



Detecting the Unknown: A Guide to Threat Hunting

v2.0
March 2019



CYBER SECURITY
PROGRAMME



Executive Summary

The National Cyber Security Strategy 2016-2021 details the UK government's investment in cyber security, with the vision for 2021 that the UK will be secure and resilient to cyber threats while prosperous and confident in the digital world. To achieve this, government departments are currently investing in improvements to their own cyber security to meet the Minimum Cyber Security Standard (MCSS), published by the Cabinet Office in June 2018; however, departments should take the opportunity to start investing in the mobilisation and development of their Threat Hunting capabilities.

Threat Hunting, often referred to as *Incident Response without the Incident*, is an emergent activity that comprises the proactive, iterative, and human-centric identification of cyber threats that are internal to an Information Technology network and have evaded existing security controls. Departments that operate a Threat Hunting capability will improve their security posture and hence reduce risk, as malicious activity can be identified earlier on in an attack, thereby minimising the opportunity for adversaries to disrupt, damage or steal.

Departments must create an enabling environment for their Threat Hunting function, by providing enablers such as Cyber Threat Intelligence, relevant data from across the estate, and appropriate investment in people, processes and tools. A joined-up approach to Threat Hunting should be taken across HM Government, where collaboration ensures that the improvements to our collective cyber security from Threat Hunting are greater than that of each department's own efforts, while helping to develop the next generation of the UK's defenders.

This guide, produced via a literature 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 a capability and a profession.

This guide's key findings are:

- Operate a SOC-based Threat Hunting capability to reduce risk, via the appointment of a Threat Hunting Lead, implementation of a formalised process such as our Extended Hunting Loop, and adoption of our Capability Maturity Model to aid development
- Enable the Threat Hunting function to improve the Return on Security Investment, via adoption of a standardised framework such as MITRE's ATT&CK™ for Enterprise, by appropriately investing in the development of people, and by providing essential data visibility
- Leverage HM Government to develop the Threat Hunter role by collaborating between departments, setting common standards for departments and suppliers, and collectively developing the Threat Hunting profession

Table of Contents

1	Introduction.....	6
2	An Overview of Threat Hunting.....	8
3	Operating a Threat Hunting Capability.....	10
3.1	Capability Maturity Model.....	10
3.2	People	14
3.3	Process.....	17
3.4	Tools.....	23
4	Creating an Enabling Environment.....	25
4.1	Utilising Cyber Threat Intelligence for Threat Hunting	25
4.2	Improving Data Visibility with MITRE's ATT&CK	27
4.3	Investment Priorities	32
4.4	Metrics	34
5	Leveraging HM Government	35
5.1	Collaboration	35
5.2	Setting the Standard	37
5.3	Development of the Profession	40
6	Conclusion	41
6.1	Operate a SOC-based Threat Hunting Capability to Reduce Risk	41
6.2	Enable the Threat Hunting Function to Improve the ROSI	42
6.3	Leverage HM Government to Develop the Threat Hunter Role	42
7	Appendices	44
7.1	Appendix I – Contributors	44
7.2	Appendix II – NCSP Funded Publications	45
7.3	Appendix III – Bibliography	47

List of Figures

Figure 1 – The SANS Sliding Scale of Cyber Security	6
Figure 2 – The Lockheed Martin Cyber Kill Chain®	8
Figure 3 – The Pyramid of Pain	9
Figure 4 – Threat Hunting Capability Maturity Model	13
Figure 5 – The Hunting Loop	17
Figure 6 – The Diamond Model	19
Figure 7 – The Extended Hunting Loop	22
Figure 8 – Hunt Tracking	24
Figure 9 – The Threat Intelligence Lifecycle	25
Figure 10 – MITRE’s Cyber Attack Lifecycle and ATT&CK Matrix for Enterprise*	29
Figure 11 – Capability Scope Comparison	46

List of Tables

Table 1 – Adversary Tactics	27
Table 2 – Data Sources	30
Table 3 – Example Threat Hunting Metrics	34

1 Introduction

The National Cyber Security Strategy (NCSS) 2016-2021 (1) details the UK government's investment in cyber security, with the vision for 2021 that the UK will be secure and resilient to cyber threats while prosperous and confident in the digital world. To achieve this, the UK needs to: defend against cyber threats and respond to incidents; deter hostile action in cyberspace; develop the cyber security industry and talent required to overcome future threats and challenges; and pursue international action to shape cyberspace. This is underpinned by the creation of the National Cyber Security Centre (NCSC) (2) in 2016 to act as the authority on the UK's cyber security, as well as investment of £1.9 billion over the five years of the NCSS as laid out in the Strategic Defence and Security Review 2015 (3).

Government departments are currently investing in improvements to their own cyber security to meet the Minimum Cyber Security Standard (MCSS) (4) published by Cabinet Office in June 2018; this is the cyber security baseline that the government expects departments to adhere to and exceed wherever possible. Much of this investment is into the development of Security Operations Centres (SOCs) that are responsible for: detecting and responding to threats; increasing resilience; identifying and addressing negligent or criminal behaviours; and deriving business intelligence about user behaviours (5). Investment is also being targeted at developing Cyber Threat Intelligence (CTI) capabilities that can provide the organisation with actionable (i.e. accurate, relevant and timely) intelligence on threat actor's targets, motivations, infrastructure, and capabilities.

To complicate development of these SOCs and CTI functions, many of the departments have complex and ageing Information Technology (IT) estates with high levels of technical debt. For example, investment is often sunk into legacy systems to extend their lifespan, or sub-optimal architectural design choices are made to work around the constraints of these systems

SOC activities and CTI functions fall into the Active Defence and Intelligence categories respectively, of SANS' Sliding Scale of Cyber Security (6). Shown in Figure 1, the Sliding Scale is a model to visualise the continuum of actions and investments that contribute to cyber security. Generally, investment should be prioritised starting on the left of the scale, before moving along to the right.

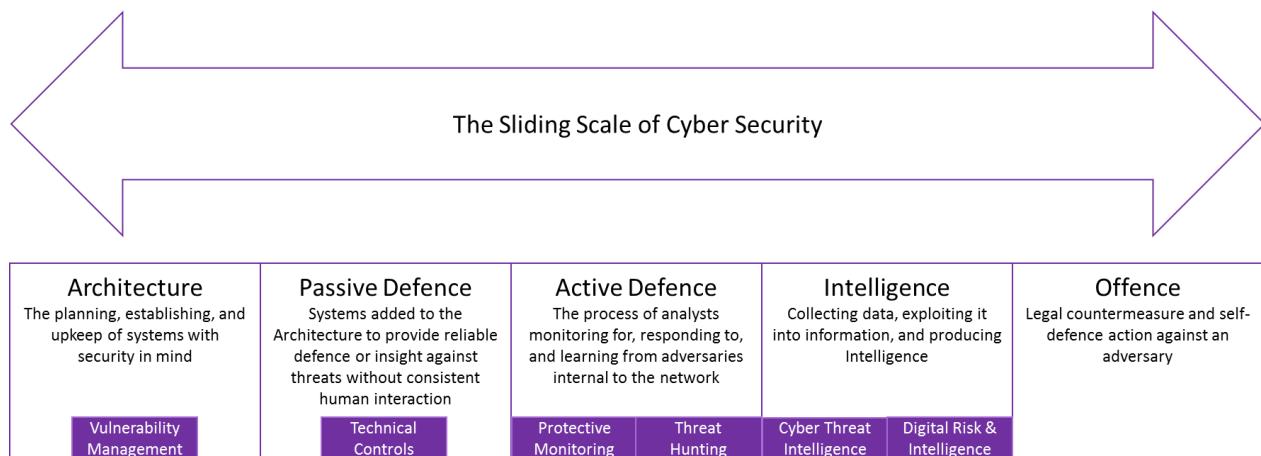


Figure 1 – The SANS Sliding Scale of Cyber Security

Threat Hunting, often described as *Incident Response without the Incident*, sits within the Active Defence phase of the Sliding Scale. As Threat Hunting is an Active Defence, departments first need to sufficiently mature their Architecture (e.g. Vulnerability Management), Passive Defence (e.g. Technical Controls such as firewalls) and other Active Defences (e.g. Protective Monitoring), while operating a mature Intelligence capability will add value to and enable Threat Hunting.

Many SOC analysts already actively search for threats within their network, albeit often in an unstructured and informal manner, but according to the SANS 2017 Threat Hunting Survey (7), only 35.3% of the 306 organisations surveyed (of which 14.4% were government) hunted on a continuous basis. This rose to 43.2% of the 600 organisations surveyed in the SANS 2018 Threat Hunting Survey (8). Additionally, in the 2017 survey, only 4.6% of respondents were using externally published guidance, suggesting little existed in terms of industry good practice for Threat Hunting. Structured Threat Hunting performed on a frequent basis provides an effective means of reducing risk across the organisation, and this report proposes a Capability Maturity Model (CMM) to track and prioritise areas for development.

This guide has been produced via a literature review and engagements with public and private sector organisations (listed in Appendix I) and will outline good practice Threat Hunting for government departments, to aid with the mobilisation and subsequent development of their Threat Hunting capabilities. Our recommendations are targeted for SOCs, departments, and HM Government as a whole, with the report structured as follows:

- **Section 2** provides an overview of Threat Hunting as a capability and introduces key concepts for the remainder of the report
- **Section 3** is targeted for **security managers**, such as heads of SOCs, and outlines the capability required from people, processes and tools for a government department's SOC to operate a basic but competent Threat Hunting capability and hence reduce risk
- **Section 4** is targeted at **security executives**, such as Chief Information Security Officers (CISOs), and outlines the activities that a government department should take at the enterprise-level to enable their Threat Hunting function and hence improve its efficiency and effectiveness
- **Section 5** is targeted at **decision makers within cross-government functions**, and outlines steps that can be taken across HM Government to improve collaboration, set a common baseline, and professionalise the threat hunter role and hence improve our collective security

2 An Overview of Threat Hunting

Based on our research and engagements, we have defined Threat Hunting as:

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

When performed by skilled people who are equipped with the necessary enablers, processes and tools, departments that implement a Threat Hunting capability will be better able to identify and remediate threats, therefore improving their security posture and risk profile.

Proactivity is the key difference between Threat Hunting and other security activities such as Protective Monitoring. SOC analysts tend to take a reactive approach, responding to alerts raised by Security Information and Event Management (SIEM) solutions, before following set triage workflows and then handing over to the appropriate team for remediation e.g. Incident Response (IR) or Vulnerability Management. Threat hunters, on the other hand, are driven by their own curiosity and intuition to hypothesise about potential threats. These hypotheses are then tested within their network, with threat hunters pivoting off each discovery, following wherever their investigation takes them (within their scope). Detailed further in Section 3.3, this process of hypothesis generation and testing is iterated: if a hunt hypothesis is not proven true, then the hunters can move on to test new hypotheses; if the hunt hypothesis is proven true, then the IR process takes over to contain and remediate the threat. Following that, details of any novel adversary activity discovered should be provided to the CTI team, while the successful hunting procedure itself should then be automated or codified for future use, for example as a SIEM detection rule. An unproven hypothesis does not necessarily confirm the absence of malicious activity; instead, further data or analytic functionality may be required, so relevant hypotheses should be re-tested as the capability matures.

The Lockheed Martin Cyber Kill Chain® (9), as depicted in Figure 2, is a representation of the phases of a cyber-attack, starting with reconnaissance of the target e.g. by analysing their digital footprint, and resulting with actions on objectives, e.g. stealing, disrupting or destroying assets, which may result in financial and/or reputational damage to the target. By adopting a proactive approach, adversaries can be detected from the delivery phase of the Cyber Kill Chain onwards, as this is the point that the network has been breached. Additionally, reactive Protective Monitoring capabilities can only identify “known known”¹ threats, while proactive Threat Hunting capabilities can identify “known unknown” and “unknown unknown” threats posed by Advanced Persistent Threats (APTs).



Figure 2 – The Lockheed Martin Cyber Kill Chain®

¹https://en.wikipedia.org/wiki/There_are_known_knowns. A phrase used by former United States Secretary of Defence Donald Rumsfeld during a news briefing on 12 February 2002.

FireEye's M-Trends 2018 (31)

- Global median dwell time in 2017 was 101 days
- Range of less than 7 days to over 2,000 days
- Median is 175 days within Europe, the Middle East and Africa (EMEA)
- Only 62% of intrusions initially detected by the organisations themselves

Indicators of Compromise (IOCs) are the artefacts that if detected on a network suggest malicious activity has occurred – these are the “known knowns”. Tactics, Techniques and Procedures (TTPs) are the abstract descriptions of adversary behaviour that IOCs indicate, so are the “known unknowns”, as the methodology is understood, but any subsequent IOCs are not known prior to the attack. Zero-day exploits² would therefore be the “unknown unknowns”.

Detection of adversary activity earlier in the Cyber Kill Chain can be tracked as dwell time (time from compromise to detection), which is a key metric for Threat Hunting. Shorter dwell times reduce the possibility of adversaries damaging the Confidentiality, Integrity or Availability (CIA Triad) of the organisation’s information systems.

Most automated network and endpoint security controls utilise signature and rule-based alerting for IOCs, which only detect the “known knowns” such as previously described malware hashes for example; however, while easy to detect, it is trivial for adversaries to overcome (e.g. by changing a single bit in the malware binary file). Targeting TTPs is significantly more difficult, but very tough for adversaries to overcome as it would require them to use an entirely different methodology. This concept of ease of detection vs. difficulty caused for the adversary can be represented in the Pyramid of Pain (POP), as defined by David Bianco (10), shown in Figure 3. Whilst hunters can benefit from leveraging automation during their hunts, the reliance of current technology on rules and signatures means it is not possible to fully automate Threat Hunting.

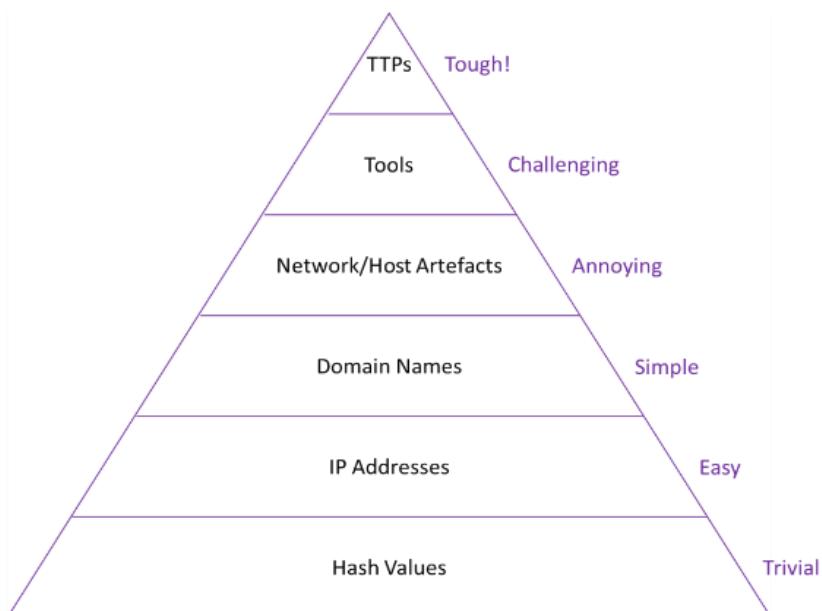


Figure 3 – The Pyramid of Pain

²[https://en.wikipedia.org/wiki/Zero-day_\(computing\)](https://en.wikipedia.org/wiki/Zero-day_(computing)). A zero-day vulnerability is a computer-software vulnerability that is unknown to those who would be interested in mitigating the vulnerability.

3 Operating a Threat Hunting Capability

This section is targeted for security managers, such as heads of SOCs, and outlines the capability required from people, processes and tools for a government department's SOC to operate a basic but competent Threat Hunting capability and hence reduce risk. The main pre-requisites are CTI ingestion, high data visibility, and appropriate investment, and these are detailed in Section 4.

3.1 Capability Maturity Model

Our Threat Hunting Capability Maturity Model is shown in Figure 4 and comprises five levels of maturity from Level 1 (Initial) to Level 5 (Optimising); these are then broken down into the sub-capabilities of People, Process and Tools. Organisations that are focused on reaction (e.g. Protective Monitoring) are represented at Level 1, whereas a HM Government department that operates a basic but competent proactive Threat Hunting capability is represented at Level 3.

This CMM can be used to assess the maturity of an organisation's Threat Hunting capability at a given point in time and aid the prioritisation of efforts to continuously improve. However, it should not be used to compare maturity between organisations, as each will have its own unique circumstances and context.³

RECOMMENDATION 1: To aid assessment of organisational performance and identify areas for improvement, adopt a standard framework such as our Threat Hunting Capability Maturity Model.

3.1.1 Level 1 – Initial

Level 1 describes an organisation that performs little or no Threat Hunting, and instead has a reactive stance, relying on alerts generated by SIEM tools and other defensive infrastructure. Threat Hunting occurs rarely, if at all, and is ad-hoc and basic; it is performed by existing staff e.g. SOC analysts, and on their own initiative. Normal systems behaviour is not well understood.

As Threat Hunting is not implemented as a formal capability, little exists in terms of recruitment or training plans, performance management, or career development. Hypothesis generation is unstructured, and little or no documentation is recorded. Data visibility is minimal, with a lack of understanding of the current data and the subsequent gaps in coverage.

3.1.2 Level 2 – Managed

At Level 2, steps have been taken to start implementing a proactive Threat Hunting capability. Existing staff are occasionally led on hunts by a dedicated and experienced Threat Hunting lead, with the focus on targeting IOCs at the bottom of the POP. Normal systems behaviour is moderately understood, and tactical CTI (as defined in Section 4.1) is ingested and automatically matched against logs.

³ <https://www.ncsc.gov.uk/blog-post/maturity-models-cyber-security-whats-happening-iamm>

Recruitment, training, performance and career development are all informally managed. Hypotheses are prioritised by the lead and only basic documentation is recorded, using standard office suite functionality (e.g. Microsoft Word and Excel). Data visibility is moderate for key areas, with a basic understanding of the data available.

3.1.3 Level 3 – Defined

Level 3 is the minimum level required for a HM Government department to operate a competent Threat Hunting capability and start realising benefits. A team of dedicated hunters, led by the Threat Hunting lead, follow a formal Threat Hunting process and hunt on a frequent schedule, with the focus on targeting IOCs in the middle of the POP, using techniques such as statistical analysis. Normal systems behaviour is adequately understood for key systems to allow identification of abnormal activity.

Plans for recruitment, training and career development are all formally documented, with performance expectations defined. Hypothesis and hunt information is recorded in a central knowledge repository, and workflow management tools are used to track workloads and progression. Data visibility includes key sources and types. Successful hunt procedures are automated, where possible, while identified IOCs are provided to the CTI and Protective Monitoring functions for the development of the subsequent SIEM detection rules.

3.1.4 Level 4 – Quantitatively Managed

At Level 4, the Threat Hunting capability is well established, and utilises quantitative metrics to improve performance and show benefit. The Threat Hunting team is supplemented by SOC analysts on a rotational basis, both to increase the resources available to hunt, but also to develop and motivate the wider SOC staff. Hunting is very frequent, and targets IOCs at the top of the POP (i.e. adversary TTPs), while making use of visualisation techniques. Mission critical systems are identified, contributing towards the hunters understanding of the organisational context and therefore starting to develop their situational awareness.

At this level, succession plans are in place for key roles, and performance is tracked at a team level using metrics. Manual risk scoring techniques e.g. Crown Jewel Analysis (CJA) are used to aid hypothesis generation, and dashboards are utilised to aid collaboration and reporting. Data visibility is moderate across all relevant areas of the estate, and there is a good level of understanding.

3.1.5 Level 5 – Optimising

Level 5 represents an advanced Threat Hunting capability that should be aspired to. At this level, the Threat Hunting team is fully integrated into the wider SOC and organisation in terms of resourcing, recruitment, training and performance, with action plans created to mitigate any underperformance. Threat hunters have extensive experience, possess situational awareness, and are demonstrably valued by the organisation. The capability is forward-looking and can adapt quickly to changes in the environment.

Automated risk scoring is leveraged using machine learning, with horizon scanning maintained for future technological developments. Hunts are occurring continuously, with successful analytics and discovered IOCs shared across the community, while the knowledge repository and workflows are integrated with the wider SOC. Data visibility is high across all relevant areas of the estate and is very well understood.

The Threat Hunting team is recognised as a great place to work and is seen as a leader in the field by other organisations.

Threat Hunting Capability Maturity Model	Level 1 INITIAL	Level 2 MANAGED	Level 3 DEFINED	Level 4 QUANTITATIVELY MANAGED	Level 5 OPTIMISING
People 	<ul style="list-style-type: none"> Existing SOC analysts Resourcing needs not known Training needs not known Performance not managed Lack of career development plan Normal systems behaviour not sufficiently understood 	<ul style="list-style-type: none"> Threat Hunting lead Informal view of resourcing Informal view of training Performance is qualitatively managed Career development informally managed Normal systems behaviour is moderately understood 	<ul style="list-style-type: none"> Dedicated threat hunters Formal recruitment plan Formal training plan Performance expectations defined with role profiles Formalised career development plan Normal systems behaviour is fully understood 	<ul style="list-style-type: none"> SOC analysts rotated for L&D Succession plans in place Training completion tracked Metrics utilised for team performance Mission critical systems identified 	<ul style="list-style-type: none"> Teams integrated across SOC Resourcing needs integrated Training needs integrated Improvement plans to address underperformance Situational awareness
Process 	<ul style="list-style-type: none"> Hypothesis generation is unstructured <i>Hunts occur ad-hoc, if at all</i> <i>Little or no data collected</i> Little understanding of anomalies indicative of malicious activity Abnormalities not routinely searched for 	<ul style="list-style-type: none"> CTI and Domain Expertise used to generate hypotheses and prioritisation by lead Hunts occur occasionally <i>Moderate data collection from key areas</i> <i>Basic threat feeds with IOCs utilised</i> Targeting of IOCs at bottom of POP 	<ul style="list-style-type: none"> Formal hunting process Hunts occur regularly <i>High data collection from key areas</i> <i>CTI and previous experience used to detect malicious activity</i> Targeting of IOCs in middle of POP 	<ul style="list-style-type: none"> Manual risk scoring e.g. Crown Jewels Hunts occur frequently <i>Moderate data collection from most of estate</i> <i>CTI tailored to organisation</i> Targeting of IOCs at top of POP 	<ul style="list-style-type: none"> Automated risk scoring e.g. machine learning Hunts occur continuously <i>High data collection from full estate</i> Hunt analytics and IOCs shared across community Automated TTP and campaign tracking
Tools 	<ul style="list-style-type: none"> <i>Reactive SOC tools</i> Little or no automation Little or no documentation produced 	<ul style="list-style-type: none"> Basic searching via text or SQL-like queries <i>Automatic matching of IOCs</i> Documentation using basic office suites 	<ul style="list-style-type: none"> Statistical analysis techniques Library of hunt procedures automated on regular schedule Central workflow and knowledge repository tools Lab environments used to aid hypothesis generation and testing 	<ul style="list-style-type: none"> Visualisation tools utilised, and analytics tested for effectiveness Library of hunt procedures automated on frequent schedule Dashboards utilised 	<ul style="list-style-type: none"> Machine learning is leveraged, with horizon scanning maintained Library of hunt procedures automated continuously Central workflow and knowledge repository are integrated and shared

Note: Items in *italics* are not strictly part of a Threat Hunting capability, but are essential prerequisites and enablers.

Figure 4 – Threat Hunting Capability Maturity Model

3.2 People

3.2.1 Skills and Experience

The key difference between a SOC analyst (performing Protective Monitoring) and a hunter is that of a proactive approach. However, many skills and traits required for Threat Hunting also overlap with the roles of other defenders such as SOC analyst, CTI analyst, and incident responder. Based on our research and engagements, hunters require the following core skills:

- A mindset of curiosity
- Log analysis and general analytical skills
- Understanding of normal network behaviour
- Understanding of normal endpoint user and application behaviour
- Understanding of the threat landscape and the use of CTI
- System administrator experience across Windows / Linux / common security products

These core security operations skills give the hunters knowledge of the capabilities and limitations of the security controls on their network such as Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS), knowledge of log types and collection sources, and an applied understanding of core information security concepts such as the CIA Triad.

Understanding of CTI and the associated threat landscape adds benefit by allowing hunters to ingest intelligence and focus their effort where the threat is greatest, enabling the generation of good hypotheses to test. Threat hunters that understand the threats they face will also be able to feed requirements back into the CTI function, helping to ensure all received intelligence is actionable. Furthermore, hunters should have self-awareness of common cognitive biases⁴ such as confirmation bias, to prevent effort being wasted or incorrect conclusions being drawn.

At a more advanced level, hunters require Digital Forensic and Incident Response (DFIR) skills and experience that would overlap with those required for IR. These skills allow for complex investigations involving live data, or data held in memory, on endpoints and across the network, rather than simply analysing collected logs. DFIR skills include:

- Endpoint forensics
- Network forensics
- Malware analysis

⁴ https://en.wikipedia.org/wiki/Cognitive_bias. The systematic pattern of deviation from norm or rationality in judgment.

Endpoint forensics, also called computer forensics, follows a process of acquisition, examination, and analysis of the endpoint, before reporting on the facts and providing an opinion of the data; this can aid hunters in finding and documenting evidence of threats on specific endpoints and storage media. If specific guidelines that protect the authenticity and integrity of digital media are followed, then any evidence can be admissible in court if later required.

Network forensics is the analysis of network traffic to identify signs of intrusion within the network, such as artefacts created during lateral movement or data exfiltration activities undertaken by the adversary. While endpoint forensics looks at acquired images, so can be performed reactively, network data is often lost once transmitted so network forensics needs to be performed proactively to ensure the required data is captured. Advanced Persistent Threats (APT) may have the skills to hide evidence of their intrusion on endpoints via destruction or tampering of logs, meaning that network forensics may be the only analysis capable of detecting these threats.

Malware analysis is a highly specialised skill that aims to determine the origin and purpose of an identified instance of malware. This analysis is usually either static, where the binary file is reverse engineered without executing it, or dynamic, where the malware is executed in a sandbox environment to observe its behaviour. Malware analysis is of most benefit when investigating novel malware, as previously identified malware will likely have IOCs available that should be provided to the defensive teams by the CTI function. As a highly specialised skill, it may be necessary to outsource malware analysis to a commercial provider of these services.

Finally, top tier threat hunters would possess situational awareness⁵ that allows them to actively defend against adversaries by perceiving threats and vulnerabilities in context. This is often expressed as a hunch that something “just doesn’t look right” on their network and therefore warrants further investigation. True situational awareness is gained from years of experience and empowers hunters to make timely and accurate decisions.

3.2.2 Resourcing the Capability

A cyber security skills gap currently exists within the UK’s Critical National Infrastructure (CNI), which includes government as one of its thirteen sectors. This gap is described within the Joint Committee on the National Security Strategy’s Second Report of Session 2017-19 (10) as being predominantly caused by a scarcity of individuals who have the required skills, an inability to match highly competitive reward packages offered by the private sector, and a lack of gender diversity that limits the size of the talent pool.

This struggle to recruit cyber security staff was echoed by the organisations we engaged with, particularly for specialist roles such as threat hunters. Staff retention is also an issue faced by many organisations. Research by the Cyentia Institute (11) found that 1 in 4 SOC analysts are dissatisfied with their job, while 1 in 3 are actively looking for other job opportunities. One of the reasons cited was a disconnect between expectations of working in a SOC and the day-to-day reality, with examples such as unclear career paths and tedious or repetitive duties.

⁵ <https://en.wikipedia.org/wiki/Situation Awareness>. The perception of environmental elements and events with respect to time or space, the comprehension of their meaning, and the projection of their future status.

Correspondingly, Protective Monitoring was found to be the most time-consuming activity performed, in contrast to Threat Hunting, which was one of the least commonly performed activities.

As a Threat Hunting capability develops in maturity, an increasingly greater proportion of the SOC analysts will be taking a proactive rather than reactive stance, which will provide the variety and challenge clearly sought by these skilled individuals.

The first step to resourcing a Threat Hunting team and moving from the initial level towards a more managed approach should be the recruitment or training of a dedicated Threat Hunting lead. This lead role is essential in providing direction and technical expertise to other SOC staff that will allow them to start hunting in a structured manner. For example, SOC analysts can hunt on an occasional basis e.g. during any periods of low activity, under the direction and guidance of the lead. As a varied, challenging and enjoyable activity, this will help to improve morale within the SOC, while simultaneously improving the analysts' understanding of the network and adversary behaviour.

RECOMMENDATION 2: Recruit or train a Threat Hunting lead as a responsible individual to enable development of the capability.

After recruiting a Threat Hunting lead, the next step of maturity is to form a team of fully dedicated hunters, who will have the necessary skills and experience to focus on proactive hunting for threats, without the time-consuming distraction of alert triaging. Further benefit is gained by rotating other SOC staff into the Threat Hunting team either on a short-term basis e.g. SOC analysts for a month at a time, or simply when available due to workload. In this manner, Threat Hunting can be used as a tool to train and engage staff. Threat Hunting was ranked as the most enjoyable of 12 common SOC activities and was perceived highly on the level of expertise required, the value gained for the organisation, and the variation within the activity; Protective Monitoring was perceived lower in these fields (11). Another example is utilising incident responders when they are not dealing with an incident, seeing as the skills required for IR and Threat Hunting are similar.

RECOMMENDATION 3: Rotate SOC analysts into the Threat Hunting team for learning and development purposes.

3.3 Process

While skilled threat hunters are key to a successful Threat Hunting capability, it is also important that a formal hunting process is followed to ensure consistency and efficiency across all hunts. A widely accepted approach to the process underlining Threat Hunting is Sqrrl's Hunting Loop (12), which has four stages that define the iterative method to be taken. Efficiently iterating through the loop enable hunts to be quickly automated so that hunters can focus on testing the next hypothesis. The following subsection will provide further detail on the activities within each stage.

3.3.1 Create Hypotheses

Threat Hunting begins with hunters asking questions, such as *“How would a threat actor infiltrate our network?”* These questions then need to be broken down into specific and measurable hypotheses that state what threats may be present in the network and how they can be identified. Hypotheses cannot be generated by tools, and instead must be derived from the hunter's observations based upon CTI, situational awareness, or domain expertise (13).

Hypotheses must also be testable i.e. the hunters must have the required data visibility and tools at their disposal to look for the hypothesised evidence of malicious activity. More data types allow the investigation of more techniques, while more data sources expand the arena in which to hunt. These three main sources of observations will be explored further, and hypotheses can be derived from any combination of them.

Hypotheses generally tend to focus on detecting either a specific threat actor, tool, or technique. Examples for each are shown in the box opposite.

3.3.1.1 Cyber Threat Intelligence

Cyber Threat Intelligence in the form of IOCs and TTPs are vital sources from which hunters can make observations and subsequently form their hypotheses. IOCs can be directly searched for within the network, but this is better suited for automated tools. Threat hunters should instead form hypotheses from the results of these searches, or by the abstraction of information from IOCs rather than the actual IOCs themselves.

Examples of hypotheses that can be drawn from IOCs include the locations that they may be found on the network, or methods that threat actors may use to obfuscate their activities. While IOCs can assist in quickly generating hypotheses, the goal should be to base hypotheses on TTPs with further context provided by assessments of the geopolitical and threat landscapes.

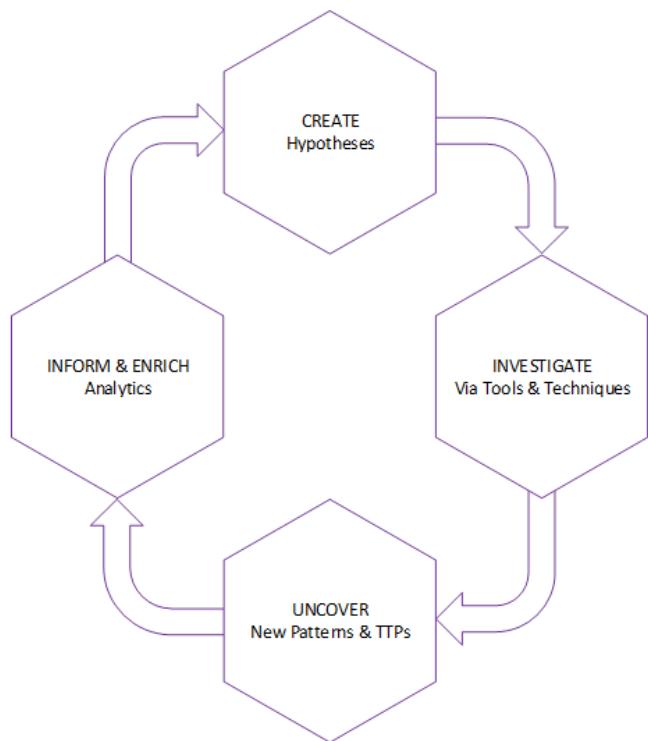


Figure 5 – The Hunting Loop

Example Hypotheses

Threat Actor:

An organisational threat assessment identified Lazarus Group as a high priority threat. Techniques attributed to this threat actor are detailed within MITRE's ATT&CK Navigator.

We therefore hypothesis that if this threat actor is present in our network, we would be able to detect evidence of multiple techniques being deployed, in a manner consistent with their known attack paths.

Tool:

CTI and our situational awareness suggests that our organisation is currently vulnerable to a variant of the WannaCry ransomware, as SMBv1 is still used.

We therefore hypothesis that if our network is infected with WannaCry, we will see an increase in the rate of file renaming.

Technique:

Lateral Movement, via *Exploitation of Remote Services*, can be performed by exploiting vulnerability MS17-10. Specifically, this can be done via the Metasploit framework with a module that uses a Server Message Block (SMB) request of a specific size to attempt compromise.

We therefore hypothesise that we can see evidence of this technique being used by isolating this SMB request in our network logs.

3.3.1.2 Situational Awareness

Situational awareness, as previously defined, is the perception of environmental elements and events with respect to time or space, the comprehension of their meaning, and the projection of their future status. In the context of Threat Hunting, this effectively means having visibility into and understanding of the network so that any significant changes can be quickly identified. Hypotheses

Crown Jewels Analysis

CJA comprises the following steps:

- Identifying the organisation's core missions
- Mapping the missions to the assets and information upon which they are reliant
- Documenting the network assets utilised
- Constructing attack graphs that determine dependencies, analyse attack paths and rate vulnerabilities for severity

can then be generated for the types of activities that threat actors could perform within the network. Automation should be utilised to assist hunters with monitoring assets and data flows via the use of dashboards, reporting and risk scoring, to highlight trends and anomalies in a visual manner.

Crown Jewels Analysis is a process for identifying those cyber assets that are most critical to the accomplishment of an organisation's mission (14), which in turn enables hunters to better focus their hypothesis generation and data collection efforts.

When generating hypotheses from situational awareness or conducting CJA, non-technical assets such as people and business processes should also be considered.

3.3.1.3 Domain Expertise

Domain expertise relates to the hunter's experience – each hunter has a unique background and set of experiences and skills that influence their hypothesis generation, ranging from previous security incidents they've been involved in and learnt from, to anecdotes they've heard from colleagues. A good hunter should have knowledge of both the organisation's network and the threats faced, so domain expertise can be viewed as both CTI and situational awareness in a historic context.

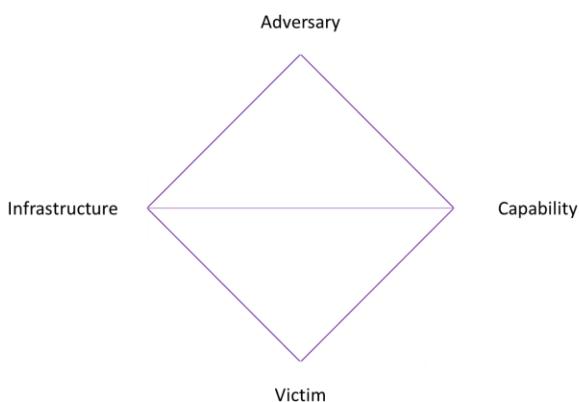


Figure 6 – The Diamond Model

Where previously utilised CTI and situational awareness may not be immediately relevant, it contributed to the hunter's current mindset and capability. However, as previously discussed, hunters also need to be aware of the cognitive biases that can arise from their past experiences to ensure that good decisions are made, and accurate conclusions drawn. The use of models, such as the Diamond Model of Intrusion Analysis (15), shown in Figure 6, aids hunters in structuring data to help overcome biases.

As the hunters conduct hunts and develop their skills, they should ensure the knowledge gained and lessons identified are appropriately documented in a centralised repository, described further in Section 3.4.2, so that this information is available across the function for other hunters to consume and learn from. This can also be shared across the government community to improve the collective security.

3.3.2 Investigate via Tools and Techniques

Once observations have led to hypotheses being generated, these then need to be tested using all the relevant tools and techniques at the threat hunters' disposal. Data visibility should be maximised by increasing the collection coverage across the estate into a centralised repository, and the data types collected should include IDS/IPS logs, Authentication, Authorisation and Auditing (AAA) logs, Domain Name Server (DNS) logs, network traffic flow, endpoint logs and SIEM alerts. Note, it may not be feasible nor desirable to centrally collect all data, and in those cases the threat hunters will need to directly access the relevant devices. Data sources will be discussed further in Section 4.2.

Existing SOC tooling, such as a SIEM platform, can be used to query the data, from basic searching to more advanced data science techniques, while visualisation can aid threat hunters in identifying anomalies and anomalous patterns. Linked Data (16) is a method for publishing structured data so that it can become interlinked and searched using semantic queries⁶. Both raw and Linked Data Analysis techniques should be used to identify patterns across disparate data sets to aid detection of adversary activity.

⁶ https://en.wikipedia.org/wiki/Semantic_query. Semantic queries allow for queries and analytics of associative and contextual nature.

Hunters can also make use of lab environments to aid the testing of hypotheses. This allows the hunters to emulate adversaries and use their tools and techniques to better understand how they can detect them - which would of course be potentially damaging on production systems with live data. This lab environment can also be used for the learning and development of junior analysts within the SOC.

Threat Hunting Techniques

Searching is the most basic method of querying the collected data. The search criteria should be specific enough so that the results returned are not unmanageable, while also general enough so that no adversary activities are missed. Characters such as wildcards (*) can be utilised within queries as required.

Clustering is a form of statistical analysis that separates groups (clusters) of similar data points from a larger set based on specific characteristics, whereas **grouping** identifies when multiple unique data points appear together based on specific criteria, for example, multiple events occurring in a specific time window. The main difference is that grouping requires an explicit set of data points as input. Both are useful for outlier detection.

Stack counting, or stacking, is the application of frequency analysis to large sets of data to identify outliers. FireEye provide a good overview of the technique with worked examples (32).

Machine learning uses algorithms and statistical models to progressively improve performance of a specific task; for Hunting, that is identifying anomalous data that could indicate adversary activities. In supervised machine learning, a set of training data is fed into the algorithm with each data point labelled with the desired output e.g. both normal and anomalous data labelled as such. Unsupervised machine learning is provided with unlabelled data, so the algorithm uses techniques like clustering and grouping to categorise the outputs instead.

3.3.3 Uncover New Patterns and TTPs

The outcome of testing a hypothesis is that the presence of malicious activity is either proven, or not proven. If not proven, this does not necessarily mean that no malicious activity is present – instead, current data visibility and tooling may not allow the required investigation to take place, in which case improving the capability should be a priority so that the hypothesis can later be re-visited. Alternatively, the threat hunter may simply have not recognised any anomalies within the data that indicated the presence of malicious activity.

When a hypothesis is proven, and malicious activity is identified, the Computer Security Incident Response Team (CSIRT) should be notified and the incident management process takes over. At this point, assuming the two teams are different entities, the Threat Hunting team would assist the CSIRT with their investigation; once remediated, the threat hunters can then move onto refining and automating their successful procedure, such as with new detection analytics for a SIEM platform.

Whether or not the hypothesis is proven, non-malicious but suspicious or risky configurations or behaviours may be identified, such as unpatched or misconfigured systems, or logging blind spots. This information can be passed onto the relevant teams e.g. Vulnerability Management for remediation.

3.3.4 Inform and Enrich Analytics

Where possible, successful hunts should be automated to maximise the efficient use of the Threat Hunting team's time and to limit them from continuously repeating the same hunts. This can be done in many ways, such scheduling a saved search, developing a new analytic within existing tools, or providing feedback to a supervised machine learning algorithm. Enrichment may take the simpler form of just providing a new IOC for matching, or by writing a new SIEM rule for reactive detection. The quicker that a hunt can be automated, the less repetition would be required of the threat hunters, and the quicker their curiosity and skills can be put towards testing new hypotheses.

Care should be taken to ensure that any automated hunts are reliable and continue to add value. Once automated, each analytic should be tested for its accuracy and precision, which can be done in several ways, such as a red team⁷ performing the technique in question and checking that the analytic reliably detects their activity.

The analytics, once live, should be monitored for any issues for a limited period with the hunters on-hand to support, before being formally handed over to the Protective Monitoring team to own. However, the analytics should be assessed periodically to ensure they still add value and are relevant to the organisation e.g. whether any changes to the organisational architecture means that some analytics are testing for TTPs that are no longer possible.

RECOMMENDATION 4: Periodically assess the relevance of automated analytics to the organisation.

3.3.5 The Extended Hunting Loop

Sqrrl's Hunting Loop provides an excellent process for individual hunters to follow; however, it does not include a few key areas that are crucial for a Threat Hunting capability to develop. Our Extended Hunting Loop in Figure 7 builds on and adapts the original model by additionally including:

- Inputs for Threat Hunting, such as the observations required for hypothesis generation
- The activities performed by the Threat Hunting lead, such as prioritising hypotheses for testing and making decisions on resourcing
- Use of workflow management and knowledge repository tools
- Additional outcomes of hypothesis testing i.e. hypothesis not unproven, and non-malicious but suspicious/risky behaviour identified
- Handover to CSIRT when malicious activity is detected

RECOMMENDATION 5: Adopt a formalised process, such as our Extended Hunting Loop, to aid operationalisation of mature Threat Hunting processes.

⁷ https://en.wikipedia.org/wiki/Red_team. A red team is an independent group that challenges an organisation to improve its effectiveness by assuming an adversarial role or point of view.

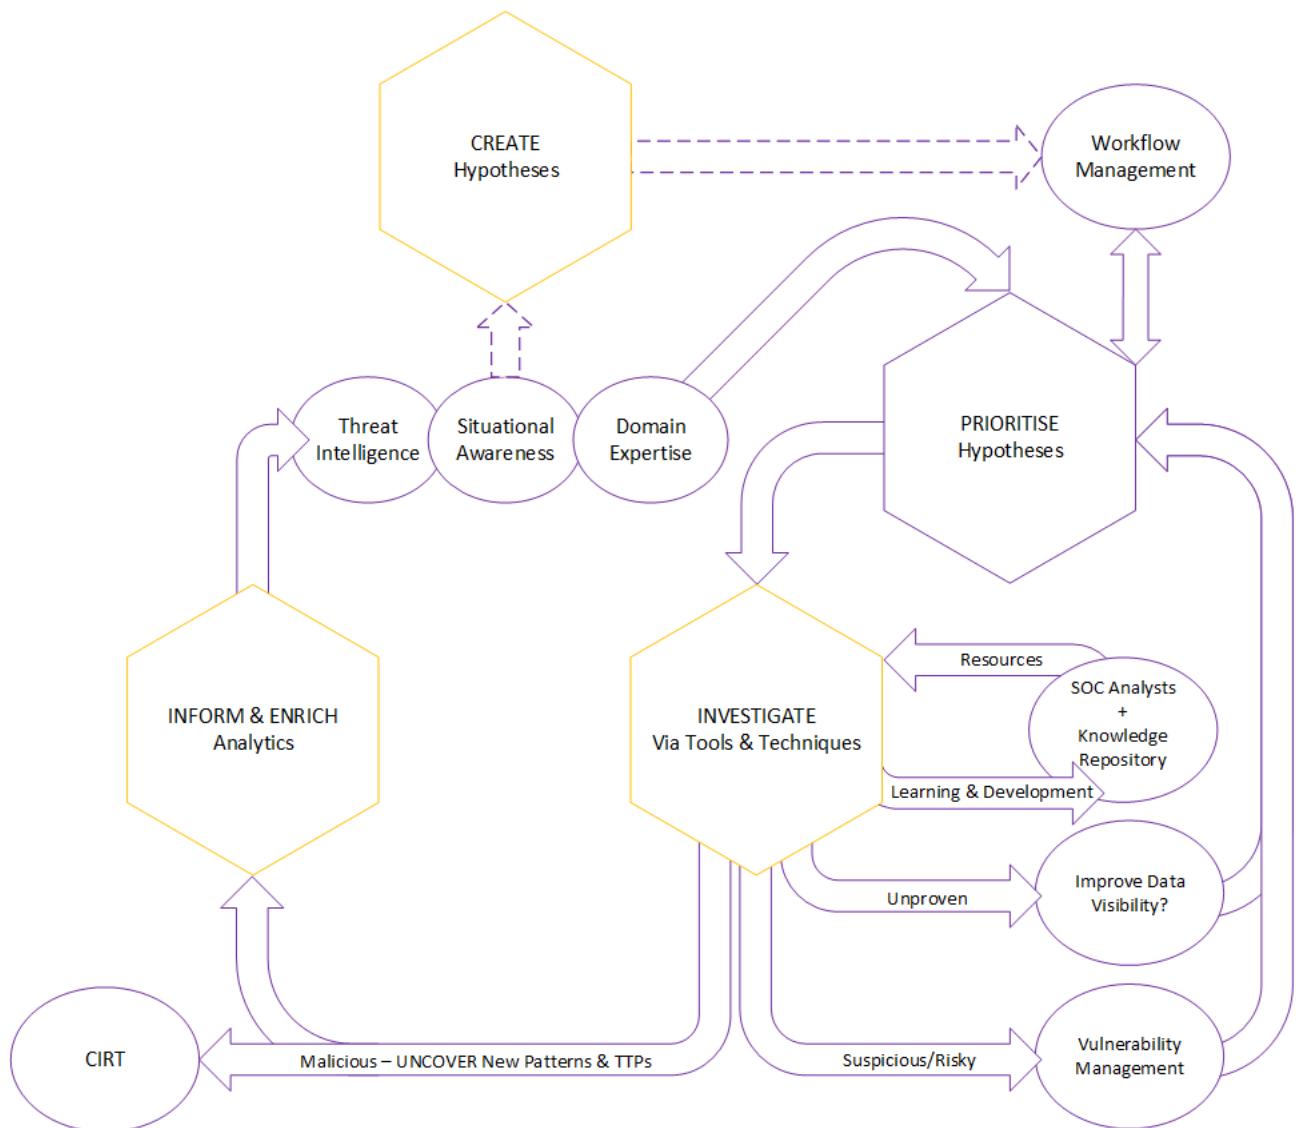


Figure 7 – The Extended Hunting Loop

3.4 Tools

While skilled people and effective processes are the critical factors for a successful Threat Hunting capability, tooling is of course still required to collect and interrogate data, automate analytics, and work collaboratively.

3.4.1 Authoring of Analytics

Existing tooling within the SOC, such as a SIEM platform or technical controls, can be utilised and leveraged by the threat hunters to identify malicious activity, and successful analytics may then be automated or used to enrich reactive monitoring, as previously discussed.

3.4.2 Collaborative Working

Many organisations we engaged with spoke of the benefit of maintaining documentation end-to-end throughout the Threat Hunting process, as this improves consistency and efficiency within the team. This documentation should be stored in a manner accessible to the entire Threat Hunting team, as this will assist knowledge transfer and enable different hunters to pick up where another left off if needed. A tool that provides a workflow management view should be utilised in unison with a knowledge repository.

One example of a workflow management view is that of Epics and Stories. Epics are bodies of work that can be broken down into specific tasks, which are the Stories. The use of these concepts helps structure workloads, and progress can then be tracked via the use of a Kanban board⁸ for each Story type, as seen in Figure 8 opposite. In the context of Threat Hunting, each Epic could be a tactic from MITRE's Adversarial Tactics, Techniques and Common Knowledge (ATT&CK™) for Enterprise (17) (discussed further in Section 4.2), with Story types of hypotheses and associated hunts. Individual Hypothesis Stories can then be tracked through a Hypothesis Kanban with example stages such as: Initial, for basic hypothesis ideas; Development, for adding detail and assessing the scope and dependencies; Production, for hypotheses ready for testing; and Retired, for hypotheses that are no longer relevant. Hypothesis Stories in Production could then lead to the creation of an associated Hunt Story that progresses through To-Do, In-Progress, and Outcome. Additionally, many of these tools produce visual dashboards to aid reporting, which can be utilised by the Threat Hunting team to track performance and prove their value. Metrics are discussed in Section 4.4.

RECOMMENDATION 6: Utilise a workflow management tool to prioritise and track workload.

Alongside the workflow management view, a collaborative knowledge repository would, for example, allow hunters to share ideas, discuss hunt procedures and challenges, and share analytics. Integration of these tools with other teams will allow effective handover from discovering malicious behaviour, or suspicious/risky activity, to the appropriate teams for remediation.

RECOMMENDATION 7: Utilise a central repository to share knowledge and lessons learnt.

⁸ https://en.wikipedia.org/wiki/Kanban_board. Kanban boards visually depict work at various stages of a process using cards to represent work items and columns to represent each stage of the process.

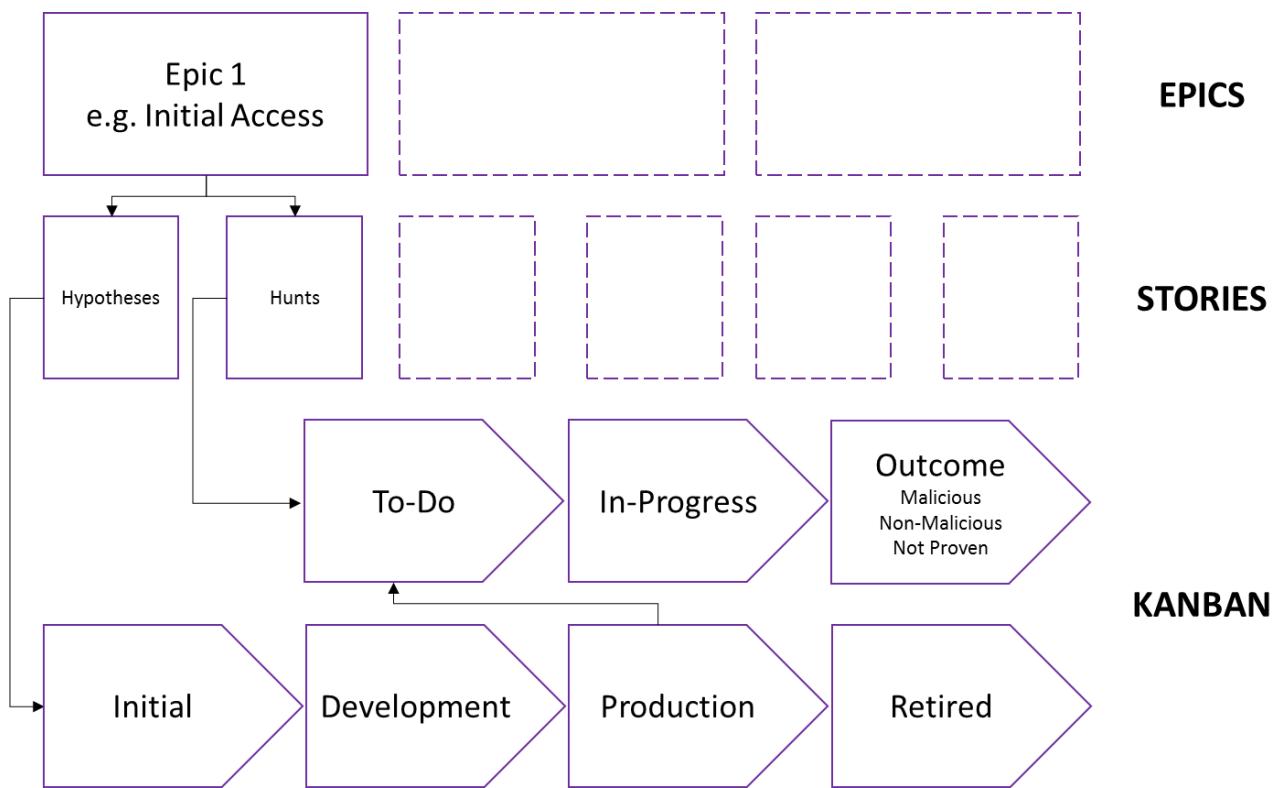


Figure 8 – Hunt Tracking

Note that the process described in Figure 8 closely relates with the rules and analytics development process that supports reactive monitoring solutions. Care should be taken to liaise with the teams developing these rules and use cases to minimise any overlap of effort.

4 Creating an Enabling Environment

This section is targeted at security executives, such as CISOs, and outlines the activities that a government department should take at the enterprise-level to enable their Threat Hunting function and hence improve its efficiency and effectiveness.

4.1 Utilising Cyber Threat Intelligence for Threat Hunting

As one of the key enablers for generating Threat Hunting hypotheses, the ability to consume CTI is essential for the Threat Hunting team. CTI is produced via a five-phase Threat Intelligence Lifecycle:

- **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
- **Processing** refers to the actions 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. Feedback from customers then contributes towards setting the direction.

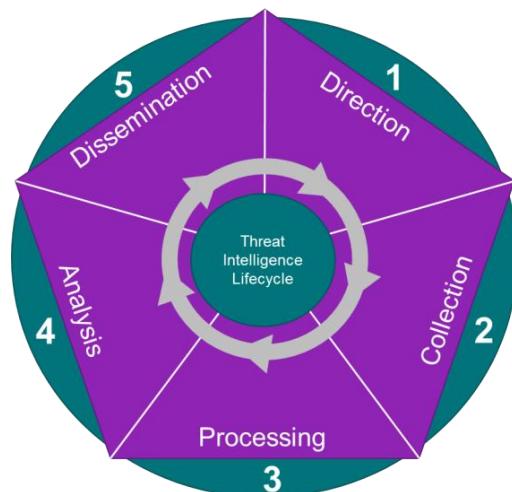


Figure 9 – The Threat Intelligence Lifecycle

While the Threat Hunting team may have little involvement in the Collection, Processing and Analysis phases (as these would fall into the responsibilities of the CTI team), they should be involved in the Direction and Dissemination phases; as customers of the CTI function the Threat Hunting team should provide them with direction and requirements to ensure that the intelligence received is actionable.

CTI falls into three categories, these being Strategic, Operational and Tactical. Strategic CTI is high-level and normally details threat trends or campaigns at a geo-political level. Operational CTI details specific threat actors and their TTPs, while tactical CTI is more technical in nature and consists of IOCs. Threat hunters will benefit most from operational CTI, as the detail on adversaries' TTPs and tools will enable the hunters to generate hypotheses. Strategic CTI is aimed at assisting business decisions, while tactical CTI ingestion should be automated by a Threat Intelligence Platform (TIP) and subsequently matched against logs to detect known attacks.

A good source of operational CTI is the *Alerts and Advisories* group on the Cyber Security Information Sharing Partnership (CiSP), discussed in Section 5.1.2, where organisations share knowledge about TTPs they are currently observing; this in turn aids hunters to generate associated hypotheses relevant to their organisation.

RECOMMENDATION 8: Provide threat hunters with the actionable Cyber Threat Intelligence they require to generate relevant and testable hypotheses.

4.2 Improving Data Visibility with MITRE's ATT&CK

At the most basic level, Threat Hunting needs people to conduct the Threat Hunting, and data for them to hunt through. Because of this, data visibility is a key enabler for any successful Threat Hunting capability. MITRE's ATT&CK for Enterprise provides a framework that describes the methodologies used by adversaries during a cyber-attack. It is represented in a matrix consisting of 11 tactics, with each tactic category containing a list of the associated techniques, as seen in Figure 10. The tactics and their descriptions are listed in Table 1.

Tactic Name	Tactic Description
Initial Access	The initial access tactic represents the vectors adversaries use to gain an initial foothold within a network.
Execution	The execution tactic represents techniques that result in execution of adversary-controlled code on a local or remote system. This tactic is often used in conjunction with initial access as the means of executing code once access is obtained, and lateral movement to expand access to remote systems on a network.
Persistence	Persistence is any access, action, or configuration change to a system that gives an adversary a persistent presence on that system. Adversaries will often need to maintain access to systems through interruptions such as system restarts, loss of credentials, or other failures that would require a remote access tool to restart or alternate backdoor for them to regain access.
Privilege Escalation	Privilege escalation is the result of actions that allows an adversary to obtain a higher level of permissions on a system or network. Certain tools or actions require a higher level of privilege to work and are likely necessary at many points throughout an operation. Adversaries can enter a system with unprivileged access and must take advantage of a system weakness to obtain local administrator or SYSTEM/root level privileges. A user account with administrator-like access can also be used. User accounts with permissions to access specific systems or perform specific functions necessary for adversaries to achieve their objective may also be considered an escalation of privilege.
Defence Evasion	Defence evasion consists of techniques an adversary may use to evade detection or avoid other defences. Sometimes these actions are the same as or variations of techniques in other categories that have the added benefit of subverting a particular Defence or mitigation. Defence evasion may be considered a set of attributes the adversary applies to all other phases of the operation.
Credential Access	Credential access represents techniques resulting in access to or control over system, domain, or service credentials that are used within an enterprise environment. Adversaries will likely attempt to obtain legitimate credentials from users or administrator accounts (local system administrator or domain users with administrator access) to use within the network. This allows the adversary to assume the identity of the account, with all of that account's permissions on the system and network and makes it harder for defenders to detect the adversary. With sufficient access within a network, an adversary can create accounts for later use within the environment.
Discovery	Discovery consists of techniques that allow the adversary to gain knowledge about the system and internal network. When adversaries gain access to a new system, they must orient themselves to what they now have control of and what benefits operating from that system give to their current objective or overall goals during the intrusion. The operating system provides many native tools that aid in this post-compromise information-gathering phase.
Lateral Movement	Lateral movement consists of techniques that enable an adversary to access and control remote systems on a network and could, but does not necessarily, include execution of tools on remote systems. The lateral movement techniques could allow an adversary to gather information from a system without needing additional tools, such as a remote access tool.
Collection	Collection consists of techniques used to identify and gather information, such as sensitive files, from a target network prior to exfiltration. This category also covers locations on a system or network where the adversary may look for information to exfiltrate.
Exfiltration	Exfiltration refers to techniques and attributes that result or aid in the adversary removing files and information from a target network. This category also covers locations on a system or network where the adversary may look for information to exfiltrate.
Command and Control	The command and control tactic represents how adversaries communicate with systems under their control within a target network. There are many ways an adversary can establish command and control with various levels of covertness, depending on system configuration and network topology. Due to the wide degree of variation available to the adversary at the network level, only the most common factors were used to describe the differences in command and control. There are still a great many specific techniques within the documented methods, largely due to how easy it is to define new protocols and use existing, legitimate protocols and network services for communication.

Table 1 – Adversary Tactics

MITRE's ATT&CK Matrix loosely maps to the latter stages of MITRE's version of the Cyber Kill Chain®, called the Cyber Attack Lifecycle (18), also shown in Figure 10. ATT&CK was widely adopted across all organisations we engaged with and has also been formally adopted by the Cross-Government CTI Working Group to provide a consistent terminology and framework.

Aside from providing a common framework, the ATT&CK Matrix can be of great use to hunters by encouraging them to ask questions such as *“Can we currently detect the Drive-by Compromise technique within the Initial Access tactic if used against us by an adversary?”* Exercises can be used here to assess whether any given technique can be detected or not, with the blue team⁹ (i.e. hunters) requesting the red team to perform a specific technique, or alternatively the red team can perform a set of techniques without the blue team's knowledge to see what they can detect. MITRE provides detection information for each technique that can be used to aid hypothesis generation, and data source information that can be used to check whether the existing data visibility is sufficient, and if not, then provide focus and justification for subsequent data collection efforts and any associated cost.

Additionally, the ATT&CK matrix can be of use when assessing data visibility – 50 data sources are required to enable detection of all 223 currently described techniques. Table 2 lists the data sources and the number of techniques that each source contributes to detection of (in a necessary but not sufficient manner i.e. most techniques require multiple data sources). This can help prioritise the organisation's data collection efforts, for example, look to ensure *Process monitoring* data is collected across the whole estate before expending effort to collect *WMI Objects* data. However, consideration also needs to be given to the scope of collection and the individual systems in question. For example, *Process monitoring* data can be collected with relative ease from in-house systems but would be more challenging to collect from outsourced systems e.g. Software-as-a-Service (SaaS) or third-party organisations within the supply chain.

RECOMMENDATION 9: Adopt MITRE's ATT&CK Matrix for Enterprise to aid hypothesis generation and data visibility tracking.

⁹ [https://en.wikipedia.org/wiki/Blue_team_\(computer_security\)](https://en.wikipedia.org/wiki/Blue_team_(computer_security)). A blue team is a group of individuals who perform an analysis of information systems to ensure security, identify security flaws, verify the effectiveness of each security measure, and to make certain all security measures will continue to be effective after implementation.

Detecting the Unknown: A Guide to Threat Hunting

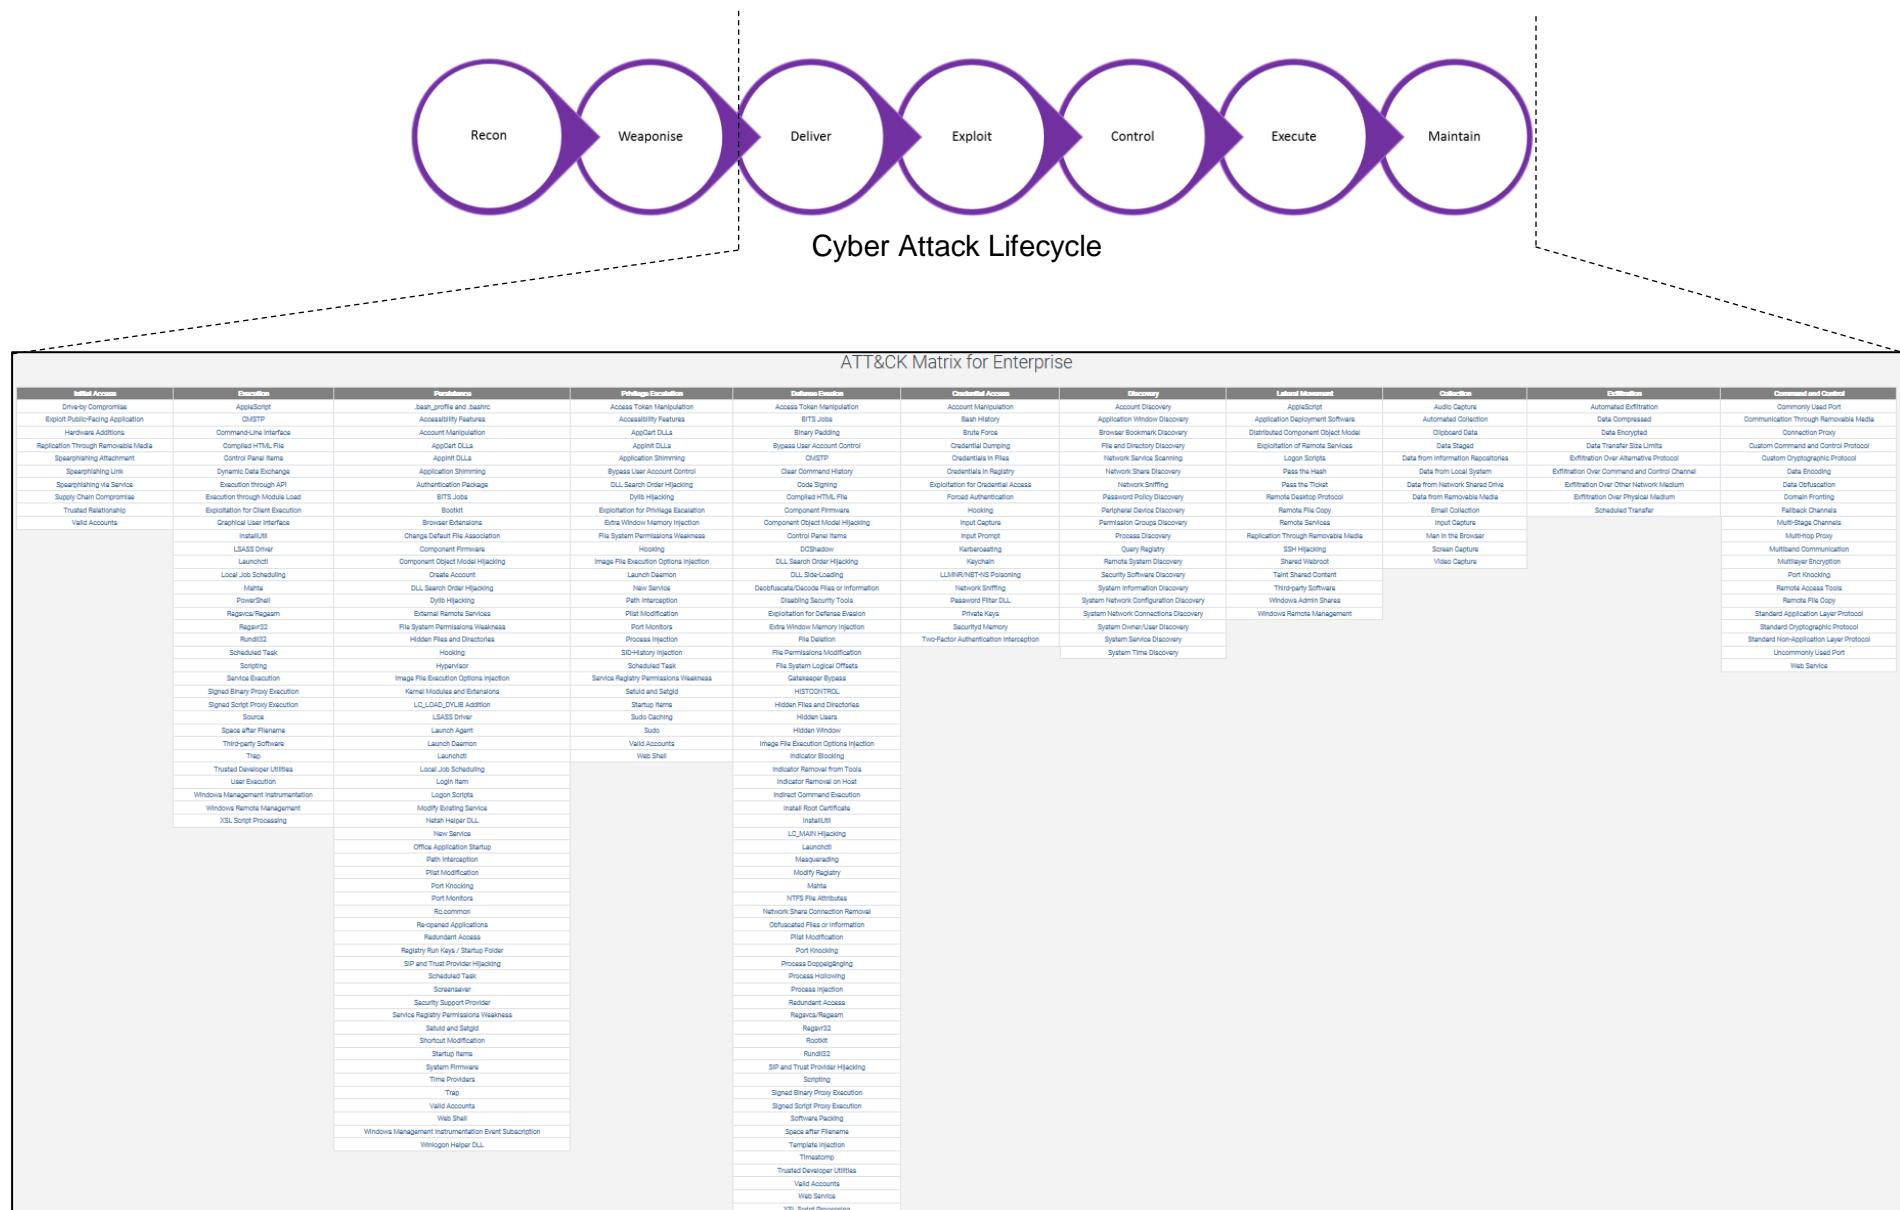


Figure 10 – MITRE’s Cyber Attack Lifecycle and ATT&CK Matrix for Enterprise*

* Accurate at time of writing – see <https://attack.mitre.org/> for the most current version

Data Source	Techniques	Data Source	Techniques
Process monitoring	155	Web proxy	4
File monitoring	89	Windows Error Reporting	4
Process command-line parameters	85	Host network interface	3
API monitoring	39	Services	3
Process use of network	36	Third-party application logs	3
Windows Registry	34	BIOS	2
Packet capture	32	Detonation chamber	2
Authentication logs	28	Environment variable	2
NetFlow/Enclave NetFlow	24	Mail server	2
Binary file metadata	18	MBR	2
DLL monitoring	17	Web logs	2
Network protocol analysis	17	Access Tokens	1
Windows event logs	15	Asset Management	1
Loaded DLLs	12	Browser extensions	1
Malware reverse engineering	9	Component firmware	1
System calls	9	Digital Certificate Logs	1
SSL/TLS inspection	8	Disk forensics	1
Anti-virus	7	DNS records	1
Data loss prevention	6	EFI	1
Network intrusion detection system	6	Named Pipes	1
Application Logs	5	PowerShell logs	1
Email gateway	4	Sensor health and status	1
Kernel drivers	4	VBR	1
Network device logs	4	Web application firewall logs	1
User interface	4	WMI Objects	1

Table 2 – Data Sources

A visual method of representing current data visibility is via the use of a heatmap, with a good instructional example detailed in a blogpost by Roberto Rodriguez (aka Cyb3rWard0g) (19). That example scores each technique based on the amount of data collected, quality of the data collected, data science techniques used, etc. To track this, MITRE’s ATT&CK Navigator¹⁰ can be utilised. The Navigator is an interactive view of the Matrix that allows each technique to be colour-coded, and the output can be exported to Excel or as JavaScript Object Notation (JSON).

While the Threat Hunting team should fully adopt the ATT&CK Matrix and embed it into all aspects of their Threat Hunting process, they (or even the SOC in general) will likely have only limited influence over what data and logs they receive, as this will often fall under the remit of the individual System Owners. Therefore, the organisation should adopt a policy that all new systems will send logs to the SOC’s central repository, and then on-board logs from existing systems, perhaps by standing up a centrally-funded on-boarding project as necessary. This can incentivise the business units to assist the data collection effort, as they would only need to fund the Business-as-Usual (BAU) costs, rather than the more expensive on-boarding costs.

RECOMMENDATION 10: Take steps at the enterprise-level, such as policy enforcement, to ensure the Threat Hunting team has the data visibility required to defend the organisation.

¹⁰ <https://mitre-attack.github.io/attack-navigator/enterprise/>

ATT&CK Heatmap Example

As a worked example of a heatmap, take for instance an organisation that only collects the top ten data sources listed in Table 2. We will use a slightly simpler scoring system than Cyb3rWard0g, shown below, which is mapped from our CMM. Each technique is scored based on the Hunting team's visibility of the required data sources.

Maturity Level	Description	Colour
1	Little or no data is collected.	Red
2	A moderate variety of data types is collected from key areas.	Yellow
3	A high variety of data types is collected from key areas.	Yellow
4	A moderate variety of data types is collected from across the estate.	Green
5	A high variety of data types is collected from across the estate.	Green

The table below shows the data sources required for each technique within the Initial Access tactic. Let's assume the organisation has visibility of the top five sources from across the estate (**bold**) but that they only have visibility of the other five sources from key areas within the estate (*italics*). Each technique would then be given a relevant detection maturity level and associated colour.

Initial Access	Initial Access
Drive-by Compromise	Drive-by Compromise
Network device logs	Detonation chamber
Network intrusion detection system	DNS records
<i>Packet capture</i>	Email gateway
Process use of network	Mail server
SSL/TLS inspection	<i>Packet capture</i>
Web proxy	SSL/TLS inspection
	Web proxy
Exploit Public-Facing Application	Spear-phishing via Service
Application logs	Anti-virus
<i>Packet capture</i>	SSL/TLS inspection
Web application firewall logs	Web proxy
Web logs	
Hardware Additions	Supply Chain Compromise
Asset Management	File monitoring
Data loss prevention	Web proxy
Spear-phishing Attachment	Trusted Relationship
Detonation chamber	Application logs
Email gateway	<i>Authentication logs</i>
File monitoring	Third-party application logs
Mail server	
Network intrusion detection system	
<i>Packet capture</i>	

4.3 Investment Priorities

As a human-centric capability, the focus for investment should be on people and processes before tooling. Recruitment and training of skilled threat hunters should be a high priority for the organisation, because while tools and technology enable Threat Hunting, they are not sufficient by themselves. However, this section assumes that departments already meet the baseline of the MCSS, discussed further in Section 5.2.1, and that basic SOC capabilities such as Protective Motioning are already operating at a competent maturity – if this is not the case, then investment should be prioritised in these areas first.

For Threat Hunting, the priority of investing in people was echoed by the organisations we engaged with, and is reflected in the SANS 2018 Threat Hunting Survey, where 29.9% of respondents prioritised staffing, and 19.8% prioritised training, meaning 49.7% prioritised investment in people (8) over services or technology.

While this paper has not assessed any specific training courses, care should be taken to ensure any training procured delivers specific Threat Hunting knowledge, as opposed to re-branded blue teaming or CSIRT courses – while these would still add value to your Threat Hunting team, they would not necessarily explore the specific processes or playbooks required. At the time of writing, the authors of this paper are aware of only a limited number of Threat Hunting-specific courses, but offerings on the market should be continuously assessed. Additionally, the benefit gained from on-the-job training, or internally developed courses, should not be ignored.

RECOMMENDATION 11: Prioritise the recruitment and training of skilled threat hunters.

The logical result of focusing investment on your people is that less investment is then available for tools. However, the organisations we engaged with were predominantly leveraging existing tools – both commercially procured and Free and Open Source Software (FOSS). Again, this is reflected by the SANS 2018 Threat Hunting Survey, with 90.3% of organisations using existing infrastructure tools, 61.9% developing tools in-house, 47.8% using FOSS hunting tools, and only 32.5% procuring commercially available Threat Hunting tools (8). This represents the maturity of the Threat Hunting tools and services on offer by security vendors within the market. While this paper has not assessed any specific tools, care should be taken to properly understand the features on offer rather than taking any sales or marketing material at face-value, ensuring the solutions are designed with a proactive stance at their core, rather than merely being reactive offerings re-branded as ‘Threat Hunting’.

As such, we would encourage organisations looking to invest in Threat Hunting-specific tooling to first look at implementing FOSS alternatives; this will assist with refining the requirements for any commercial product before going to market, and hence improve the Return on Security Investment (ROSI). Additionally, FOSS tools are commonly supported, and mitigate the risk of being ‘locked-in’ with any specific vendor.

RECOMMENDATION 12: Apply caution to Threat Hunting tooling investment, instead leveraging existing tools and FOSS.

None of the organisations we engaged with were utilising Threat Hunting-specific tooling, highlighting that the market is fairly immature at the time of writing; however, technology develops rapidly and horizon scanning (systematically investigating evidence about future trends) should be maintained to benefit from future advances.

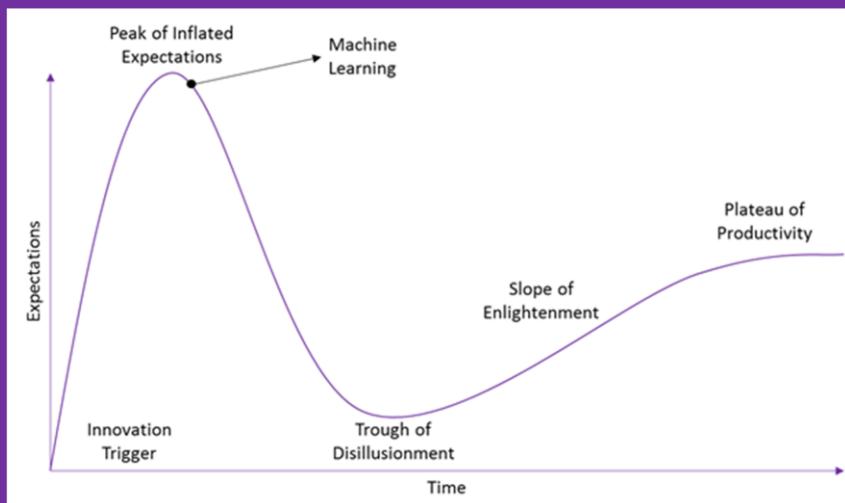
RECOMMENDATION 13: Maintain horizon scanning for future tooling e.g. machine learning solutions.

Of the organisations we engaged with, those that did wish to invest in tooling were primarily focused on solutions that improve data visibility to better enable Threat Hunting, rather than Threat Hunting-specific tooling itself. A priority we repeatedly heard was for the deployment of Endpoint Detection and Response (EDR) solutions that allow greater visibility into endpoint data. This reiterates the earlier discussion that data visibility is key for effective Threat Hunting to occur.

By assessing the current data visibility and identifying detection gaps, the SOC can then start to build a business case for investment in additional tooling e.g. EDR solutions, with metrics from Hunting being used to highlight the benefits of a proactive approach.

Horizon Scanning for Machine Learning

Taking machine learning for example – as of July 2018, Gartner included it in their Hype Cycle for Data Science and Machine Learning (33) in the ‘Peak of Inflated Expectations’ stage, represented below, predicting 2-5 years until it reaches the ‘Plateau of Productivity’. Essentially this means that machine learning has now been implemented by early adopters, with mixed success, and the prediction is that it will be 2-5 years until it is widely implemented, and its application and benefits are well understood.



https://en.wikipedia.org/wiki/Hype_cycle. The hype cycle is a branded graphical presentation developed and used by the American research, advisory and information technology firm Gartner to represent the maturity, adoption, and social application of specific technologies.

4.4 Metrics

Most of the organisations we engaged with did not have any metrics for their Threat Hunting capability, either due to a difficulty in identifying metrics of value, or through a belief that as Threat Hunting needs to be a flexible process, metrics in general were not suitable. However, there are useful metrics that can provide a measurement of performance to help drive improvements and can also evidence the ROSI to senior managers within the organisation, helping to build the business case for further investment (financial and time) in your people and tools. Below is an example set of metrics that could be adopted:

Metric Description	Metric Type
Number of incidents identified proactively (vs. reactively)	Trend, Comparison
Number of vulnerabilities identified proactively (vs. vulnerability assessments)	Trend, Comparison
Dwell time of proactively discovered incidents (vs. reactively)	Trend, Comparison
Containment time ¹¹ of proactively discovered incidents (vs. reactively)	Trend, Comparison
Effort per remediation of proactively discovered incidents (vs. reactively)	Trend, Comparison
Data coverage (data types and coverage of estate)	Percentage
Hypotheses per MITRE ATT&CK tactic	Pie Chart
Hunts per MITRE ATT&CK tactic	Pie Chart
Incidents per MITRE ATT&CK tactic	Pie Chart
Percentage of successful hunts that result in a new detection analytic or rule	Service Level
Sensitivity and specificity of analytics or rules derived from hunts (true & false positive rates)	Service Level

Table 3 – Example Threat Hunting Metrics

Ultimately, the value of any metric is how useful it is to the recipient, often a senior manager such as a CISO, so all metrics should be developed in collaboration between the Threat Hunting team and relevant senior managers.

RECOMMENDATION 14: Adopt organisationally-relevant metrics, such as our example set, to drive improvements and evidence the ROSI over time.

¹¹ Time from detection to remediation

5 Leveraging HM Government

This section is targeted at decision makers within cross-government functions, and outlines steps that can be taken across HM Government to improve collaboration, set a common baseline, and professionalise the threat hunter role and hence improve our collective security.

5.1 Collaboration

Many of the Threat Hunting teams we engaged with operated in the isolation of their own organisation. Instead, greater collaboration should be encouraged between organisations so that the community can collectively benefit from each other's experiences and lessons learnt, to better defend against malicious threats.

5.1.1 Cluster Security Units

The Government Transformation Strategy (20) has the vision to transform the relationship between citizens and the state by 2020. One aspect of achieving this is the Transforming Government Security Programme (TGSP), which will deliver a common Target Operating Model (TOM) across the four Cluster Security Units (CSU)¹². Led by the Government Security Group (GSG), this will assist in fostering collaboration by developing a common security framework and profession to support all of government in meeting the minimum level of security.

Each CSU is responsible for defining a set of security service offerings that can then be adopted and rolled out by all CSUs to their customers. Currently, the service offerings are more generally focused on physical security, with plans to explore a full cyber catalogue in the future. We recommend that GSG include Threat Hunting as a future cyber security offering for development and delivery by the CSUs, as this is an ideal existing mechanism to collaborate, define good practice and implement Threat Hunting across government.

The Government Transformation Strategies Objectives

- Continue to deliver world-class digital services and transform the way government operates
- Develop the right skills and culture among our people and leaders, and bring together policy and delivery
- Build better workplace tools and processes to make it easier for public servants to work effectively
- Make better use of data
- Create, operate, iterate and embed good use of shared platforms and reusable business capabilities to speed up transformation

RECOMMENDATION 15: When defining the full cyber catalogue, the GSG should include Threat Hunting as a security service offering from the CSUs.

¹² Cluster 1 is led by HM Revenue and Customs (HMRC); Cluster 2 is led by the Home Office (HO); Cluster 3 is led by the Department for Work and Pensions (DWP); and Cluster 4 is jointly led by the Ministry of Defence (MOD) and the Foreign and Commonwealth Office (FCO).

5.1.2 Sharing of CTI

As a minimum, any potential CTI discovered during a hunt should be processed and provided to the organisation's CTI function for dissemination, both internally and externally. The CTI can be shared in a machine-readable fashion via standardised languages and protocols such as Structured Threat Information eXpression (STIX) and Trusted Automated eXchange of Indicator Information (TAXII) respectively (21). Unstructured CTI, such as reports on specific Threat Hunting procedures, can be shared on CiSP¹³ for example, which is managed by the NCSC.

Launched in March 2013, 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."*

Aside from CTI, we would also encourage organisations to share their hunt hypotheses, procedures, playbooks and analytics with each other. While a specific hypothesis or analytic may not be of direct use from one organisation to the next due to differences in estate architecture, threat landscape, etc., it may help to stimulate discussion and produce new hypotheses or hunt ideas.

RECOMMENDATION 16: Share relevant CTI and knowledge gained from Threat Hunting across the community.

5.1.3 Cross-Government Working Groups

Many of the organisations we engaged with attend Cross-Government Working Groups for CTI and/or Security Monitoring. These groups provide an opportunity to update the community on developments at each organisation, and to share best practice. We would strongly recommend the incorporation of Hunting either into an existing group, or alternatively the establishment of a new group focused on Threat Hunting. Additionally, one organisation we engaged with hosted occasional hackathons-type events with different government organisations. Lasting for 1-2 days each, the events were used as an opportunity to share knowledge and develop hypotheses and hunt ideas together. This idea could be expanded to include cross-industry sharing, or even Threat Hunting conferences.

RECOMMENDATION 17: Set up a Cross-Government Working Group for Threat Hunting and run hackathon-type events.

¹³ <https://www.ncsc.gov.uk/cisp>

5.2 Setting the Standard

5.2.1 The Minimum Cyber Security Standard

In June 2018, the Cabinet Office published the Minimum Cyber Security Standard (4). This is the minimum set of cyber security standards that the government expects departments to adhere to and exceed wherever possible. It will be continually revisited to incrementally ‘raise the bar’ and address any new threats or vulnerabilities that arise.

While the MCSS details outcomes, rather than specific implementations, Standard 8 (DETECT), subsection a) states that *“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.”*, while subsection d) states that *“Attackers attempting to use common cyber-attack techniques should not be able to gain access to data or any control of technology services without being detected.”* The references to ‘known threats’ and ‘common cyber-attack techniques’ implies a reactive stance searching for IOCs that is better suited for Protective Monitoring, as opposed to the proactive searching for advanced unknown threats that is better suited to Threat Hunting.

Currently, only a small number of organisations are performing Threat Hunting at a competent maturity; this is reflected within the MCSS as it does not reference Threat Hunting either in name or principle. However, as Threat Hunting across departments matures, its outcomes (such as reduced dwell time through a proactive approach) should be included in a future iteration of the MCSS. Departments must however, be given sufficient sight of this to allow them to prepare.

RECOMMENDATION 18: Include the outcomes from Threat Hunting in a future iteration of the MCSS.

Additionally, the MCSS includes reference to Cyber Essentials¹⁴ (Standard 1, subsection d) as a method to demonstrate appropriate diligence by suppliers of third-party services. This is very much a minimum and does not cover the entirety of the MCSS, let alone include Threat Hunting. However, departments can define the level of assurance they require from their suppliers as part of a risk-based approach and should consider the requirement for third-parties to perform Threat Hunting of their own networks. Suppliers of critical services, where extra assurances such as ISO27001 certification are often already sought should, over time, be required to Threat Hunt on their own networks. By Threat Hunting, third-parties will better protect themselves against compromise and hence improve the security of your organisation’s supply chain; assessment of their capability should therefore be incorporated into the supplier management process.

RECOMMENDATION 19: As part of a risk-based approach to procurement and supplier management, consider requiring third-party suppliers of critical services to Threat Hunt on their own networks.

¹⁴ <https://www.cyberessentials.ncsc.gov.uk/>. Cyber Essentials helps guard against the most common cyber threats and demonstrates a commitment to cyber security.

5.2.2 The Cyber Assessment Framework v2.0

The NCSC published the Cyber Assessment Framework (CAF) v2.0 in October 2018 (22) to support the UK's implementation of the European Union's (EU) Network and Information Systems (NIS) Directive, which in turn aims to improve cyber security in Operators of Essential Services (OES). The NIS Directive defines Competent Authorities (CA) that assess OES against the NIS principles (23). In the UK, the CAs cover water, energy, digital infrastructure, the health sector, transport, and digital services providers, and include several government organisations, such as the Department for Environment, Food and Rural Affairs (DEFRA).

The CAF is broken down into four objectives, A to D, with each detailing Indicators of Good Practice (IGP). Objective C (detecting cyber security incidents) is broken down into two principles (24):

- **C1. Security monitoring:** The organisation monitors the security status of the networks and systems supporting the delivery of essential services to detect potential security problems and to track the ongoing effectiveness of protective security measures
- **C2. Proactive security event discovery:** The organisation detects, within networks and information systems, malicious activity affecting, or with the potential to affect, the delivery of essential services, even when the activity evades standard signature-based security prevent/detect solutions, or when it is not possible to use signature-based detection, for some reason

Like the MCSS, the CAF details outcomes rather than specific implementations. C1 takes a reactive stance, focusing on the monitoring coverage, security of logs, the generation of alerts, the identification of security incidents, and the monitoring tools and skills required. However, C2 takes a proactive stance and very much alludes to Threat Hunting in principle, even if not by name. C2.a (system abnormalities for attack detection) covers defining examples of the abnormalities in systems' behaviour to aid detection of malicious activity – this can effectively be achieved via adoption of MITRE's ATT&CK framework and via ingestion of relevant CTI. C2.b (proactive attack discovery) covers the understanding of sophisticated attack methods and normal system behaviour to proactively search for malicious activity – effectively performing Threat Hunting.

In the UK, it is up to each sector's CA to decide if and how the CAF is implemented, as the NCSC has no regulatory role under NIS. To broaden implementation of Threat Hunting as a capability, we recommend that the CAs implement the CAF (and specifically Objective C) where possible, and that other organisations not bound by NIS look to CAF as an example of a proactive standard.

RECOMMENDATION 20: OESs could operate a Threat Hunting capability to evidence Objective C2 of the EU's NIS Directive

5.2.3 Commercial Considerations for Threat Hunting

As Threat Hunting requires intimate knowledge of the network in question, there is a commonly held view that it is not possible to perform it as an outsourced capability e.g. by a managed SOC utilising their own people, processes and tools. Of the organisations we engaged with, only one operated an outsourced SOC, but they did not Threat Hunt due to the simple fact that the commercial arrangements did not cater for it. While an in-house Threat Hunting capability is preferable as the team would have flexibility not afforded by commercial arrangements, an outsourced capability is certainly possible if the hunters are provided the data visibility and access they require – this is no different to the requirements of an in-house SOC and Hunting team.

We would recommend that any government organisation that looks to outsource their SOC should include requirements for a Threat Hunting capability and use our CMM, or similar, to appropriately define this in the commercial documentation. Any lessons learnt from such a process should then be shared across the community, such as via CiSP.

RECOMMENDATION 21: Include requirements for Threat Hunting in future commercial arrangements for outsourced SOC functions, and share the lessons learnt.

Future engagement with the Crown Commercial Service (CCS) could be considered if a significant number of public organisations wish to procure outsourced SOC functions that provide Threat Hunting capabilities, to define the best commercial framework/approach to facilitate this.

RECOMMENDATION 22: Work with the CCS to define the best route to market for outsourced SOC functions that provide Threat Hunting capabilities.

Finally, commercial arrangements should be in place with suppliers that manage infrastructure on behalf of your organisation e.g. Software/Platform/Infrastructure as a Service providers, to ensure they provide the appropriate data sources and visibility to the SOC, allowing the Threat Hunting team to operate within these systems.

RECOMMENDATION 23: Include commercial requirements for service providers to supply the SOC with the necessary data visibility required for Threat Hunting.

5.3 Development of the Profession

Section 3.2.2. articulated the difficulty faced by organisations across the industry with recruiting and retaining skilled cyber professionals, particularly for specialist roles such as threat hunters. Work is being carried out by the government to ensure the existence of a pipeline of talented individuals.

5.3.1 UK Cyber Security Council

In July 2018, the Department for Digital, Culture, Media and Sport (DCMS) launched a consultation on the development of the cyber security profession in the UK (25). The definition of cyber security taken by DCMS is the 19 draft Knowledge Areas from the Cyber Security Body of Knowledge (CyBOK), currently being developed by UK academics led by Bristol University (26).

Delivery of these objectives would be driven by a new and independent UK Cyber Security Council. The consultation closed in August 2018, and in December 2018 DCMS issued a Request for Proposal for the design and delivery of this council (27). Applications are due in February 2019, with working aiming to commence in May 2019.

We recommend that the teams developing the CyBOK and of the UK Cyber Security Council engage with Threat Hunting teams across government to ensure that Threat Hunting as a distinct capability is appropriately represented within the Cyber Security Profession, including via the adoption of this papers recommendations.

RECOMMENDATION 24: The teams developing the CyBOK and the UK Cyber Security Council should consider recognising Threat Hunting as a distinct domain.

5.3.2 The Government Security Profession Unit

The Government Security Profession Unit (GSPU) aims to bring together security professionals working in government to help them gain the skills and knowledge they need to carry out their roles (28). There are five Profession Frameworks currently being defined – these are: Cyber, Technical, Personnel, Physical, and Business Continuity Management (BCM). These Profession Frameworks will establish job families, and then job roles and the corresponding skills frameworks within each. The job families for Cyber are:

- Operational Security
- Security Architecture
- Risk Management
- Governance, Strategy and Policy

Objectives for the Profession

The DCMS consultation's proposed objectives for the profession to deliver by 2021 are:

- Professional Development (e.g. Royal Chartered status)
- Professional Ethics
- Thought Leadership and Influence
- Outreach and Diversity
- Developing the Next Generation

RECOMMENDATION 25: The GSPU should define threat hunter as a distinct role within the Operational Security job family.

6 Conclusion

In an ever increasingly digital and connected world, the cyber threat facing most organisations is growing. While the threat profile is different for each organisation, UK government departments will undoubtedly have their defences tested by advanced and persistent threat actors, which may not be detected or prevented by technical controls and reactive monitoring.

To detect these unknown and advanced threats, departments should now start moving towards a proactive stance by operating a Threat Hunting capability, and hence improve their security posture and reduce their cyber risk. This capability enables malicious activity to be identified earlier on in an attack, thereby minimising the opportunity for adversaries to disrupt, damage or steal.

To research this guide, we conducted an extensive literature review and held engagements with nine government bodies, including the NCSC, and three industry partners. This allowed us to understand the current capability across HM Government, and define a target capability through our Capability Maturity Model.

This guide provides recommendations for SOCs, government departments, and across HM Government, to detect unknown malicious activity through development of Threat Hunting as both a capability and a profession.

Before operating a proactive detection capability, government departments must create an enabling environment by ensuring they meet the requirements of the MCSS, and by providing the necessary enablers for their Threat Hunting function such as actionable Cyber Threat Intelligence, relevant data from across the estate, and appropriate investment in people, processes and tools.

6.1 Operate a SOC-based Threat Hunting Capability to Reduce Risk

The following recommendations will aid SOCs in building the capability required from people, processes and tools to operate a basic but competent Threat Hunting capability and hence reduce risk:

- To aid assessment of organisational performance and identify areas for improvement, adopt a standard framework such as our Threat Hunting Capability Maturity Model
- Recruit or train a Threat Hunting lead as a responsible individual to enable development of the capability
- Rotate SOC analysts into the Threat Hunting team for learning and development purposes
- Periodically assess the relevance of automated analytics to the organisation
- Adopt a formalised process, such as our Extended Hunting Loop, to aid operationalisation of mature Threat Hunting processes
- Utilise a workflow management tool to prioritise and track workload
- Utilise a central repository to share knowledge and lessons learnt

6.2 Enable the Threat Hunting Function to Improve the ROSI

The following recommendations will aid government departments to enable their Threat Hunting function and hence improve its efficiency and effectiveness:

- Adopt MITRE’s ATT&CK™ Matrix for Enterprise to aid hypothesis generation and data visibility tracking
- Take steps at the enterprise-level, such as policy enforcement, to ensure the Threat Hunting team has the data visibility required to defend the organisation
- Prioritise the recruitment and training of skilled threat hunters
- Apply caution to Threat Hunting tooling investment, instead leveraging existing tools and free and open source software
- Maintain horizon scanning for future tooling e.g. machine learning solutions
- Adopt organisationally-relevant metrics, such as our example set, to drive improvements and evidence the return on security investment over time

6.3 Leverage HM Government to Develop the Threat Hunter Role

The following recommendations will aid cross-government functions to improve collaboration, set a common baseline, and professionalise the threat hunter role and hence improve our collective security:

- Provide threat hunters with the actionable Cyber Threat Intelligence they require to generate relevant and testable hypotheses
- When defining the full cyber catalogue, the Government Security Group should include Threat Hunting as a security service offering from the Cluster Security Units
- Share relevant Cyber Threat Intelligence and knowledge gained from Threat Hunting across the community
- Set up a Cross-Government Working Group for Threat Hunting and run hackathon-type events
- Include the outcomes from Threat Hunting in a future iteration of the Minimum Cyber Security Standard
- As part of a risk-based approach to procurement and supplier management, consider requiring third-party suppliers of critical services to Threat Hunt on their own networks
- Operators of Essential Services could operate a Threat Hunting capability to evidence Objective C2 of the European Union’s Network and Information Systems Directive
- Include requirements for Threat Hunting in future commercial arrangements for outsourced SOC functions, and share the lessons learnt
- Work with the Crown Commercial Service to define the best route to market for outsourced SOC functions that provide Threat Hunting capabilities

- Include commercial requirements for service providers to supply the SOC with the necessary data visibility required for Threat Hunting
- The teams developing the Cyber Security Body of Knowledge and the UK Cyber Security Council should consider recognising Threat Hunting as a distinct domain
- The Government Security Professional Unit could define threat hunter as a distinct role within the Operational Security job family

Coordinated investment in Threat Hunting across SOCs, departments, and HM Government can lead to improvements in our collective security, while helping to develop the next generation of the UK's defenders.

7 Appendices

7.1 Appendix I – Contributors

The authors would like to thank the following organisations, and the individuals we engaged with, for their input into this report between October 2018 and January 2019:

- BAE Systems Applied Intelligence
- Bank of England
- BT
- Cabinet Office
- Cluster 2 Security Unit
- Department for Work and Pensions
- Foreign and Commonwealth Office
- Government Digital Service
- HM Revenue and Customs
- Home Office
- National Cyber Security Centre
- Transport for London

7.2 Appendix II – 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 (NCSP).

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
- Detecting the Unknown – A Guide to Threat Hunting
- Controlling Your Exposure – A Guide to Digital Risk and Intelligence

7.2.1 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 guide 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.

7.2.2 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 literature review and engagements with public and private sector organisations, provides recommendations for SOCs, government departments, and across HM Government, to detect unknown malicious activity through development of Threat Hunting as both a capability and a profession.

7.2.3 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 provides recommendations as to how and why government departments and HM Government as a whole, can better understand and control their digital footprint through developing a Digital Risk and Intelligence capability. Recommendations are provided at three levels; threat intelligence team level, government department level, and cross-government function level. These recommendations are also provided in the context of short, medium and long-term goals.

7.2.4 Full Capability Adoption

We recognise that each of these publications recommends dedicated resources 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 MITRE's Cyber Attack Lifecycle:

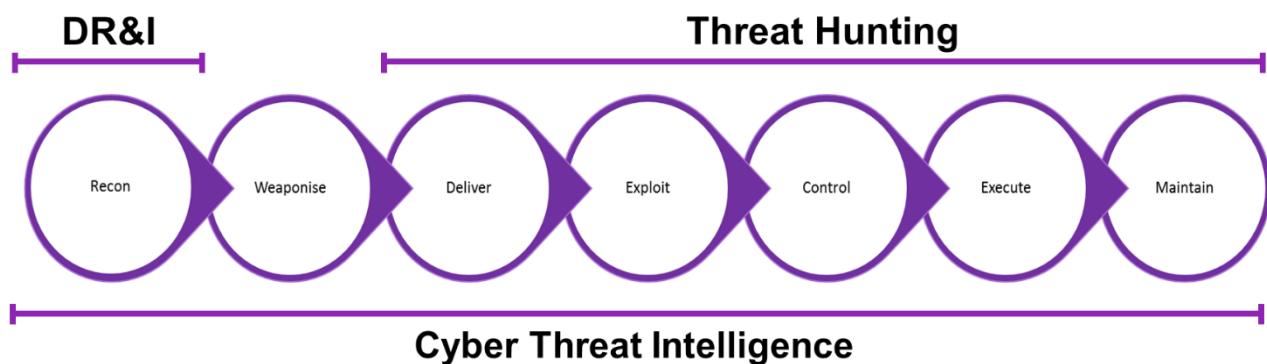


Figure 11 – Capability Scope Comparison

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 based on its cost versus business benefit.

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

- All three capabilities are subservient to each of the outcomes described 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. Each department should prioritise their investment in these capabilities based upon their own requirements and organisational context
- 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 as 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 guides specifically.

7.3 Appendix III – Bibliography

1. **HM Government.** National Cyber Security Strategy 2016 to 2021. GOV.UK. [Online] 1 November 2016. <https://www.gov.uk/government/publications/national-cyber-security-strategy-2016-to-2021>.
2. **National Cyber Security Centre.** The National Cyber Security Centre. NCSC.GOV.UK. [Online] <https://www.ncsc.gov.uk/>.
3. **HM Government.** National Security Strategy and Strategic Defence and Security Review 2015. GOV.UK. [Online] 23 November 2015. <https://www.gov.uk/government/publications/national-security-strategy-and-strategic-defence-and-security-review-2015>.
4. **Cabinet Office.** Minimum Cyber Security Standard. Gov.UK. [Online] June 2018. <https://www.gov.uk/government/publications/the-minimum-cyber-security-standard>.
5. **National Cyber Security Centre.** Security operations centre (SOC) buyers guide. NCSC.GOV.UK. [Online] 24 September 2016. <https://www.ncsc.gov.uk/guidance/security-operations-centre-soc-buyers-guide>.
6. *The Sliding Scale of Cyber Security.* Lee, Robert M. s.l. : SANS Institute, 2015.
7. *The Hunter Strikes Back: The SANS 2017 Threat Hunting Survey.* Lee, Rob and Lee, Robert M. s.l. : SANS Institute, 2017.
8. *SANS 2018 Threat Hunting Survey Results.* Lee, Robert M. and Lee, Rob T. s.l. : SANS Institute, 2018.
9. **Lockheed Martin.** The Cyber Kill Chain. Lockheed Martin. [Online] <https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html>.
10. **Joint Committee on the National Security Strategy.** *Cyber Security Skills and the UK's Critical National Infrastructure.* s.l. : House of Lords and House of Commons, 2018.
11. **Cyentia Institute.** *Voice of the Analyst Study.* 2018.
12. **Sqrrl Data.** *A Framework for Cyber Threat Hunting.* 2016.
13. **Lee, Robert. M and Bianco, David.** *Generating Hypotheses for Successful Threat Hunting.* s.l. : SANS Institute, 2016.
14. **The MITRE Corporation.** Crown Jewels Analysis. MITRE. [Online] <https://www.mitre.org/publications/systems-engineering-guide/enterprise-engineering/systems-engineering-for-mission-assurance/crown-jewels-analysis>.
15. *The Diamond Model of Intrusion Analysis.* Caltagirone, Sergio, Pendergast, Andrew and Betz, Christopher. 2013.
16. **World Wide Web Consortium.** Linked Data. W3C. [Online] <http://www.w3.org/standards/semanticweb/data>.
17. **Strom, Blake E., et al.** *MITRE ATT&CK: Design and Philosophy.* s.l. : The MITRE Corporation, 2018.
18. **The MITRE Corporation.** ATT&CK for Enterprise. attack.MITRE.org. [Online] <https://attack.mitre.org/resources/enterprise-introduction/>.

19. **Rodriguez, Roberto.** How Hot Is Your Hunt Team? *Cyber Wardog Lab.* [Online] July 2017. <https://cyberwardog.blogspot.com/2017/07/how-hot-is-your-hunt-team.html>.
20. **Cabinet Office.** *Government Transformation Strategy 2017 to 2020.* s.l. : Gov.UK, 2017.
21. **OASIS Open.** CTI Documentation. [Online] <https://oasis-open.github.io/cti-documentation/>.
22. **T, Kevin.** Introducing the Cyber Assessment Framework v2.0. *National Cyber Security Centre.* [Online] October 2018. <https://www.ncsc.gov.uk/blog-post/introducing-cyber-assessment-framework-v20>.
23. **Department for Digital, Culture, Media & Support.** *Security of Network and Information Systems. Guidance for Competent Authorities.* 2018.
24. **National Cyber Security Centre.** CAF - Objective C. *National Cyber Security Centre.* [Online] April 2018. <https://www.ncsc.gov.uk/guidance/caf-objective-c>.
25. **Department for Digital, Culture, Media & Sport.** *Implementing the National Cyber Security Strategy - Developing the Cyber Security Profession in the UK.* 2018.
26. **University of Bristol.** The Cyber Security Body of Knowledge. *CyBOK.org.* [Online] <https://www.cybok.org/>.
27. **Department for Digital, Culture, Media & Sport.** *Request for Proposals - A New UK Cyber Security Council. Annex A - Application Process and Guidance for Applicants.* 2018.
28. **Government Security Profession.** Gov.UK. [Online] <https://www.gov.uk/government/organisations/government-security-profession>.
29. **Krensky, Peter and Hare, Jim.** *Hype Cycle for Data Science and Machine Learning,* 2018. s.l. : Gartner, 2018.
30. *Threat Hunting: Open Season on the Adversary.* **Cole, Eric.** s.l. : SANS Institute, 2016.
31. **Bianco, David.** The Pyramid of Pain. *Endpoint Detection & Response.* [Online] March 2013. <http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>.
32. **FireEye.** *M-Trends 2018.* 2018.
33. *Hunt Evil: Your Practical Guide to Threat Hunting.* **Sqrrl.**
34. **FireEye.** An In-Depth Look Into Data Stacking. *FireEye.* [Online] 2012. <https://www.fireeye.com/blog/threat-research/2012/11/indepth-data-stacking.html>.

[BLANK PAGE]



© Crown copyright 2019

This publication is licensed under the terms of the Open Government Licence v3.0 except where otherwise stated. To view this licence, visit nationalarchives.gov.uk/doc/open-government-licence/version/3 or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk.

Where we have identified any third-party copyright information you will need to obtain permission from the copyright holders concerned.