fit for purpose due to the lack of clarity between designations.

However, TLP fundamentally has value because it is widely adopted in the CTI industry. Alternatives, such as using the Government Security Classifications (GSC), do not provide flexibility in dissemination as the rules surrounding it are strict and enforced, though each department should choose a handling scheme that meets their requirements. Extra caveats can be added by the owner, if required.

If unclear what the sharing restrictions on a marked object are, or for permission to share it more widely, **ask the owner**.

8.1.2.3 Suitability of the Government Security Classification for CTI

The Government Security Classifications (18) are the replacement for the Government Protective Marking Scheme (GPMS). The GSC indicates the sensitivity of information and specifies the required baseline security controls (administrative, physical and logical) to appropriately protect assets.

The GSC has three classifications as described below:

- OFFICIAL = the majority of information that is created or processed by the Public Sector. This includes routine business operations and services, some of which could have damaging consequences if lost, stolen or published in the media, but are not subject to a heightened risk profile. Some OFFICIAL information is particularly sensitive and requires additional controls to enforce the 'need to know' principle – this should be marked OFFICIAL-SENSITIVE;

- SECRET = very sensitive information that justifies heightened protective measures to defend against determined and highly capable threat actors. For example, where compromise could seriously damage military capabilities, international relations or the investigation of serious organised crime;

- TOP SECRET = Her Majesty's Government's (HMG) most sensitive information requiring the highest levels of protection from the most serious threats. For example, where compromise could cause widespread loss of life or else threaten the security or economic wellbeing of the country or friendly nations.

The GSC is designed to ensure that government information is appropriately protected. This means that mandatory controls are required, proportional to each classification. This is, by design, an impediment to sharing of information and there is no flexibility in the controls required. It is therefore recommended that departments do not regularly mark CTI with the GSC, and instead use the TLP unless mandated to by operational restrictions.

8.1.3 Cyber Security Information Sharing Partnership

Launched in March 2013, CiSP now sits under the management of the NCSC, a part of the Government Communications Headquarters (GCHQ). CiSP is an online sharing portal described as "a joint industry and government initiative set up to exchange cyber threat information in real time, in a secure, confidential and dynamic environment, increasing situational awareness and reducing the impact on UK business."

UK companies responsible for administration of an electronic communications network in the UK can register, as can companies sponsored by a government department, regional cyber PROTECT police officer, or an industry champion. Once the department is registered, staff then need to individually register for access.

Membership of CiSP provides the ability to securely engage with other government departments and industry partners to seek advice and learn from each other. Discussion is encouraged from beginner through to expert level, and this collaboration helps to provide an early warning of threats and improves members' ability to protect their assets. CiSP also provides free access to network monitoring reports.

CiSP as an active community provides a useful forum to share and collect intelligence content and provides a functional platform to do so. We recommend all departments maintain a presence on CiSP and utilise its services.

8.2 Technical Intelligence Dissemination

Machine-Readable Threat Intelligence (MRTI) feeds (such as MISP JSON, Snort, YARA or STIX) disseminate IOCs to defensive infrastructure, allowing for automatic ruleset or configuration changes in near real-time. Most MRTI formats are based on eXtensible Mark-up Language (XML) or a derivative format, which use tags to provide information readable by both humans and machines.

8.2.1 Integration with Defensive Infrastructure

Once a collection of technical content has been marked for integration into defensive infrastructure, there are several ways to forward content. The most effective is for defensive capabilities to regularly pull content lists from trusted locations, either via an API or via the regular saving of documents in machine-readable format. MISP can deliver content in a variety of formats depending on that required by defensive infrastructure.

As noted in the 2017 Ponemon report (14), integration with defensive infrastructure varies in effectiveness between infrastructure type. Next Generation Firewalls (NGFW) and Unified Threat Management solutions are easier to integrate with, whereas endpoint security and IDS or IPS devices tend to be more challenging; however, the ability to integrate with defensive infrastructure will also vary by vendor.

Additionally, CTI feeds should not be the single source of truth, instead they should complement existing rulesets rather than override them. Additional rulesets may be required to prevent false positives from affecting the network. For example, if the department relies on a cloud platform range, a specific whitelist rule should be in place for that infrastructure to prevent legitimate activity from being blocked.

8.2.2 Rolling Rule Changes

Many IOCs are only actionable for a limited timeframe, so adding old indicators into defensive infrastructure blacklists has diminishing returns and may result in legitimate content being blocked in the future. In addition, most defensive infrastructure has a performance limit on the number of IOCs which can be processed.

Indicators should be prioritised based upon their relevance to the department, as discussed in previous sections, with associated rules having specific retention periods. The ability to apply retention periods on rulesets can, for example, be delivered either through creation of rule update scripts from a MISP server, or directly through your defensive infrastructure ruleset management process. Sightings of indicators is also a useful metric to define how long to maintain a rule in the defensive infrastructure. If an indicator is sighted, it implies that the threat actor is still actively using the relevant infrastructure and therefore is still a threat. As such, it is appropriate to then reset the retention period. For example:

- High priority indicators (i.e. those most relevant to the department and with high confidence) can be ingested and applied to the defensive ruleset for one year, and directly reviewed after the year to validate their continued use based upon hit rate, current intelligence and indicator context. Sighted indicators should also be marked high priority

- Medium priority indicators are ingested and applied to the defensive ruleset for three months, then automatically deleted unless the indicator has been sighted

- Low priority indicators are ingested and applied for one month, then automatically deleted unless sighted

This structure allows for indicators to be efficiently managed and the thresholds between high, medium and low can be varied depending on infrastructure capability and departmental requirements. Timescales listed above are indicative and should be tailored to your departments infrastructure. Once indicator management has been streamlined, the next step is to enable the sharing of sightings back to other departments, via a TIP such as MISP.

8.2.3 External dissemination

In line with the discussion in section 8.1.1, it is recommended that the sharing of indicators is proactive, but within trusted circles. MISP provides both the ability to share events created by the owner to subscribers to a feed and to mark indicators as “sighted” to add further context to event content. Whilst there is not currently an active sharing network for technical content which everyone can contribute to, it is anticipated that this will become available through NCSC in the near future.

9 Continuous Improvement

To ensure that a CTI function provides and demonstrates value to the business, it must continuously examine the quality and usefulness of its outputs. For each product type that the CTI function produces, an appropriate performance metric should be set.

9.1 Deliverable Metrics

For the selection of products listed in Table 4, this section proposes a potential metric that can be used to measure its quality.

Strategic products are the most challenging to create effective metrics and feedback for. When delivering content which is designed to provide awareness, the value of the delivery is regularly delayed and may not be easy to measure.

Strategic Product	Metric
Annual threat assessments.	Scoring from a templated feedback report can provide useful feedback. Key metrics to include are whether the customer found the product useful, and whether they took direct action as a result.
Quarterly/six monthly briefings to the SMT.	We recommend recording the number of times that a risk is modified, or decision taken on the basis of information provided by the CTI team. This can then provide a view of whether CTI briefings have a tangible impact on SMT's decision making.

Table 5 – Strategic Metrics

Operational products better lend themselves to metrics, as each is a regularly repeated deliverable resulting in tangible actions.

Operational Product	Metric
Project specific threat assessments.	We would recommend that delivery metrics are designed around the time to respond to requests, and the number of requests per year. Metrics should also be created on the quality of assessments, such as how many project risks were better informed as a result of the assessment.
Department relevant writeups on key events (e.g. WannaCry, Spectre/Meltdown, etc.).	These deliverables are primarily intended to brief the SMT on the department's position regarding cyber threats in the news or assessed CTI sources. The metrics for strategic deliverables may be used.
Threat actor and campaign reports.	The metric for this deliverable is proposed to be the number of times information gathered through a deep dive campaign report is subsequently used by the CTI team or by defenders within a fixed timescale. If a report surfaces no useful information, then decisions should be taken in context as to whether it was of genuine value to produce.

Table 6 – Operational Metrics

Metrics for tactical products must be specific and should measure the quality of information provided (be that IOCs or otherwise), rather than the volume.

Tactical Product	Metric
Incident support.	Measurement of support to incident management is difficult. We recommend a metric based around the number of occasions where the CTI function has delivered content to facilitate a more effective remediation. This should be standardised against the number of incidents investigated.
Enriched IOC feed to defenders.	Measurement of the enriched IOC feed can be achieved by taking statistics from defensive infrastructure, such as the number of positive alerts, and the ratio of positive alerts to false positives.

Table 7 – Tactical Metrics

9.2 Key Performance Indicators

Key Performance Indicators (KPIs) for the CTI function should be created to ensure that it can be held accountable to non-technical stakeholders as required. Therefore, it is proposed that the CTI function is measured on the speed of its delivery following request and the volume of quality content produced (see above for metrics of quality). These metrics will vary based on the size and capability of the CTI function and should therefore be agreed between the CTI lead and the relevant SMT representative. Associated technical functions within the department, such as architecture or assurance, may also have KPIs that can be leveraged.

9.3 Additional Considerations

CTI functions should evolve not only with their adversaries, but also with their customers. As the CTI capability continues to grow, the team should be aware of new development projects being delivered, to ensure that the CTI priorities continue to be aligned with the department and its objectives.

For example, the following changes to government IT ways of working may change the threat intelligence required by a department:

- Bring Your Own Device (BYOD)
- Adoption of cloud
- Containerisation technologies
- Novel technologies, as yet unknown

10 Organisation

One of the key areas of delivering a CTI function is understanding how the capability fits into the department, and the skills that are required. Security capabilities regularly sit within sub-optimal locations in the organisational structure – often for historic or political reasons. Standing up a CTI function provides a useful opportunity to restructure and bring traditional security and cyber security together, gaining value from their integration.

10.1 Structure

Many functional CTI team structures exist, and each department should look to their requirements for guidance. However, we recommend two routes to maturity and discuss the pre-requisites needed to ensure that the capability provides maximum value as it grows and develops.

The two major routes to integrating a CTI function into a department are:

- Expand existing security or intelligence teams within a department to cover CTI requirements
- Stand up a dedicated CTI function within the CSOC

10.1.1 Existing Intelligence Structures

Use of existing intelligence structures may be most appropriate to departments with a specific reason to maintain an in-house intelligence capability, e.g. required for Ministerial protection.[5] For these departments, a CTI capability is a natural extension of physical threat intelligence, and can use existing reporting, dissemination and triage techniques. Likewise, it does not require much additional training or technical analysis to upskill existing intelligence analysts in the cyber arena – the opposite is much more challenging.

10.1.2 CSOC Stand-up

For some departments, building a threat intelligence capability within the CSOC provides a natural extension of existing CSOC responsibilities and is therefore the most sensible location to stand up a CTI capability. Existing CSOC analysts may have intelligence expertise and may be analysing indicators already through threat hunting or incident response responsibilities.

Regardless of the capabilities location, the CTI team should ideally be dedicated resources.

KEY POINT: CTI is a standalone capability and should report independently of protective monitoring.

[5] Note this includes most large government departments and many Financial Times Stock Exchange (FTSE) 250 companies. If you are not aware of an in-house intelligence capability, speak to your Head of Security as it is likely that their teams are already gathering intelligence in some capacity.

10.2 Roles and Responsibilities

10.2.1 Leads and Analysts

One model for a CTI capability is the use of two essential roles – analysts and leads. Analysts are primarily tasked with collecting, processing and analysing intelligence content prior to internal dissemination. Leads are focussed on direction, setting strategy, managing the team and leading on content dissemination. Table 7 lists the key responsibilities between analysts and leads.

Analyst	Lead
Analyses OSINT content and threat feeds.	Manages analysts, e.g. tasking, ensuring quality of deliverables, etc.
Maintains automated CTI collection, e.g. management of a TIP.	Reports to SMT on team performance.
Creates reports on specific events for dissemination to relevant parties.	Sponsors reports where required.
Creates regular updates for general dissemination.	Dictates strategy and priorities.
Advises stakeholders (such as system owners) of the wider threat landscape.	Agrees requirements with customers.

Table 8 – Roles and Responsibilities

10.2.2 Managing Analysts Effectively

In many large departments, there is a need for multiple analyst resources to effectively deliver CTI tasks; for a large department (10,000+ employees and a large threat profile) a CTI team may require four or more analysts. These analysts can be split across all tasks that the team needs to complete or can be given a specific focus or specialism. For example, a function can either:

- Focus each analyst on either strategic, operational or tactical analysis;
- Focus each analyst on different threat actor groups.

We recommend that each analyst is tasked with gaining a deep understanding of a different threat actor or actor group, as this provides the maximum amount of contextual understanding when analysing intelligence and creating deliverables. This also means that when an event is attributed to one of these threat actors, specialist knowledge is immediately available. This structure is practised in a number of mature departments, including the US National Security Agency's Cyber Security Threat Operations Centre. (16)

Additionally, it is useful to have an analyst who will lead the technical configuration and management of the TIP, and fully understands the defensive infrastructure of the estate. This role will also work closely with CSOC. Whilst this may only be a single resource, it is beneficial to have multiple members of the team familiar with this role to mitigate a single point of failure. We therefore recommend the team structure outlined in Figure 11 (section 12.5).

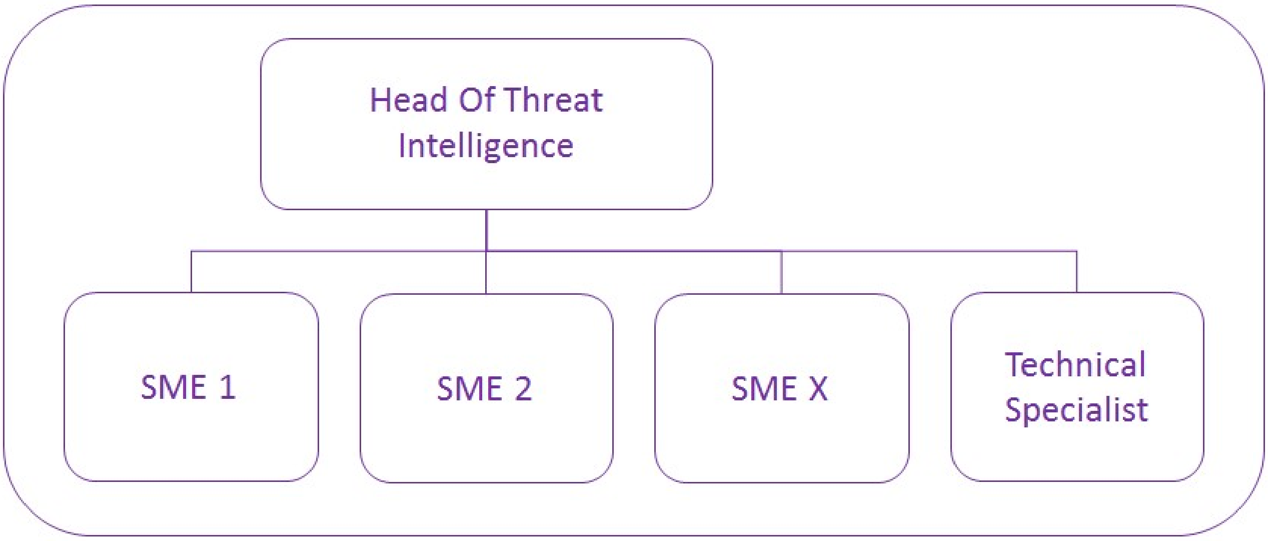

Figure 11 – CTI Team Structure

10.3 Resourcing

Resourcing roles in cyber security is challenging across all sectors, with a shortage of individuals possessing the necessary skills and expertise. One option is for a department to partner with a Managed Service Provider (MSP) to work alongside in-house staff. There is also an option to have personnel rotate responsibilities and to expose them to other parts of the department. Another option could be to recruit personnel from the department's IT support team, e.g. using a network administrator as a CTI analyst.

To gain the best value from a CTI team, there should be a minimum of one experienced professional from both a cyber security background and an intelligence background. There are six main avenues to pursue when recruiting. We have loosely ranked them in order of difficultly to recruit, and in salary they are likely to command:

Rank	Avenue
1	**Existing CTI professionals.** Recruiting existing CTI professionals is clearly the most preferable option, however they are limited in number, and command significantly higher salaries in the private sector than can be provided by government.
2	**Cyber security professionals from other disciplines.** As for Existing CTI professionals, cyber professionals more broadly are likewise expensive and in short supply. Even if a resource is recruited with existing cyber experience, it may be that their ability to analyse intelligence is limited and they do not provide proportionate value.
3	**Individuals with a background in intelligence.** Recruiting and training individuals from an intelligence background in cyber is an easier task than training individuals from a cyber background in intelligence. The core skill is to communicate uncertainty clearly, which can be readily provided by the intelligence sector. Individuals with a military intelligence background are also not as uncommon as cyber professionals.
4	**Outsource the CTI function.** Outsourcing of a CTI function may represent an attractive option due to the minimal effort required by the department; however, there are two main issues with this approach. The first is requirements definition, covered in Section 3, as the outsourced function will not have the business insight and engagement

Rank	Avenue
	required to perform this effectively. In addition, there is a consideration that outsourced CTI providers have varying levels of maturity in their offerings. The NCSC are currently working with the Industry 100 (19) to improve this, so if you are considering a managed service, talk to your NCSC representative for advice.
5	**Contractors.** Contractors are regularly used across government to meet the needs of the business where capability does not exist in house. Contracting may represent a viable option for CTI; however it carries significant expense and the department must clearly define responsibilities. In addition, the contractor market also has a shortage of the relevant skills and expertise, so recruitment will be challenging. CTI functions also benefit from having resources stay for long periods, as patterns in data related to APT activity may take many months to materialise, longer than the span of an average contract tenure.
6	**Individuals with no relevant background.** Recruiting individuals with no relevant experience e.g. Civil Service Fast Streamers is easier for the public sector; however, they will require significant training in order to be effective members of a CTI team. This has had mixed success when applied to other areas of cyber security – strong leadership is critical to ensure that inexperienced staff meet their objectives. Outside graduate recruitment, Intelligence Analyst apprenticeships are currently being developed as part of a Home Office programme (20), which can also be explored.

Table 9 – Resource Ranking

For the public sector, it may be challenging to recruit CTI professionals as pay regulations prevent offering a competitive salary. This extends to recruiting cyber security professionals from other disciplines, with the additional challenge of having to train those individuals in CTI practices.

Regardless of recruitment route, limiting attrition of skilled staff will be a challenge, again due to the inability of the public sector to offer competitive salaries in an aggressively competitive market. Fundamentally, there is always another department or organisation willing to pay higher salaries. This may be mitigated to a certain degree by specialist pay agreements for some roles within the Civil Service.

KEY POINT: Use a range of recruiting strategies to support the CTI function.

10.3.1 Diversity and Inclusion

As with all sectors in cyber security, there is currently a significant lack of diversity in backgrounds and experience. When recruiting members of the CTI team, a range of backgrounds should be explored, particularly those individuals without the traditional Science, Technology, Engineering and Maths (STEM) disciplines. Key areas such as international relations or geo-politics are both vital to contextualising CTI. The NCSC are making significant efforts to encourage diversity and inclusion in cyber; talk to your representative to better understand how you can improve in this area.

10.4 Training

It was not within the scope of this paper to assess any training courses. However the market has limited offerings for the training of CTI professionals and there is no comprehensive entry-level course for CTI analysts. The most relevant courses include:

- SANS FOR578: Cyber Threat Intelligence (21), which focusses on structured analytical techniques to better process and manage intelligence.
- CREST Cyber Threat Intelligence certification.
- HMG Cyber Threat Intelligence apprenticeship.

A number of CTI vendors provide courses to their internal staff, some of which may reach the market in the near future. More general intelligence training, specifically for the required writing and communication style, is available through several sources; specifically within government the Cabinet Office provides intelligence readership training across departments depending on need. If required, further information can be sourced from the authors.

11 Conclusion

Creating a CTI capability is a significant investment, but if done well can provide significant value to a department. This guide has reviewed the key areas for leads and analysts to operate a CTI capability, as well as additional elements of continuous improvement and organisation. This section reviews the key conclusions which have been highlighted throughout the paper.

11.1 Strategy and Planning

CTI functions have commonly grown organically from existing cyber or native intelligence capabilities in departments, and commonly lack direction or accountability. To ensure that a CTI function consistently delivers value to a department, firstly the threat profile of that department should be understood. This allows for the department to target the most significant threats to it, and gather information related to those threats.

KEY POINT: Before starting a cyber threat intelligence programme, generate a mature cyber threat assessment.

Defining a delivery framework with objectives and key deliverables allows the CTI function to focus its operations and ensure that all content analysed and reported is relevant to the department.

KEY POINT: CTI is like any other operational capability – teams need a delivery framework including a strategy, operational objectives and management.

11.2 Targeted Collection

The key challenge raised by all collaborating partners on this paper was that there is too much information available to effectively process and make actionable. In order to tackle this problem, we recommend that departments focus on collecting content from sources which are known to be reliable, and to focus on content which is relevant to their department.

KEY POINT: Target intelligence collection on priority threat actors.

Publicly available sources (known as OSINT) can prove to be immensely important to a CTI function, and CTI analysts should make use of these. Examples of these OSINT sources include social media, news sites and WHOIS data.

KEY POINT: Multiple sources of OSINT must be used, while ensuring they represent value in proportion to the time invested for analysis.

11.3 Automation

In addition to targeting collection, the problem of threat intelligence feeds providing too much information can be remediated through automation. We recommend that threat feeds are consumed and processed automatically, scoring each feed entry according to its reliability, relevance and credibility. This enables analysts to spend time analysing the intelligence which is of best quality, and drive as much value from their time as possible.

KEY POINT: When using threat intelligence feeds, automate your analysis.

In order to limit the scale of the intelligence collection problem, we recommend consuming a trusted threat feed, maturing business processes for consuming intelligence provided then expanding the scope to other intelligence sources.

KEY POINT: Begin threat feed analysis using a trusted feed.

11.4 Pilot Using Open Source

Threat intelligence vendors currently vary widely in the maturity of their offerings, and it is common for vendors to be unclear about the specific capabilities of their product. This, combined with customers with little knowledge of CTI platforms means it is common for purchased solutions to not meet expectations.

We recommend that departments pilot a threat intelligence platform capability using open source offerings (i.e. the MISP platform) to better understand what capabilities and features a TIP would provide benefit to a department. Once piloted, departments should be in a much better position to distinguish between vendor offerings, ensure product purchases offer value for money and are tailored to help mature the CTI capability.

KEY POINT: Consider an open source pilot to validate TIP requirements.

11.5 Reporting

Reporting refers to the authority and autonomy of the CTI function within a governance structure. CTI functions provide enriched intelligence content to all levels of a department and provide a service which enhances the effectiveness of defensive capabilities. Whilst it may be a natural extension of existing CSOC operations, CTI is a fundamentally separate function, and should be allowed to develop independently of a CSOC.

KEY POINT: CTI is a standalone capability and should report independently of protective monitoring.

11.6 Recruitment

Recruitment for the lead and analyst roles within CTI is a challenge, particularly in government where existing salary bands restrict recruitment of specialist resources, without using alternative funding sources. Additionally, if an adequate salary is not paid to recruited staff, they may leave government for the private sector. We recommend exploring recruitment from not just the existing cyber sector, but also exploring whether individuals from the wider intelligence community or those with non-traditional skillsets could deliver effectively.

KEY POINT: Use a range of recruiting strategies to support the CTI function.

11.7 Feedback

We are grateful to all individuals and departments who offered their time and expertise to the authors of this paper to bring it to completion. This paper marks a point in time view of CTI best practice across UK government which is likely to change and evolve as technology and the threat landscape changes. We welcome feedback on the paper from any source, please provide it to:

ctiplaybook@homeoffice.gov.uk

12 Appendices

12.1 Authors and Collaborators

This guide was written by the following individuals:

Author	Role
Rob Flanders	Lead Author
Lucy Johnson	Lead Reviewer
Matthew Trevelyan	Technical Author
Anna Whitmore	Business Author
Lisa Lesowiec	MISP Architect
Rajinder Tumber	Technical Author

Table 10 – Authors

We are grateful to the following departments who provided expertise and insight which enabled the production of this guide:

Government Departments:

- Cabinet Office;
- Centre for the Protection of National Infrastructure;
- Department of Work and Pensions;
- Foreign and Commonwealth Office;
- Government Digital Service;
- Her Majesty's Revenue and Customs;
- Home Office;
- Ministry of Defence;
- National Cyber Security Centre;
- Transport for London.

External Organisations:

- BAE Systems Applied Intelligence;
- Bank of England;

- Lloyds Banking Group;
- Orpheus Cyber.

12.2 Acronyms

Acronym	Definition
API	Application Programming Interface
APT	Advanced Persistent Threat
BAU	Business As Usual
BU	Business Unit
CMDB	Configuration Management Database
CISO	Chief Information Security Officer
CiSP	Cyber Security Information Sharing Partnership
CTI	Cyber Threat Intelligence
CVE	Common Vulnerabilities and Exposure
DNS	Domain Name System
DWP	Department of Work and Pensions
FIS	Foreign Intelligence Service
GSC	Government Security Classification
HMG	Her Majesty’s Government
IDS	Intrusion Detection System
IOC	Indicator Of Compromise
IP	Internet Protocol
IPS	Intrusion Prevention System
IR	Incident Response
ISAC	Information Sharing and Analysis Centre
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
MISP	Malware Information Sharing Platform

Acronym	Definition
MRTI	Machine Readable Threat Intelligence
NATO	North Atlantic Treaty Organisation
NCSC	National Cyber Security Centre
NCSP	National Cyber Security Portfolio
NOC	Network Operations Centre
OASIS	Organisation for the Advancement of Structured Information Standards
OCG	Organised Criminal Group
OS	Operating System
OSINT	Open Source Intelligence
REST	Representational State Transfer
RFI	Request for Information
SANS	The SANS Institute
SDO	STIX Domain Object
SRO	STIX Relationship Object
SIEM	Security Information and Event Management
SME	Subject Matter Expert
SMT	Senior Management Team
CSOC	Cyber Security Operations Centre
SME	Subject Matter Expert
SMT	Senior Management Team
STIX	Structured Threat Information Expression
TAXII	Trusted Automated Exchange of Intelligence Information
TIP	Threat Intelligence Platform
TLP	Traffic Light Protocol
TTP’s	Tactics, Techniques and Processes
UI	User Interface

12.3 MISP vs Other Technologies

As part of the technical research for this guide we considered a number of technologies. MISP was down selected as the ideal candidate as it is open source and contains sufficient capability deliver a proof-of-concept technical CTI platform.

In conversations with industry partners it was noted that during down selection of threat intelligence platforms MISP performed (from a technical standpoint) extremely well. However, as a enterprise class platform, it failed to meet non-functional requirements. MISP has difficulty scaling and providing a corporate supporting platform for a significant sized department is difficult.

A full breakdown of MISP compared to other CTI platforms can be sourced from various partners – please contact the authors if interested.

12.4 Additional CTI Capabilities

A number of activities related to CTI were excluded from the scope of this paper. The overarching reason for limiting the scope of this paper was time – CTI is very broad subject and in-depth exploration of all areas would be an unachievable exercise in the time available. The areas excluded from scope, and the rationale of their exclusion is provided below.

12.4.1 Threat Hunting

Threat hunting is the proactive, iterative and human-centric identification of threats that are internal to the network and have evaded existing security controls. This activity has been excluded as the intention of it is to search for existing malicious activity on your infrastructure, rather to analyse external content from the wider threat landscape. Whilst we wholeheartedly agree with departments performing this activity in a mature environment, exploration of it here would not be sufficiently in depth to provide appropriate guidance.

A number of firms have had some success in using rule based or machine learning driven technology to better deliver threat hunting capability. These capabilities are likewise excluded from the scope of this guide as it was not the intention of this exercise to do a supplier analysis.

12.4.2 Digital Risk and Intelligence

Digital Risk and Intelligence (DR&I) is the process of monitoring, detecting and remediating publicly available information, through the control of an organisation's digital footprint. This includes:

- What sites or IP addresses are public facing
- Which public facing cloud services the department is using (e.g. public facing login to a SaaS application)
- What information is available on the internet about a department
- What information about key members of the department is available on the Internet
- Which credentials owned by members of the department have been leaked (e.g. in data breaches)
- What content related to the department exists within darknet forums or markets

Digital Risk and Intelligence has not been included in this guide for two reasons:

1. This area overlaps with threat intelligence exploration but the fundamental goal of it is to identify department information in the public domain, not to understand the activities of threat actors.

2. The amount of capability development required to deliver an effective digital risk and intelligence capability requires significant additional effort and integration on top of that in a core threat intelligence capability.

Given the complexity and limited overlap, this topic has been excluded from this guide. However, a paper on threat hunting and digital risk intelligence to complement this guide is being delivered as a parallel activity.

12.4.3 Sample Analysis and Reverse Engineering

Reverse engineering is the activity of taking known or suspected malicious files and attempting to extract useful information from them. This could include IP addresses or domains with which the malware communicates, hashes of malicious files, rogue processes spawned etc.

Malware analysis has been the topic of a number of books, courses and tutorials for many years, and it would be naive to attempt to include such a complex topic within this guide.

Whilst malware analysis can produce information which can be processed as threat intelligence (such as indicators of compromise), it is a complex activity which requires a specific skillset. This guide describes the core capabilities required to start a threat intelligence function and malware analysis is not required as part of that core capability, however there is no reason why a mature capability could not be bolstered by expertise in this area.

12.5 Models for CTI Organisational Positioning

Below are two possible models for CTI integration into departmental organisational structures. The authors recommend each organisation find the best fit for CTI in line with their own requirements. These structures are given as potential starting points for discussion only. As noted in section 10, every organisation with which we collaborated had a different organisational structure; there is not a wrong answer to this question.

12.5.1 CTI within Department Security

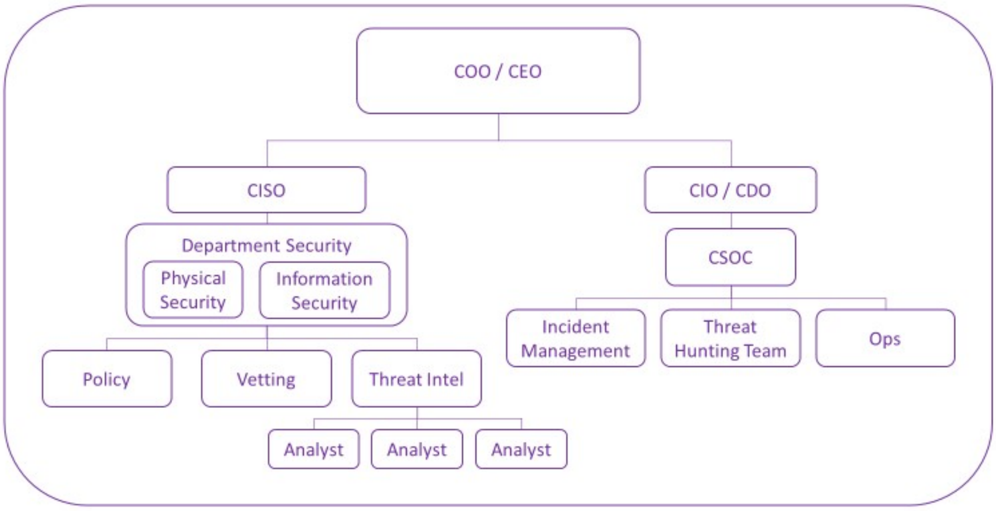

Figure 12 – CTI within Department Security

This structure demonstrates a scenario where the CTI function reports to the Chief Information Security Officer (CISO), rather than the Chief Information Officer (CIO)/Chief Digital Officer (CDO). Whilst CTI fundamentally requires a technical capability, this organisational structure may be used when its objectives are more aligned to those of the departmental security function, rather than technical intelligence.

12.5.2 CTI Within CSOC

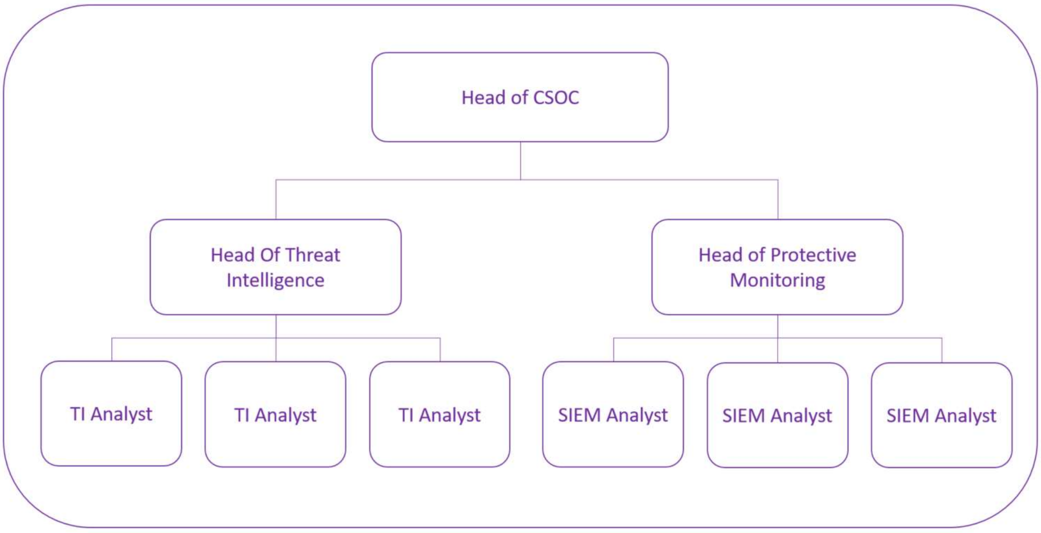

Figure 13 – CTI within CSOC

This structure shows a structure where the CTI function reports to the Head of CSOC. This structure is best used where the priorities of the CTI function are technically focussed, and requires extensive integration with existing CSOC functions.

The following should be noted as part of this approach:

- SIEM analysts and threat hunters are not the same individuals as those performing CTI. Whilst a shared role structure is favoured by some departments, it is recommended that these roles are kept separate. This is to ensure that the terms of reference for each role are correctly adhered to, and prevents analysts prioritising activity based upon personal interest rather than business requirements
- Splitting the reporting lines of each team is deliberate, as this provides a secondary layer of separation. As noted, the CTI function provides information to the defenders of a department; it is not a defensive function in itself. As such, the strategic objectives of the CTI and Protective Monitoring functions will be complementary but distinct
- Using this structure allows for close collaboration between the CTI analysts and SIEM analysts, which enables positive relationships, and should be encouraged wherever possible. In particular those with responsibility for technical intelligence should work closely with CSOC analysts
- The Head of CTI should also attempt, wherever possible, to work with Departmental Security and other CTI Leads within their cluster to mature their capability

12.6 Bibliography

1. **NCSC / Cabinet Office.** *Minimum Cyber Security Standard.* London : publishing.service.gov.uk, 2018. Version 0.1.

2. **MITRE.** MITRE Cyber Attack Lifecycle. *MITRE.* [Online] attack.mitre.org/resources/enterprise-introduction.

3. **Gartner.** Market Guide for Security Threat Intelligence Products and Services. *Gartner.* [Online] https://www.gartner.com/doc.3765965/market-guide-security-threat-intelligence.

4. **FireEye.** Threat Intelligence Use Case Series. *FireEye.* [Online] https://www.fireeye.com/content/dam/fireeye-www/products/pdfs/pf/intel/sb-incident-responder-profile.pdf.

5. **Recorded Future.** Threat Intelligence Use Cases. *Recorded Future.* [Online] https://www.recordedfuture.com/threat-intelligence-use-cases.

6. **CESG.** *Information Risk Management (IS1 & 2).* s.l. : CESG, 2012. Version 4.0.

7. **BAE Systems.** *Intelligence Led Threat Mitigation.* s.l. : BAE Systems, 2017.

8. **Caltagirone, Sergio, Pendergast, Andrew and Betz, Christopher.** *The Diamond Model Of Intrusion Analysis.* s.l. : Active Response, 2013.

9. **Lockheed Martin.** The Cyber Kill Chain. *Lockheed Martin.* [Online] [Cited: 27 09 2018.] https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html.

10. **MISP.** Home Page. *MISP Threat Sharing.* [Online] http://www.misp-project.org/index.html.

11. **OASIS.** Introduction to STIX. *OASIS Open Github.* [Online] https://oasis-open.github.io/cti-documentation/stix/intro.

12. —. APT1 JSON. *OASIS Open Github.* [Online] https://oasis-open.github.io/cti-documentation/examples/example_json/apt1.json.

13. **Mandiant.** FireEye. *Fireeye.com.* [Online] https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf.

14. **Ponemon Institute.** *The Value of Threat Intelligence: The Second Annual Study of North American and United Kingdom Companies.* s.l. : Ponemon Institue, 2017.

15. **MOD.** *Joint Doctrine Publication 2-00 Understanding and Intelligence Support to Joint Operations.* s.l. : MOD UK, 2011. Third Edition.

16. *Innovative Approaches to Cyber Defence.* **Hogue, Dave.** Manchester : NSA, 2018.

17. **FIRST.** TRAFFIC LIGHT PROTOCOL (TLP). *FIRST Standards Definitions and Usage Guidance — Version 1.0.* [Online] FIRST. https://www.first.org/tlp/.

18. **Cabinet Office.** *Government Security Classifications.* London : Cabinet Office, May 2018.

19. **NCSC.** Industry 100. *NCSC.* [Online] https://www.ncsc.gov.uk/information/industry-100.

20. **Intelligence Analyst. *Institute for Apprenticeships.* [Online] 2018. [Cited: 01 10 2018.] https://www.instituteforapprenticeships.org/apprenticeship-standards/intelligence-analyst/.**

21. SANS. FOR578: Cyber Threat Intelligence. *SANS.* [Online] https://uk.sans.org/course/cyber-threat-intelligence.

22. NCSC. National Cyber Security Strategy. [Online] https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/567242/national_cyber_security_strategy_2016.pdf .

23. The Verge. [Online] https://www.theverge.com/2018/4/2/17189880/john-mcafee-bitcoin-cryptocurrency-twitter-ico.

24. Kernkamp, Eric Luiijf and Allard. *Sharing Cyber Security Information.* s.l. : Global Conference on Cyber Space 2015, 2015.

25. NCSC. NCSC Risk Management Collection. *NCSC.* [Online] NCSC, 14 December 2017. [Cited: 21 07 2018.] https://www.ncsc.gov.uk/guidance/risk-management-collection.

26. SANS. SANS Threat Intelligence. *SANS.* [Online] https://uk.sans.org/course/cyber-threat-intelligence.

28. MITRE. Adversarial Tactics, Techniques & Common Knowledge. *Mitre Partnership Netwok.* [Online] [Cited: 27 09 2018.] https://attack.mitre.org/wiki/Main_Page.

29. BBC. Marcus Hutchins Expo. *BBC News.* [Online] https://www.bbc.co.uk/news/av/world-middle-east-39907885/cyber-attack-hero-malware-tech-expert-and-the-kill-switch.

30. Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX™). *STIX Project.* [Online] http://stixproject.github.io/getting-started/whitepaper/.

31. MWR. *Threat Intelligence: Collecting, Analysing, Evaluating.* s.l. : CPNI, 2015.

32. Active Response. [Online] http://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf.

33. [Online] https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html.

34. ThreatConnect. ThreatConnect - STIX-TAXII. *ThreatConnect.* [Online] https://threatconnect.com/stix-taxii/.

35. Wikipedia. Confirmation Bias. [Online] https://en.wikipedia.org/wiki/Confirmation_bias.

36. —. Unconcious Bias Training. *Wikipedia.* [Online] https://en.wikipedia.org/wiki/Unconscious_bias_training.

37. MISP Project. MISP. *MISP.* [Online] http://www.misp-project.org/.

38. VirusTotal. How it works. *VirusTotal.* [Online] https://support.virustotal.com/hc/en-us/articles/115002126889-How-it-works.