SWIFT INSTITUTE

SWIFT INSTITUTE WORKING PAPER NO. 2016-002

THE CYBER SECURITY ECOSYSTEM: DEFINING A TAXONOMY OF EXISTING, EMERGING AND FUTURE CYBER THREATS

JASON FERDINAND

RICHARD BENHAM

PUBLICATION DATE: OCTOBER 2, 2017

The Cyber Security Ecosystem: Defining a Taxonomy of Existing, Emerging and Future Cyber Threats

Dr Jason Ferdinand with Professor Richard Benham

Abstract

This report presents a new Cyber Threat Taxonomy, Cyber Attack Taxonomy and Knowledge-based Cyber Resilience Framework. The research shows that despite high-profile cases in the media and the push from government agencies and industry groups, the importance of understanding cyber threats and the associated risks are still not widely known (or at least not widely communicated within organisations). Adopting these models across industries would enhance our understanding of cyber security and enable organisations to improve communication, coordination, governance and recovery when managing cyber security. These taxonomies are built upon the result of a systematic literature review and empirical research, which provides an overview of the cyber security ecosystem. They incorporate the most up-to-date understanding of 'cyber harm' as an attempt to facilitate a more incisive understanding of Value at Risk in Cyberspace (VaRiC).

Introduction

Our world is becoming increasingly interconnected, with vital services and supporting infrastructures dependent upon digitised information. Any threat to the confidentiality, integrity, and availability of digitised information is thus of critical importance to nation states, organisations, and end users. Despite decades of research on the cyber threat landscape we have yet to produce a comprehensive overview of the existing, emerging, and future threats which organisations can use to ascertain (and improve) the robustness of their cyber defences and the level of cyber resilience that they have achieved. The range of threat actors and levels of cyber attack sophistication have increased over the years, prompting an ever-tighter focus of research work on specific aspects of cyber security from within academic subject areas. Dominated by computer science and law, the literature relating to the cyber threat landscape remains opaque to the non-specialist, with a proliferation of terminology that obfuscates points of commonality and similarity. It could also lead to misunderstanding, confusion, wasted and/or redundant effort, and great financial cost. Advances have undoubtedly been made in these areas, but we must also recognise that cyber threats are always evolving.

Although the primary audience for this article is banking and finance organisations, we sought to provide analysis and guidance for all organisations based on the logic that all organisations are connected to banking and finance as corporate customers. Hence the banking and finance industry have a vested interest in assisting organisations to improve their knowledge and actions in relation to cyber crime and cyber security. The starting point for our research was thus the commitment to attempt to establish a common language for cyber security to help all organisations deal with the cyber threats in their environment, and to enable meaningful discussion of these threats within and between organisations. This sought to build upon the pioneering work of Howard and Longstaff (1998), and their early attempts at establishing a common language for computer security. Howard and Longstaff developed a minimum set of 'high-level' terms, along with a structure indicating their relationship (a taxonomy), which they used to classify and understand computer security incident information (1998: iii).

To adequately understand and assess the cyber threat landscape, and to benefit from the insights from the multitude of academic, government and practitioner sources of information, we argue that it is necessary to map the cyber security ecosystem utilising the

most up-to-date taxonomies, theories, and understandings (Oborne, 2004; Pfleeger et al 2006; Das, 2015), where the cyber security ecosystem is understood to describe "the activities of creating, preventing, dealing with, and mitigating insecurity in the use of information technology" (Oborne, 2004: 12). This orientation enables us to understand the various factors and relationships in play in the environment, as well as the ability to locate recent developments in research within this framework. Within this cyber security ecosystem we can insert terminology and taxonomies to aid understanding and promote a common language.

This report therefore presents an improved understanding and taxonomy of cyber threat. A cyber security ecosystem model will contextualise the taxonomy of cyber threat and its associated relationships with cyber security maturity and cyber resilience (Ferdinand, 2015). Furthermore, to facilitate the development of dynamic cyber resilience capability, and effective policy and strategy decision making, an updated taxonomy of cyber attacks within the cyber security ecosystem is presented. This incorporates the most up-to-date understanding of cyber harm (Agrafiotis et. al., 2016) as an attempt to facilitate a more incisive understanding of Value at Risk in Cyberspace (VaRiC) that has recently been advocated by the World Economic Forum (WEF, 2015).

The decision to incorporate the idea of cyber harm into the ecosystem is undertaken to balance the focus on cyber threat. Threat-based approaches have been criticised for often being too linear, cause-and-effect based, which leads to reactive rather than proactive defence. Cause-and-effect analysis can also neglect the cascading effects of cyber attacks (Wang and Ronga, 2009), whereas an appreciation of cyber harm offsets and balances these foci of attention by considering broader non-linear effects of cyber attacks (Agrafiotis et. al., 2016).

Section one of this report commences with a brief overview of the cyber security ecosystem to orientate the reader to the complexity of the task. This is built upon a systematic literature review that explores our past understandings and experiences of cyber threat landscape, through considering the taxonomies, models, and ideal types present. Incorporating insights from cyber incidents, risk management, cyber harm, and financial management, this section captures and summarises our past in cyber security.

Section two explores how contemporary managers tasked with cyber security management understand and use these taxonomies, models, and ideal types in practice. Through dedicated focus groups we explored these current understandings and behaviours, with some surprising results. Despite efforts made within the industry, and by government and leading figures, managers in the UK appear to be still lacking knowledge and understanding of cyber threats. It is important to note here that different countries and industries therein, have significantly different cyber security postures and associated legal requirements to conform to. This study is UK centric, therefore further international comparison studies would enhance our understanding of how representative and potentially widespread the issues identified in our research are. This international comparison was beyond the scope of the current research parameters, but this report does provide a basis for subsequent comparison studies to be made. In the UK the apparent lack of knowledge and understanding of cyber threats could be the result of a lack of cyber security activities within the organisations themselves, and/or a failure to communicate the activities in cyber security that have been undertaken. Either way the results are concerning.

Section three combines the insights gained from the previous two sections to present a new cyber threat taxonomy, a new cyber attack taxonomy, and some practical recommendations as to how organisations can build and maintain cyber resilience within the dynamic cyber threat landscape by using the taxonomies presented. The report ends with section four, summarising our conclusions and recommendations for future research.

Section One: The Cyber Security Ecosystem

A great number of articles, books, and reports have been produced in relation to cyber threat and cyber attack, but most of these publications focus on individual organisations in a decontextualized manner. They do not consider the interactions and relationships between the relevant actors in the broader environment, and therefore can produce a limited account of the nature of cyber threat. In order to contextualise our understanding of cyber threats we start our considerations by briefly setting out the cyber security ecosystem (see Figure 1).

Here the organisation is embedded in a nexus of relationships with other actors in the ecosystem, including individuals and groups of cyber criminals. We will return to this ecosystem perspective in section three of the report to demonstrate how our research contributes to increased levels of security in the cyber security ecosystem, but initially it is important to note that this perspective encourages us to consider cyber threats as both direct and indirect, where the cyber threat posed is to more than digital assets and organisational reputation. Given the increased levels of attacks reported annually, especially to third parties holding information both in supply chain organisations and professional service providers, such a starting point is invaluable to understand the nature of cyber threat and the potential resources available to reduce the cyber threat to organisations.

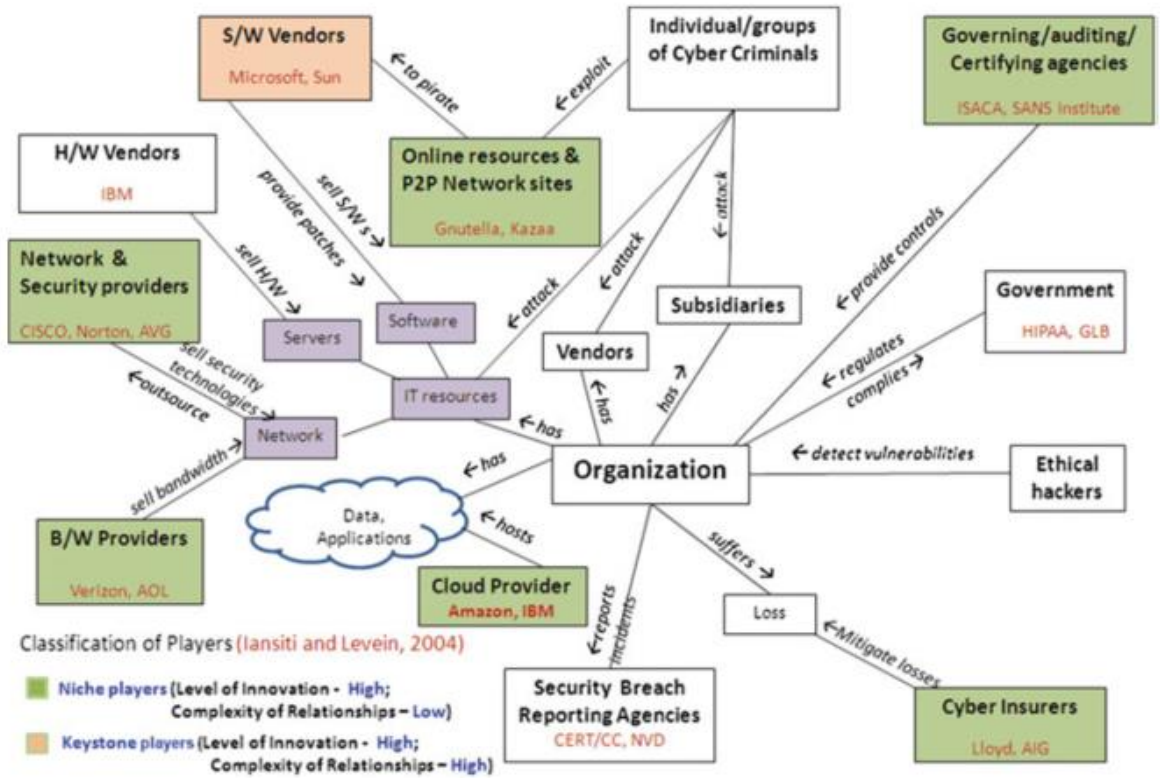

Figure 1: The cyber security ecosystem (Das, 2015: 455)

In the first stage of our research we engaged in a systematic literature review to gather all relevant information on cyber threat and response. This was done through searching four academic databases to discover the academic work already done in these areas, as well as a broader literature search of government documents, trade publications, and practitioner facing publications.

The final review of articles undertaken as a result of the systematic review focused on three areas: cyber incidents and cyber attack factors, detection and defence, and latest arguments regarding cost and harm caused by cyber attacks. The articles reviewed were selected as they capture the most salient points and examples from the 37 articles that formed our review sample to create an updated cyber threat taxonomy. For more detail on our literature review, refer to Appendix A.

Cyber incidents and cyber attack factors

Howard and Longstaff (1998) present an attempt to create a common language for computer security, and a taxonomy of computer and network incident. Here an attack is focused on a specific target within an organisation, with the intention to change that target. Attacks are described through a linear progression of five steps to achieve an unauthorised result, which allows the reader to 'plot' different pathways through this five-step process according to the motives of the attacker. Although this taxonomy is useful it does not provide detailed information regarding the motivation of the attacker, but does suggest seven types of attacker (Hackers, Spies, Terrorists, Corporate Raiders, Professional Criminals, Vandals, and Voyeurs). Thus, the taxonomy provides us with seven headline categories: Attackers, Tool, Vulnerability, Action, Target, Unauthorised Result, and Objectives. Although this taxonomy of computer and network incidents is nearly 20 years old, it is still a useful benchmark from which to map subsequent developments and improvements (see figure 2). Finally, Howard and Longstaff note that although individual attacks do occur, many attacks are sustained as 'incidents' (1998:15), and more recent work has noted that these incidents could contain multiple simultaneous attacks (Simmons et al. 2014).

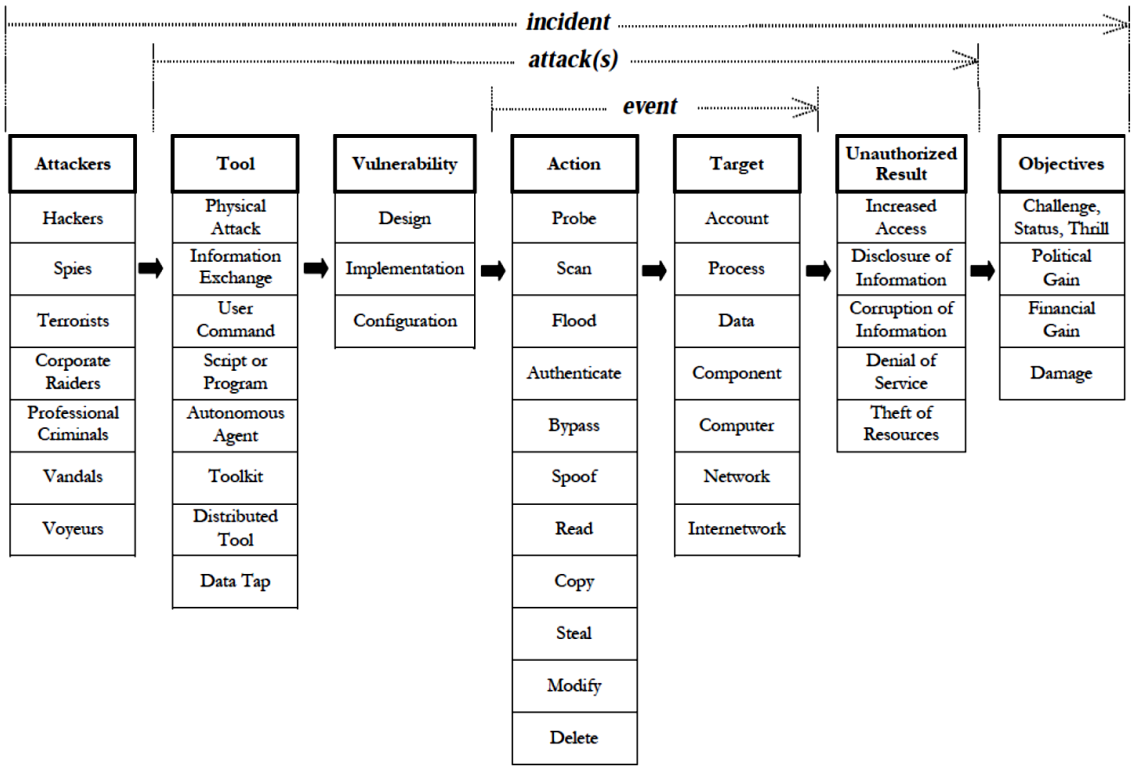

Figure 2: Computer and Network Incident Taxonomy (Howard and Longstaff, 1998: 15)

Kjaerland (2005) offers a taxonomy of 'cyber-intrusions' where attacks were examined according to their Method of Operation, Target, Source, and Impact. Each aspect contains an exhaustive description to compare government and commercial incidents. Here Kjaerland (2005) contributes most to our understanding of the motive and origin of the attacker, which has been further developed by Meyers, et al. (2009) and Hald and Pedersen (2012).

Hansman and Hunt (2005) present a taxonomy of cyber attack containing four unique dimensions, in an attempt to revive Howard and Longstaff's (1998) goal of providing consistency in language used. Hansman and Hunt's first dimension is the 'Attack Vector' which is used to classify the attack. Next is 'Target' of the attack and 'Vulnerability' as classification categories. The final dimension is 'Payload' which is used to describe and discern the effects of the attack. Each dimension contains various levels of information to supply cyber attack details.

Meyers, et al. (2009) proposed a cyber attack taxonomy similar to that of Hansman and Hunt that presents nine classes of cyber attacks. Here Meyers et al., focus on multiple attacks in a continuous development of attack taxonomies. Arguably their most valuable insight is to organise cyber adversaries into eight classes according to skill level, an approach that has recently been updated by Hald and Pederson (2012: 83) (see table 1).

New Categories	**Old Categories**
Script Kiddies	Novice
Cyber-Punks	Cyber-Punks, Virus Writers
Insiders	Internals
Petty Thieves	Petty Thieves
Grey Hats	Old Guard Hackers
Professional Criminals	Professional Criminals, Information Warriors
Hacktivists	Political Activists
Nation States	N/A, Information Warriors

Table 1: New categories of cyber hostiles

As a result of this stage of the reviewing process we can now clearly identify the cyber attack incident, a revised taxonomy of cyber hostiles, and have gained insights into levels of skill and motivations of these cyber hostile. We now turn to the sample of articles that address the detection and defence of cyber attacks, commencing with one of the most common forms of attack, the Distributed Denial of Service attack (DDoS).

Detection and defence

Mirkovic and Reihner's (2004) work focuses on Distributed Denial of Service (DDoS) attacks and related defensive mechanisms. Their comprehensive taxonomy features attack strategies and defensive countermeasures. The taxonomy of DDoS attacks is categorized by Degree of Automation, Exploited Weakness, Source Address Validity, Attack Rate Dynamics, Possibility of Characterization, Persistent Agent Set, Victim Type, and Impact on Victim. This is connected to a taxonomy of defences categorised as Activity Level, Cooperation Degree, and Deployment Location. Although there are numerous other taxonomies of attack and defence Mirkovic and Reihner's (2004) work is discussed here as an exemplar of such articles (*cf* Geng et al., 2002; Wood and Stankovic, 2004; Peng et al., 2007; Abliz, 2011; Zargar et al., 2013), in particular because they posed the question as to 'how would two different defence models perform under a given attack?'. This connects Mirkovic and

Reihner's (2004) work to later taxonomies, especially that by Abliz (2011) and Shameli-Sendi et al. (2015) on defence mechanisms, and the work of Simmons et al. (2013; 2014) discussed below.

King et al. (2009) proposed a taxonomy to capture attack pre-conditions against log anonymization. This taxonomy is based on suggesting the information required for an adversary to mount an attack. This involved the following pre-conditions and specific log properties: Structure Recognition, Fingerprinting, Known Mapping, Cryptography and Data Injection. King et al.'s taxonomy presents a novel attack classification utilising a graph to represent attack pre-conditions using nodes.

Amer and Hamilton (2010) present an Intrusion Detection System (IDS) taxonomy built upon source of audit, layout technology, structure and arrangement, data collection and processing, detection paradigm and time of detection. This led Amer and Hamilton to argue that organisations employing one IDS may have ineffective cyber security. They suggest that using their taxonomy may enable organisations to tailor their IDS according to the defined characteristics.

Wu et al. (2011) provide an attack classification for automatic response systems built upon three dimensions (Source: attack origin; technique: method used by attacker; result: outcome of the attack). Here a matrix for every attack technique is defined according to source and result to define automatic defence. This is a novel approach that is somewhat limited in that it does not specify the target of attack, nor do the authors address blended and complex attacks which may be a lot more difficult to classify and therefore automatically counter.

Li et al. (2012) present a taxonomy of basic cyber attacks in a smart grid based on four levels of attacks (device, data, privacy, network availability). This approach focuses on the target of the smart grid attacked, rather than on the incident, attacker characteristics or tools used. It is quite a basic approach, but Li et al. do match these attacks to defence activities such as Access Control, Authentication, Privacy Preservation, and Intrusion Detection.

Simmons et al (2013) present an attack-defence and performance taxonomy (ADAPT) based on a game theoretic approach to cyber security. Three classifiers: attacker, defender and performance, are identified which are actually metrics introduced in an attempt to assign values to the cost and benefit of attack and defence. Here it is important to note that the attack and defence metrics are related to the game model being employed, whereas the performance metric is intended to provide insight into how well a system is performing. This should facilitate a more informed discussion of the cost/benefit of cyber security investments (see figure 3).

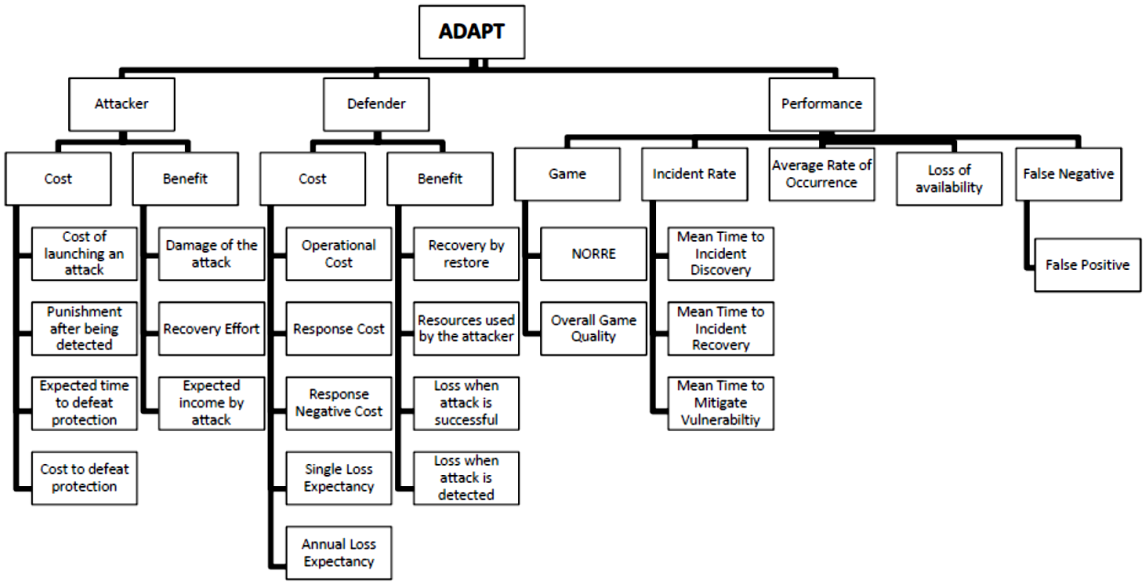

Figure 3: Attack-Defence and Performance Taxonomy (Simmons et. al. 2013: 4)

This work is useful as it brings us more up-to-date in terms of considering the costs of cyber security for the organisation concerned. Simmons continues this work with colleagues to present a cyber attack taxonomy called AVOIDIT used to identify and characterise attacks (see figure 4). This taxonomy focuses on the notion of 'blended attacks' where a hostile actor exploits one or more vulnerabilities to perform its attack. Simmons et al (2014) use the example of the Windows Server service Remote Procedure Call (RPC) stack buffer overflow vulnerability which was utilized by both 'Gimmiv.A' and 'Conficker' attacks (Porras et. al., 2009) as a result a variation of an attack used against a pre-existing vulnerability. Simmons et. al. (2014) argues that variants of an attack are not captured in previous taxonomies, and nor are the impacts of a successful cyber attack on the organisation once it has been compromised.

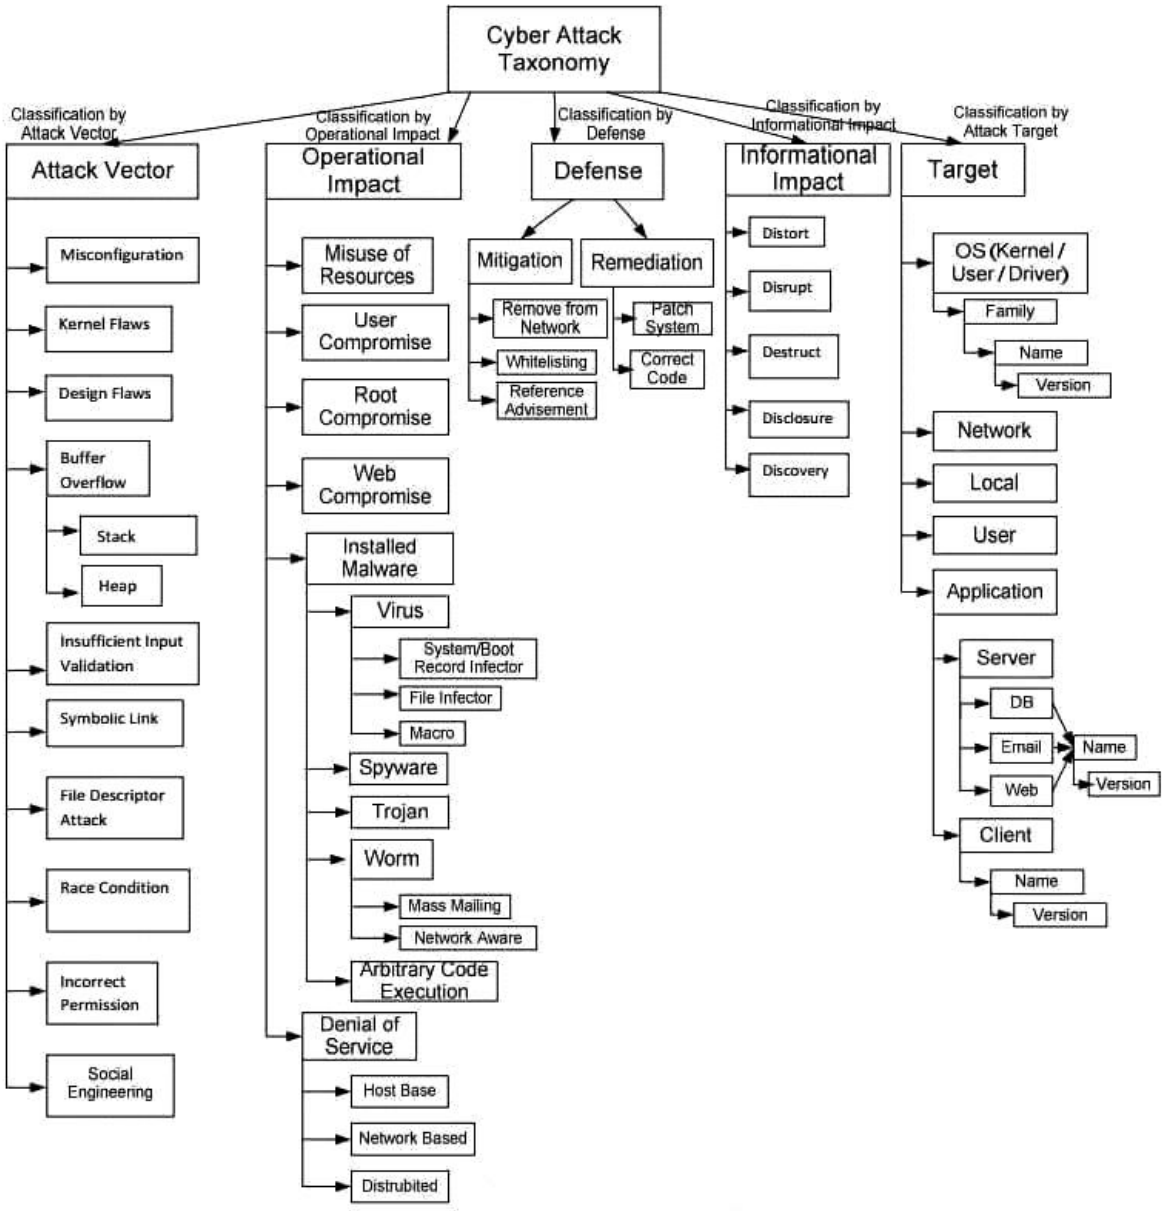

Figure 4: The AVOIDIT cyber attack taxonomy

Although Simmons et al.'s (2014) work is an improvement on previous attack taxonomies, it still demonstrates some limitations in terms of a lack of defence strategies and the incorporation of physical and virtual attack surfaces. The lack of defence strategies could be due to the range of defensive measures that are now available, whereas the lack of discussion of the connection between physical and virtual aspects of cyber attack and defence perhaps reflects the preoccupation with technology popular at the time of publication.

What AVOIDIT also offers is a more detailed breakdown of cyber attacks and suggests areas to be addressed in terms of monitoring, detection and mitigation. It can therefore be seen as a contributory update on our cumulative knowledge of cyber attacks and defence. But if we were to incorporate the physical aspects of cyber attacks, highlighted above as an omission of the AVOIDIT taxonomy, we could then address the securing of physical and virtual attack surfaces leading to prevention of cyber attack. This is an aspect addressed by Shameli-Sendi et al. (2015) through a generic defence life-cycle of four phases: prevention, monitoring, detection, and mitigation (see figure 5). Starting with the prevention phase, organisations can put in place appropriate controls at different locations to secure services and data against attack. These can be physical controls, such as security card access, and digital controls such as firewalls.

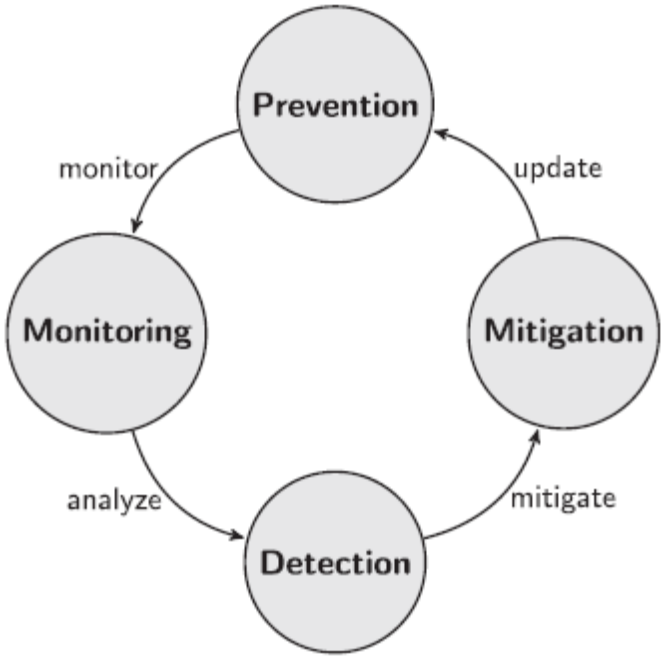

Figure 5: Cyber defence life-cycle (adapted from Shameli-Sendi et. al. 2015)

What this generic life-cycle of cyber defence encourages is a vital focus on prevention, based on an organisation's overall knowledge and understanding of potential cyber threats and actual cyber incidents. It also encourages a dynamic learning process to be established within organisations to facilitate improved cyber defences. For example, in the monitoring phase tools are deployed to gather information on attacks and the operational execution. The detection phase demonstrates how properly the IDS and live stream analysis of systems is configured, and what attacks are not being picked up. The mitigation stage selects the

correct response to attack, evaluates the seriousness of attacks, and helps managers to identify effective countermeasures to attacks to effectively handle or slow down malicious access (Walfish et al., 2010).

This ongoing process thus calls for campaign analysis of cyber threats and attacks suffered (see Hutchins et. al., 2011). It is important to note here that campaign analysis is normally associated with multi-year, Advanced Persistent Threats (APTs) that seeks to utilise vulnerabilities to perform blended and multiple/multistage attacks (Vries et al. 2012). In an ideal situation all organisations should either engage with campaign analysis directly, and/or become members of an industry association undertaking such analysis. What is needed to justify the decision to do this is the business case for cyber security, and so in the final section of this review we turn to the most recent efforts to capture and understand the costs and harm associated with cyber threats.

The costs and harm associated with cyber threats

The issue of calculating, or at least estimating with a fair degree of accuracy, the cost of cyber crime is a perennial problem in the cyber security industry. For example, leading insurers Lloyds of London estimated that in 2015 cyber crimes cost businesses as much as $400 billion a year[1], whereas in the same year research by Juniper estimated that cyber crime will cost business over $2 trillion by 2019[2]. This represents the impact of cyber crime already, and the increasing cyber crime costs over the coming years. Depending on where you look estimates of cyber crime, both now and in the future, range dramatically, and it is against this backdrop of confusion that the World Economic Forum (WEF) promoted a focus on Value at Risk (WEF, 2015).

Value at Risk (VaR) is a well-known statistical approach to risk modelling in the banking and finance industry that seeks to quantify the level of financial risk over a specific timeframe. Originally introduced in 1994 by JP Morgan, VaR has been widely accepted and is now considered as an industry standard tool to assess market risk. Investment and commercial banks use it to estimate the extent and occurrence ratio of potential loss. VaR modelling is based on the decision as to the potential for loss and the probability of the loss occurring over a set time period. Despite being widely accepted it is worth noting that there is no standard protocol for the statistical determination of a portfolio of assets, or firm-wide risks,

[1] *cf* http://fortune.com/2015/01/23/cyber attack-insurance-lloyds/
[2] *cf* https://www.juniperresearch.com/press/press-releases/cyber crime-cost-businesses-over-2trillion

especially in terms of cyber threat. This lack of a standard overarching approach to statistical assessment has led to underestimations of both the frequency and impact of risk in the industry, and has prompted calls for a more uniform approach to improve the reliability of modelling.

Thus in 2014 WEF argued for the creation of a cyber VaR that they envision will "transcend traditional investment value at risk, unifying technical, behavioural and economic factors from both internal (enterprise) and external (systemic) perspectives" (WEF, 2015: 5). WEF specify what properties a cyber VaR should have embedded in a quantification model, by arguing that organisations should understand the key risk components and the dependencies between these components, but not how to calculate it (2015:12). The key components that make up the WEF risk framework (see figure 6) to facilitate the calculation of cyber VaR are:

1. Existing vulnerabilities and defence maturity of an organisation;
2. Value of the assets;
3. Profile of an attacker.

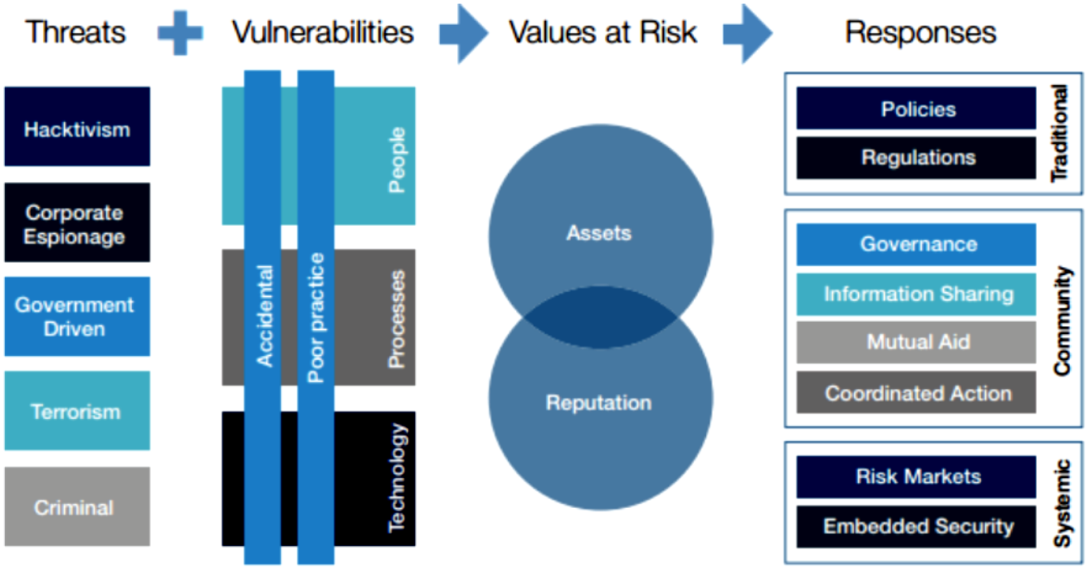

Figure 6: The WEF risk framework (WEF, 2015: 5)

This should be a relatively straightforward task for organisations to achieve if we have a standardised approach to understanding the threats, vulnerabilities, values at risk and responses available to us. However, the WEF report also asserts that the cyber VaR approach should incorporate both direct and indirect cyber security costs, as well as threat type and frequency (Hall and Ramasubramanian, 2013). This requires a more sophisticated approach to understanding the direct and indirect costs of cyber security than presented in their risk framework. What is required is a more nuanced understanding of the cost of providing cyber security, and the cost of the potential harm caused by cyber attacks. This arguably requires a better understanding of the cyber ecosystem, an overarching cyber threat taxonomy incorporating the best existing taxonomies of each aspect, and a more nuanced approach to the direct and indirect costs of cyber security. Perhaps the most recent attempt to address the direct and indirect costs involved comes from Oxford University and their cyber harm model discussed below.

From cyber threat to assets and reputation, to cyber threat and cyber harms

The challenge to produce a more nuanced account of cost associated with cyber security has been taken up in a parallel stream of research work undertaken on cyber harm at the Oxford Martin Global Cyber Security Capability Centre (2016). Here Agrafiotis et. al (2016: 2) argue that their model helps to assess VaR in Cyber (VaRiC)

> "not only in the classic sense of threat/hazard on the one hand and vulnerability/ dependency on the other, but also in terms of a more qualitative 'cascade' of harm to a wide variety of cyber-assets, stakeholders and dependencies, many of which might at first glance seem to be disconnected (actually and figuratively) from the digital world".

Agrafiotis et al. (2016:4) propose a modified version of VaRiC that aims to consider quantitative and qualitative variation in value and vulnerability in direct and indirect harm assessment. They argue that an assessment of cyber harm needs to occur before cyber risk, and therefore attempts to manage such risks. Their model reports six types of cyber harm that could be experienced by individuals, organisations, and nation states (see figure 7).

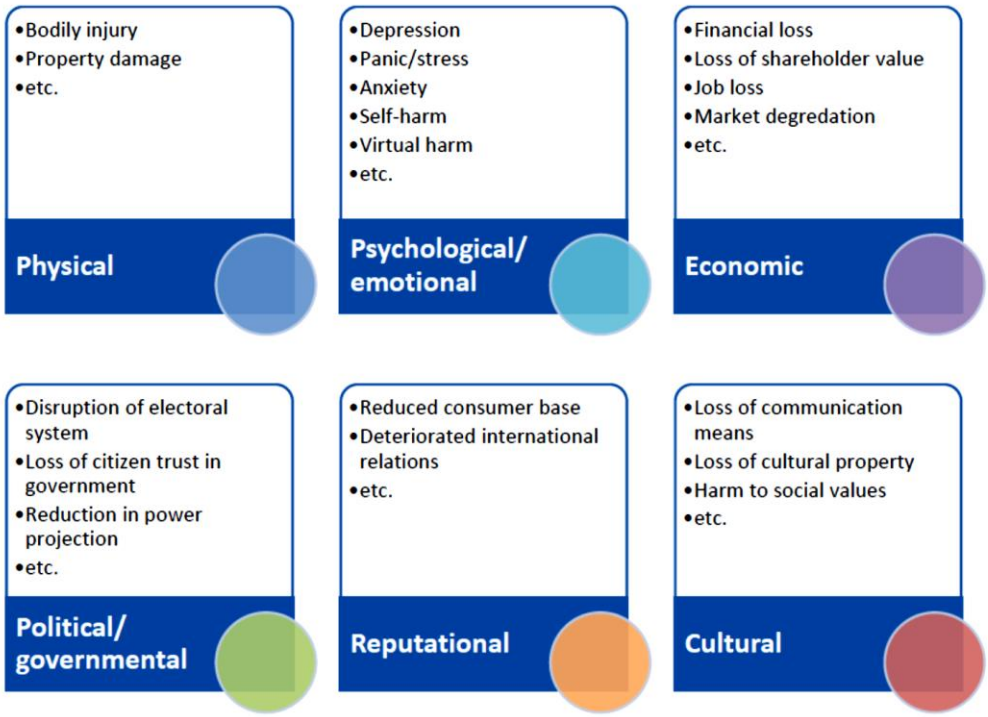

Figure 7: Types and examples of cyber harm (Agrafiotis et. al., 2016:30)

This has enriched our understanding of cyber harm, and could potentially be used to feed into a VaRiC model to provide a more accurate assessment of the risk-exposure, or more accurately 'harm-exposure', for organisations and thus strengthen cyber security. Further work is required to develop and connect these two strands of research to provide a comprehensive VaRiC approach that is of real use and value to organisations, especially those in banking and finance at the forefront of cyber threat. In terms of the existing cyber threat taxonomies and associated articles on threat actors, actual attacks, and defensive strategies and countermeasures the majority of work has been undertaken from a technical perspective that does not adequately address the governance and management of cyber security in organisations. This means that a great deal of work has yet to be done to translate the technical aspects of cyber security into managerial know how. We can start this process once we have ascertained the level of knowledge and understanding managers currently possess. This is the focus for the following section where we present our findings from some empirical research undertaken to directly ask managers what they understood personally, and their impression of their organisation's level of cyber security knowledge.

Section Two: Focus Group Evidence

In section one we have explored and reviewed the published material on cyber threat, incidents, defenses and remediation. To explore how cyber threat taxonomies are understood and utilized in organisations, a set of focus groups of managers from a range of industries were conducted in this second section of the research. Focus groups were selected as our empirical data collection method because these allow the researchers to gain insights from individual responses, in the same manner as interviews, but also additional information through interaction and debate. This enabled the researchers to assess the practical relevance and use of taxonomies within the cyber security industry, as well as ascertain the potential use-value of aligning taxonomies of threat, attack and defence measures.

In this study six focus groups were organised, each containing between 10 and 12 participants. All participants had more than five years' experience in information security management and were employed by medium to large organisations. Each focus group ran for 90 minutes and was audio recorded for subsequent analysis. Each participant signed an informed consent form agreeing to participate in the research, and in total 67 managers from a range of sectors participated (see table 2). The specific make up of each group was deliberately mixed, containing participants from different industries, employed by UK and non-UK based organisations. This decision was taken to attempt to generate a balanced and wide-ranging discussion in each focus group.

Sector	**Specific**	**Number of People**
Public	Policing	5
	Local Government	1
	National Government (MP)	1
Charity	Various	3
Financial Services	Banking	5
	Insurance	2
	Brokers	3
IT	Web Services	7
	Cyber and IS	10
	Data Management	2

Sector	Specific	Number of People
Schools	Private	1
	State	1
Engineering	National Infrastructure	2
	Manufacturing	2
Utilities	Electricity	2
	Water	5
	Gas	1
Sports Clubs	Premier League	10
	Rugby	2
	Tennis	1
Entertainment	Media	1
Marketing and Comms	Various	3
Retail	Clothing	1
	Electrical	1
Training	Technical	1
TOTAL		67

Table 2: Breakdown of participants

Process

First, a briefing session was conducted to ensure that each focus group followed the same structure and had the same understanding on the key objectives as well as the discussion guidelines. The focus group questions were open-ended and were designed to get participants to drive the discussion. Discussions adopted the structure used in Ahmad et al (2012), and focused on three main areas: 1) start of the focus group where participants got to know each other and the area of research was introduced, 2) the discussion – where focus group members were guided to discuss the topics associated with cyber threat and taxonomies, in association with strategies used to mitigate those threats, and 3) the conclusions of the session – where the researcher wrapped up the session by summarising the discussion and thereby checking to see if the debate had been captured accurately. The researchers only interrupted the discussion to guide the conversation back on to the focus

of the question, but generally allowed participants to explore the various aspects around information security strategies and taxonomies.

The recorded material was transcribed verbatim, and analysed using thematic content analysis (Krippendorff 1980; Miles and Huberman 1994), facilitated by QSRNVIVO software. Data was classified according to the research design parameters, and the two authors assessed the evidence and discussed assessments. In most instances, there was significant agreement between the authors on how the different mechanisms were classified. Additionally, a list of observations was developed for each focus group discussion pointing to subtle nuances or departures in the perspectives of the participants compared to the advice in literature.

Findings

This section presents the empirical findings from the focus groups. The discussions focused on an initial introduction to the terminology of cyber security to assess the levels of understanding and engagement with the plethora of terms used in the field. The first notable issue was with 'taxonomy' itself, which many participants were unsure about. This prompted general discussion in each of the focus groups about the language used, with the overwhelming majority stating that the terminology was confusing, overly technical, and unhelpful. Many voiced calls for the use of 'plain English' and consistency in cyber security, more practical examples, and standardised language "like terrorist alerts" (Respondent 23). Once the meaning of 'taxonomy' had been explained to the groups all participants saw the benefits of having a universal taxonomy for cyber threat, and expressed surprise this was not in place already.

It is important to recognise this early confusion as it provides a concrete example of an issue that was raised throughout the focus group sessions regarding the language used in cyber security. Leading managers from across different industries bemoaned the lack of consistency and stated repeatedly that they felt uneasy about voicing such confusion and lack of understanding for fear of looking incompetent in their role.

When asked directly about their personal, rather than their organisations', knowledge of specific cyber attacks, all participants had heard of them, and 70% felt they personally

behaved differently to different types of cyber threats. 90% of respondents, however, reported that their own organisation had no classification of a cyber risk in their registers that they were aware of. This issue was pursued in the focus groups where it was discovered that 84% said that their organisations did not classify cyber threats consciously. This was despite 20% reporting that their organisation had been successfully hacked and 2% having had their personal identity stolen. Here there appears to be a disconnect between management personal knowledge and organisational knowledge of cyber risks and cyber attacks.

It is important to note, however, that what was being reported could be individual lack of knowledge, and that in fact these organisations do have risk registers and make use of taxonomies of cyber threat. This would signal a lack of internal communication regarding cyber security, rather than the total absence of risk registers and cyber threat taxonomies. These are extremely worrying findings as it suggests that despite high-profile cases in the media, and the push from various government agencies and industry groups, the importance of understanding cyber threat and the associated risks are still not widely known, or are not widely communicated within organisations. Given the seniority of many of the participants, further exploration of this issue is required in different industries and in subsequent international comparison studies.

We attempted to discover what the participants knew about cyber attacks by asking them directly what cyber attacks their organisations had suffered, and what cyber attacks their organisations could potentially suffer from. The decision to ask specifically about cyber attacks was based on the logic that people need to know about attacks in order to understand the threat of cyber attack. If people do not know anything about cyber attacks it logically follows that their knowledge of cyber threat will be impoverished. Responses included such general comments as “the Russians are in the press a lot at the moment” (Respondent 13), “my staff do silly things” (Respondent 47), and “viruses and bugs” (Respondent 8).

When pushed for more detail all participants stated, or agreed with other group member statements, that they knew about viruses and suspicious emails, but as few as 12% understood ‘phishing’, and only 2% knew what ‘vishing’ was. 55% knew there was a distinction between an IT hack and a human hack (known as social engineering), but

struggled to provide any further information about the differences involved. This gave the strong impression that respondents had some general awareness of cyber threat, but lacked the detailed knowledge required to effectively deal with the reality of cyber threat. The consensus of opinion expressed in the focus groups was that cyber threats were increasing, that people were the weakest link in cyber security, and that technology would be introduced to improve security.

When asked for the source of such opinions respondents stated that the majority of their opinions were based on media and social media reports, and a number of respondents stating that this knowledge was a result of "things you just pick up" (Respondents 3, 21, 43, 48, 59). This again suggests that the internal communication and promotion of cyber security within organisations from across a range of industries is not at the levels required to provide cyber security.

The discussions moved on to the issue of how the individuals responded to cyber attacks, and then how their organisations respond to cyber attacks. The initial general responses to these questions was very telling, with the majority of focus groups stating "we don't know" or "we phone the IT department", both of which prompted laughter and nodding affirmations from other participants in the rooms to equal degrees. This response also prompted disclosures from some participants who had personal experience of suffering a cyber attack. Where participants had experienced this the personal anecdotes spoke of a feeling of helplessness and panic. This prompted more frank discussions of attacks and led to 23% stating that they would not know what to do in the event of a cyber attack. Although all participants either stated or affirmed agreement with statements that they would not open suspicious emails, 17% subsequently admitted to opening a suspicious email inadvertently. The overall impression from the discussion was that the majority of organisations represented either did not have plans in place, or failed to effectively communicate any such plans to managers, and therefore was not cascaded to staff.

In response to being successfully attacked the participants concerned reported that they either opened new accounts or had sought a refund for funds lost. 15% stated that they had simply written off the amount of money taken, and many others simply shrugged their shoulders. When asked specifically who they should report a cyber crime to 80% didn't know, despite 26% stating their company undertook Cyber Awareness exercises and 80%

claiming that an IT induction document was used internally to demonstrate that staff were aware of cyber security issues. This means the actual statistics quoted for cyber crime may be a lot lower than the real numbers, and that criminals may exploit this as a commercial model. Furthermore, if these percentages are a representative sample then the cyber awareness and induction sessions currently being delivered are not fit for purpose.

In terms of organisational responses to cyber attacks responses ranged from “we don’t know – it’s never happened” (Respondent 61), to “we went through it and wouldn’t want to again – it’s awful” (Respondent 26). Overall 5% confirmed they had an action plan in place, whereas 95%, including representatives of IT providers and board members, did not know what they had themselves, although 100% thought IT would be responsible and accountable for managing this. This indicated that generally companies and individuals in our sample either do not have plans in place to deal with a cyber attack, or do not know about the existence of such plans, and assume that IT will sort it out. Whilst all agreed cyber safety is a personal responsibility the majority did not report adopting this attitude at work. Interestingly most participants appeared to believe that there were fewer issues at home, and less risk, although they were not sure why. This finding deserves further study.

The final section of the focus group discussions considered what an organisation is currently doing to monitor and assess cyber threats and attacks they are exposed to, and would they value more information in the future. Participants were first asked if their organisations engaged in any analysis of cyber attacks. Responses to this question ranged from “No” (Respondents, 3, 9, 14, 18, 23, 25, 34, 51, 54, 62, 66, 67), to “as IT providers we update software to repel attacks, but don’t monitor them as they are so varied and frequent” (Respondent 21). Each of the focus groups were then asked for any participant to state if they were aware of their organisation engaging in analysis of cyber threats and cyber attacks. There were no affirmative responses. This was a very surprising response which suggests that analysis does not really happen in these organisations, or again that such information is not communicated internally. This also suggests that there is a widespread lack of information sharing between organisations, often in the same industries. No respondent reported knowledge of any organisation they knew of and interacted with who engaged in threat and cyber attack analysis. When asked the reason for this lack several respondents (Respondents 25, 28, 30, 36, 41, 42, 55, 58) stated that they believed that

other organisations probably did do this 'kind of thing', but that they would not talk about it as knowledge of the lack of analysis may be detrimental to an organisation's reputation.

The final round of questions in the focus groups concentrated on the future and what the participants would like to see develop in cyber security. Overwhelmingly participants stated that more publicity was needed for businesses and citizens and that it would be a good idea to share knowledge of cyber threats and organisational experiences, but 86% said this would not even happen in their own organisation. This suggests that even when individuals and organisations are fully aware of the need they remain hesitant to act accordingly. The potential range of reasons why this is the case is beyond the scope of this work, but unless this is addressed we are unlikely to witness any meaningful change.

90% thought it a good idea to have a real time alert system in place, and 65% felt the Police should have a central database and helpline. Although these represent the majority of views expressed, the participants were unsure as to what was currently available and how they could get access to such information in a cost-effective manner. Overall there was a palpable sense of professional managers wanting to know more about all the aspects of real cyber threats that they personally, and their organisations, face.

Summary and conclusions

The information provided above in the analysis of respondents needs to be taken in context. The sample of 67 managers from a range of industries (see table 2) should not be taken as a representative sample, as the number is too small and the selection of participants was somewhat opportunistic. Also we have to accept that respondents report their own understanding, and as was indicated earlier these self-reports may indicate personal lack of knowledge rather than the non-existence of organisational policies, procedures, and practices. That being said, the findings from this study do provide valuable insights into the current state of knowledge and understanding of cyber threats and cyber attacks from a range of managers from different industries operating in the UK. Our findings thus provide a 'snapshot' that suggests areas that need detailed further exploration.

The participant responses demonstrate that there is still a tremendous amount of work to be done to raise levels of knowledge and understanding of cyber threat in the UK.

Participants, some of whom are senior managers, responses can be summarised as wanting a more consistent approach to cyber threat to be presented in plain English to avoid confusion. Also respondents demonstrated a surprising lack of knowledge of cyber attacks, monitoring, reporting, and mitigation strategies and practices, which suggests a larger problem in cyber security. As stated earlier this could be the result of poor communication, both within and between organisations in the cyber security ecosystem. This is itself problematic and requires new approaches to improve what is currently being done in these spaces. It could also be the case that the widespread adoption of cyber security practices themselves has yet to occur, and this proposition is very concerning for cyber security professionals.

One potential explanation for this is the identifiable bias towards IT and technology in general. Respondents' overwhelming first response was to view cyber security as an IT issue, and yet after discussion the human elements of cyber security became more apparent. This is demonstrated in our findings in terms of the acknowledgement of the need to take personal responsibility, in action and communication, but a failure to do so. It also is apparent as the managers lacked knowledge and understanding despite induction courses, and in some cases cyber awareness schemes. Here the quality and quantity of training, development, and education has to be questioned. Respondents reported feeling panic, the 'awfulness' of cyber breach, and a total lack of knowledge of what to do and who to report incidents to. This needs to be urgently addressed to improve cyber security in the UK.

On the more positive side, respondents all spoke about the value of knowledge sharing, straightforward and consistent approaches to cyber threats, and the desire to know more about cyber threats and what they can do about them. These need to be nurtured and facilitated by organisations and cyber security professionals to improve our current situation. The final section of this paper provides some guidance as to what steps organisations need to take by demonstrating the utility of the cyber security ecosystem and the new taxonomies presented in section one.

Section Three: Cyber Resilience in the Face of Cyber Threats

It is clear from the review of the established knowledge, and from the focus groups' respondents, that there is a great deal of knowledge available but this is not yet being

applied in practice, and/or not *known* to be applied in practice. Our respondents described the terminology used in cyber security as confusing and inconsistent, and wanted to have a simple, and easy to understand classification of the cyber threats facing themselves as individuals and the organisations that they work for. It also became clear that a 'top line' classification of the relationship between cyber threat, actual attacks, potential defenses, and possible harm caused would be useful. Therefore what we present below is the culmination of these insights in terms of a new cyber threat taxonomy (figure 8), and a revised cyber attack taxonomy (figure 9), where the cyber attack taxonomy develops the cyber threat taxonomy by presenting the stages of cyber attack that have to be moved through to produce cyber harm. The utilisation of these taxonomies across different industries enable clear communication and facilitated understanding, thereby addressing the concerns of respondents for a clear and overarching approach to cyber threat. By using these 'top line' categories as a standard approach it becomes possible to talk, in simple terms, about cyber security and the ways to increase organisational resilience to cyber attacks. This should enable organisations to become less vulnerable to cyber attack because the organisation can build and mobilise defences according to a more nuanced understand of the threats and stages of attack. We know that there are a range of different sub-categories, or sub-taxonomies, that can be located within this design to enable each organisation to produce its' own unique map of cyber security.

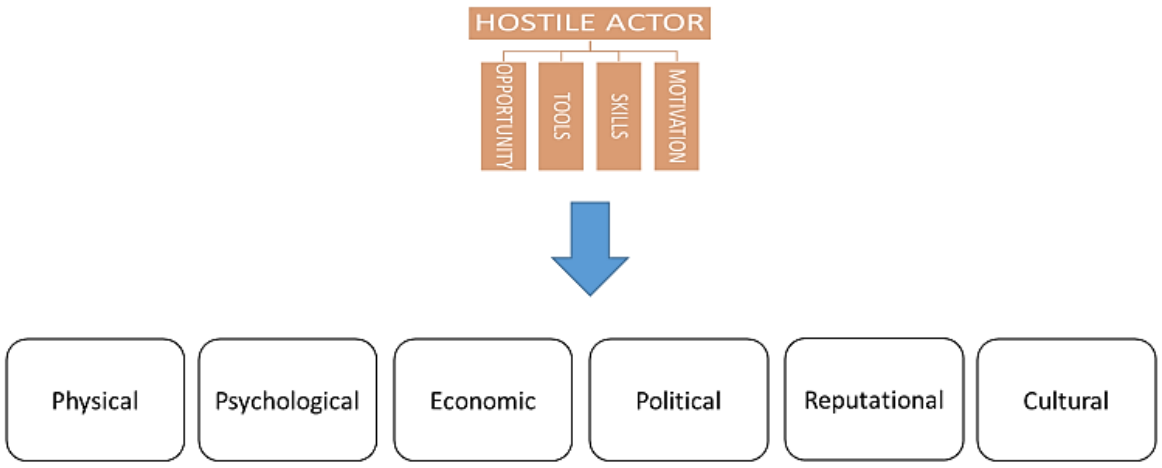

Figure 8: Cyber Threat Taxonomy

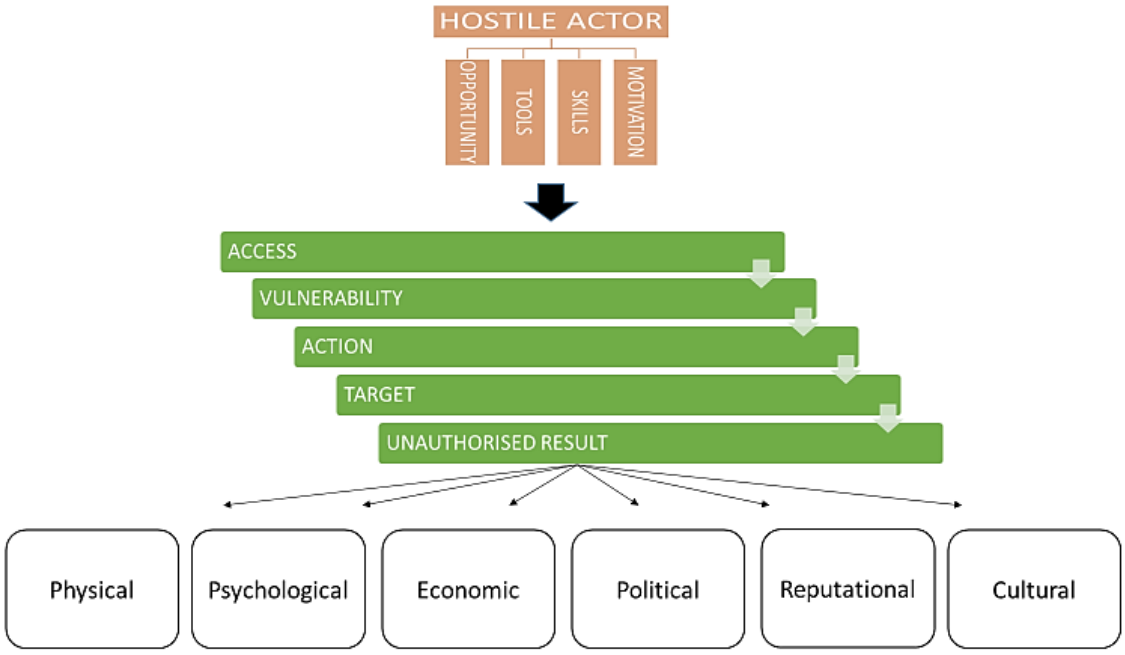

Figure 9: Cyber Attack Taxonomy

For example, we know that within the 'hostile actor' category there are eight sub-categories. Each of these actors has specific characteristics, skills, tools at their disposal, and motivations. We can insert each of these into the cyber attack taxonomy, and likewise our own organisational characteristics (defensive strategy, position, methods and operations), and our own potential targets. We can also map out the unauthorised results of successful breach, and the potential harms caused by breach. The added benefit is that this one taxonomy can be used to analyse cyber attack campaigns over time to assess how resilient our organisation is, despite changes to the particular actors and attack methods.

This idea of 'cyber resilience' is a concept that has recently emerged in the cyber security field as a result of the recognition that the traditional understanding of defence in cyberspace, built upon the notion that "a system that must defend against all possible attacks", (Schneier, 2006) is unrealistic in what is a rapidly evolving threat landscape containing increasingly sophisticated levels of cyber attacks (NIST, 2014). Existing IT security models like NIST emphasise security controls, in the face of increasing levels of cyber threat, but say little about what to do when controls fail. In the case of security control failure it is

the organisation with the failure that has to respond and recover as fast as possible and this is the aim of cyber resilience. Given the emphasis on the amount and range of cyber harm that is being threatened it is important for organisations to have a single goal to aim for, in order to assess where they currently stand in the face of cyber threat. Cyber resilience has been promoted in these terms for organisations, critical infrastructures, and nation states across the globe (Gov.UK, 2013; EU, 2013).

Organisations like the European Network and Information Security Agency (ENISA), the EU's cyber security organisation, have prompted the development of our understanding through promoting ideas such as 'end-to-end resilience' (ENISA, 2011), where organisations are expected to be able to cope with cyber incidents from minor annoyance that can be addressed through standard operations, through to major disruption that require adaptation and change. ENISA argue that at the organizational level resilience is thus "an additional dimension of holistic strategic approach…… [with the result that]….. the organisation gains a new capacity to deal with different, even sudden and extreme shocks" (ENISA, 2011: 13).

Since our research in section two suggests that even minimal levels of awareness appear to be lacking in a range of UK organisations this goal seems a long way off, but nevertheless it is incumbent upon organisations of all sizes and in all sectors to know what steps can be taken to improve their own situation. The first step is to understand the relationship between the hostile actors in the cyber ecosystem and the harm that can be caused by successful cyber attack. The cyber threat taxonomy (see figure 8) and the cyber attack taxonomy (see figure 9) enable organisations to become familiar with a simple way of understanding both the nature of threats and the process of cyber attacks. These headline categories allow organisations to insert the specific elements in each category as they relate to their own experience. This requires an organisation to engage in monitoring the cyber security ecosystem themselves, or through trusted third parties, to discover who can attack them and how. Once this task has been undertaken the organisation needs to engage in an internal audit to assess its own cyber security position, understood as its level of cyber security resilience (see table 3).

Stage 1: Non-existent Cyber Resilience	**Stage 2: Immature Cyber Resilience**	**Stage 3: Established Basic Cyber Resilience**	**Stage 4: Reactive Cyber Resilience**	**Stage 5: Fully Proactive and Reactive Cyber Resilience**
Only Generic Capability associated with 'business as usual'	Generic Capabilities	Generic Capabilities	Generic Capabilities	Generic Capabilities
	Ordinary Defensive Capability	Ordinary Defensive Capability	Ordinary Defensive Capability	Ordinary Defensive Capability
		Internal Monitoring Capability	Internal Monitoring Capability	Internal Monitoring Capability
			External Monitoring Capability	External Monitoring Capability
			Extra-ordinary Capability	Extra-ordinary Capability
			Reactive Dynamic Capability	Reactive Dynamic Capability
				Proactive Dynamic Capability
				Future Proofing
				'Hacking Back'

Table 3: A Knowledge-based Cyber Resilience Framework (Ferdinand, 2015)

As indicated in Ferdinand (2015) organisations have generic capabilities associated with operating as usual within their industries, and may have ordinary defensive capabilities (ODCs). ODCs are the front line in the cyber defence of an organisation and the most basic level required to build cyber resilience. There is some debate as to what this most basic level should contain, but a good starting is the UK 'Cyber Essentials' scheme (https://www.gov.uk/government/publications/cyber-essentials-scheme-overview). Cyber Essentials argues that the large majority of basic yet successful cyber attacks against UK businesses and citizens could be mitigated by full implementation of the following controls: "Boundary firewalls and internet gateways; secure configuration; access control; malware

protection; and patch management". These technical controls are based on cyber defence standards, such as the ISO/IEC 27000 series and the Information Security Forum's standard of good practice.

To implement these controls an organisation must thus have the capability to determine the technology in scope as part of its resourcing strategy, in association with organisational controls that take into account the human factors in cyber security. To assist in this process an organisation can use the cyber attack taxonomy (see figure 9) and the associated defensive measures that can be deployed at each step of the taxonomy.

- At the 'Access' step an organisation has to determine whether physical access and/or virtual access is possible to hostile actors. This means reviewing the physical security measures in place to assess whether physical access can be obtained. This will include policies and practices associated with security card limited access to sensitive areas, the use of USB devices, zip drives, the use of own devices whilst at work, and subcontracting arrangements. In terms of virtual access the organisation should review policies and procedures in relation to their supply chain and information sharing, password protection, whitelisting, and authentication.
- At the 'Vulnerability' step the organisation should seek to limit the vulnerabilities by considering the design, implementation and configuration of hard and soft systems, including IDS.
- At the 'Action' step each of the alternatives should be examined in order to assess what limits and controls can be put in place to stop each of these actions.
- At the 'Target' step the organisation should seek to reduce the potential availability of targets for a hostile actor. The possibilities here are numerous, and should be tailored to the specific characteristics of the organisation in question.

Once in place the organisation has reached stage 2, or 'Immature Cyber Resilience'. This would be sufficient if cyber threats were simplistic and static in nature, but we know that alongside such simple attacks the cyber threat is dynamically evolving in the cyber security ecosystem, with APTs and blended attacks being used in sophisticated campaigns against organisations (National Security Council, 2009; Gov.UK, 2013). To move on to stage 3, or 'Established Basic Cyber Resilience', organisations should carefully monitor the ecosystem and their own defences to map the changing nature, frequency, and level of sophistication

of cyber threat. If organisations do not learn about changes to threat agents and activities they will become vulnerable to sophisticated APTs. This lack of learning may also include not taking into account updates from the environment of even the most basic kind, like patch management updates.

If the organisation is not monitoring its internal environment then patch updates may be ignored, untrained staff may be allowed access, configurations can become insecure and firewalls may be bypassed without the organisation's knowledge. Organizations may thus be successfully breached without the organisation knowing about it, which may inhibit the organisation's capability to deliver products and services effectively and efficiently. The cyber attack taxonomy presented in this report can also be useful here as part of attack campaign analysis, as the identification of a cyber attack at a particular step suggests that previous steps have been ineffective. It is then a matter of reviewing the prior steps in the taxonomy to see what aspect of the organisational defences has been at fault, and correcting this to make the organisation less vulnerable to subsequent attacks. Many organisations will not wait to be attacked to assess this situation, but rather will engage a penetration testing organisation from the cyber security ecosystem to test their defences. This marks an organisation's development towards greater sophistication, and greater cyber resilience.

Organisations seeking to develop cyber resilience to a greater level of sophistication need to incorporate organizational learning about the internal environment and external ecosystem, to create, buy in, and/or reconfigure and modify internal and external monitoring functions and ODCs. This can be achieved by either engaging in deliberate learning or by employing third party specialists from the cyber security ecosystem. Organisations would also be expected to have introduced properly configured IDS and countermeasures, according to their understanding of the cyber attack taxonomy presented earlier.

In addition, organisations would as a minimum have to create response and recovery capabilities to be able to respond to cyber incidents. These can be described as extra-ordinary capabilities (EOCs), as they are not related to normal day to day activities of an organisation, and would align with other crisis related capabilities that organisations already have. Organizations would therefore have to modify and/or create crisis capabilities which would involve an emergency plan, the identification of key people to mitigate the crisis, and

training of staff to raise levels of awareness and knowledge of response policy and practice. In the event of a major cyber disruption organisations should have contingency plans in place to recover in the fastest time possible. Organisations who have reached this stage 4 are said to have achieved 'Reactive Cyber Resilience'.

The final stage 5, 'Fully Proactive and Reactive Cyber Resilience', describes a situation where an organisation is proactive in its cyber security and fully engaged in its cyber security ecosystem. Organisations at this highest level of cyber resilience proactively intervene in the ecosystem to engage in prevention of cyber attack, thereby nullifying cyber threat, are described as having Dynamic Cyber Resilience Capability (DCRC). Examples of proactive action includes staff awareness training, the creation of 'air gaps' around vital assets, cyber attack campaign analysis, employing encryption technologies, active defence such as 'hacking back' (Denning, 2014), penetration testing to assess the robustness of cyber defences, and intelligence based research on the external environment that seeks not only to respond to the environmental turbulence but also to influence the externalities. The range of possibilities, and detailed discussion, of these options is beyond the scope of this report but is available in the wider cyber security ecosystem.

Section Four: Conclusions and Recommendations

The new Cyber Threat Taxonomy, Cyberattack Taxonomy, and Knowledge-based Cyber Resilience Framework presented here provide the foundational models for a common language in cyber security. Managers can use these models to assess their own stage of development, the options available within the cyber security ecosystem, and thus make more informed decisions as to resource deployment and procurement to build cyber resilience. It also allows a manager to review the organisation's cyber resilience in relation to the NIST IT Security Maturity Model in a more nuanced way by locating the policies, procedures, implementation, testing and integration levels of the NIST model within, and across, each of the five stages of the Cyber Resilience Framework. This encourages a holistic understanding of cyber resilience that incorporates IT security, as the framework presented includes response by an organisation, through incorporating EOCs triggered when security controls have been proved to be ineffective. Adopting these models across industries would

enhance our understanding of cyber security and enable managers to improve communication, coordination, governance, and recovery when managing cyber security.

The research also shows a number of issues that need to be addressed to improve our knowledge and understanding of cyber security. Despite the high-profile cases in the media, and the push from various Government agencies and industry groups, the importance of understanding cyber threat and the associated risks are still not widely known, or are not widely communicated within organisations. The overall impression from the discussions was that most of the organisations represented either didn't have plans in place, or failed to effectively communicate any such plans to managers. Given the seniority of many of the participants further exploration of this issue is required.

We also have to recognise that if senior managers within organisations have not had information communicated to them, then there is little chance that staff would be aware. If these percentages are a representative sample then the cyber awareness and induction sessions currently being delivered are not fit for purpose. Likewise, where organisations stress personal responsibility for cyber security most of our respondents didn't report adopting this attitude at work. Overwhelmingly participants stated that more publicity was needed for businesses and citizens and that it would be a good idea to share knowledge of cyber threats and organisational experiences, but 86% said this wouldn't even happen in their own organisation. This suggests that even when individuals and organisations are fully aware of the need they remain hesitant to act accordingly. The potential range of reasons why this is the case is beyond the scope of this work, but unless this is addressed we are unlikely to witness any meaningful change.

Finally, it became clear through the research stage that organisations still lack the metrics required to make informed decisions in cyber security. Moves like the VaRiC provide a basis for developing economic measures, but the insights from the work conducted in cyber harm show that organisations need to consider far more than the economics of cyber crime. Here leading organisations and industry bodies can make a significant contribution by becoming real research partners, and by sponsoring targeted research to produce practically useful models and guidance. SWIFT Institute has taken a lead in this process but more needs to be done, especially in converting abstract and relatively inaccessible academic work into actionable knowledge for managers.

To improve this situation, beyond conducting further research, we can suggest that organisations need to become fully aware of all aspects of the cyber security ecosystem. This means not only knowledge of cyber threats from hostile actors, but also of third-party organisations able to provide the skills and capacity that is required to become fully cyber resilient.

Appendix A – Systematic Literature Review

Systematic reviews are different from traditional reviews because they adopt a replicable, scientific and transparent process. This aims to minimize bias through exhaustive literature searches of published and unpublished studies and by providing an auditable trail of decisions, procedures, and conclusions (Cook, Mulrow and Haynes, 1997). We followed the SLR process advocated by Denyer and Tranfield (2009) and Macpherson and Jones (2010) using four databases including Business Source Premier, ABI- INFORM, Emerald and ProQuest.

Our initial search concentrated on articles published in the following subject categories associated with the cyber security discipline: 'Cybersecurity', 'Cyber Security', 'Information Security', 'Cyber Threats', 'Cyber Attacks', and 'Information Security Threat' as the primary search categories. This initial search generated over 20,000 articles. We then limited our search to 2011-2016 publications which generated 8,675 returns. Third, we added the search terms 'Banking and/or Finance' which reduced the returns to 132. After assessing these articles, we further reduced the number to 113 by adding the criterion of published in English. Our first review discovered that these articles did not fully address specific threats, or taxonomies, but rather provided general information at differing levels of academic quality. It was then decided to refine our search parameters and redo the search.

Our second systematic review refined our Boolean search terms to 'cyber threat', 'cyber security threat', 'information security threat', 'computer security', and 'taxonomy'. The decision to include 'taxonomy' in our Boolean search terms was taken to conform to guidance that the production of a successful taxonomy should satisfy several requirements for its universal acceptance (Howard and Longstaff, 1998, as cited in Simmons et. al., 2014: 5), including the following typical requirements:

- Accepted – builds on previous work that is well accepted.
- Mutually exclusive – each attack can only be classified into one category, which prevents overlapping.
- Comprehensible – clear and concise information; able to be understood by experts and those less familiar.
- Complete/exhaustive – available categories are exhaustive within each classification, it is assumed to be complete.
- Unambiguous – involves clearly defined classes, with no doubt of which class

an attack belongs.

- Repeatable – the classification of attack should be repeatable.
- Terms well defined – categories should be well defined, and those terms should consist of established terminology that is compliant within the security community.
- Useful – use and gain insight into a particular field of study, particularly those having great interest within the field of study.

Although these requirements remain complete to propose a successful taxonomy, Simmons et al (2014) slightly update these requirements to include *application*. Application is presented as a requirement that further enhances the usefulness of a taxonomy, thereby enabling the taxonomy to be used in a more useful and knowledgeable manner, within a system repository. We therefore accepted this revision and included application as a criterion for our review. This produced 37 articles based on coherence to our task.

To reduce human error and bias, systematic reviews employ what are known as 'data-extraction forms'. These contain general information (such as the title, author, and publication details), study features and specific information (such as details and methods) and notes on emerging themes coupled with details of synthesis. Data-extraction forms serve at least three important functions. Firstly, the form is directly linked to the review question and the planned assessment of the incorporated studies that can include providing a visual representation of these assessments. Secondly, the extraction form acts as a historical record of the decisions made during the process. This allows third parties the opportunity to examine the process, and decisions made therein, to assess the credibility of findings. Finally, the data-extraction form is the data-repository from which the subsequent analysis will emerge (Clarke and Oxman, 2001).

To produce the data-extraction forms we examined the abstracts and conclusions of the articles returned in our search to assess the relevance of the article. Due to the complex nature of the task before us each author independently coded papers on the agreed criteria, and discussed any differences in coding before producing our final review sample of key texts summarised below.

References

Ahmad, R., Yunos, Y., Sahib, S., and M. Yusoff (2012) 'Perception on Cyber Terrorism: A Focus Group Discussion Approach', Journal of Information Security, 3, pp. 231-237.

Abliz M. (2011) 'Internet denial of service attacks and defense mechanisms', University of Pittsburgh, Department of Computer Science, Technical report.TR-11-178;2011, as cited in Shameli-Sendi et. al. (2015).

Amer, S. and J. Hamilton, (2010) "Intrusion Detection Systems (IDS) Taxonomy – A Short Review". Defense Cyber Security, 13 (2), June.Das, S. (2015) "The Cyber Security Ecosystem: Post-global Financial Crisis", in S.¬Chatterjee, S., ¬Singh, N.P., ¬Goyal, D.P., and N. ¬Gupta (eds) Managing in Recovering Markets, Springer: London, pp. 453-460.

Clarke, M, and A. D, Oxman (Eds) (2001). Cochrane Reviewers' Handbook 4.1.4 [updated October 2001], The Cochrane Library, Oxford

Cook, Mulrow, and. Haynes (1997). 'Systematic Reviews: Synthesis of Best Evidence for Clinical Decisions', Annals of Internal Medicine, 126 (5) March, pp. 376-380.

Denning, D. (2014) "Framework and principles for active cyber defense", Computers & Security, 40: 108-113.

ENISA (2011) 'Enabling and managing end-to-end resilience', available at http://www.enisa.europa.eu/act/cert/support/incident-management (accessed 3rd June 2015)

EU (2013) 'The EU Cyber Security Strategy' available at: http://www.eeas.europa.eu/policies/eu-cyber security/index_en.htm (accessed 26th June 2015)

Geng, X., Huang,Y., and A.B. Whinston, (2002) 'Defending wireless infrastructure against the challenge of DDoS attacks', Mobile Network Applications, 7(3), pp. 213–23.

Hald, S.L.N., and J. M. Pedersen (2012) "An Updated Taxonomy for Characterizing Hackers According to Their Threat Properties", in the proceedings of ICACT Conference, February.Hall, K. and Ramasubramanian, G. (2013). "How do you measure cyber risk?" World Economic Forum. Accessed at http://www.weforum.org/agenda/2013/11/how-to-measure-cyber-risk/.

Falliere N., Murchu L.O., and Chien E. (2011) Symantec "W32.Stuxnet Dossier" Version 1.4

Gov.UK (2013) 'Inside Government. The national security strategy — a strong Britain in an age of uncertainty', available at: https://www.gov.uk/government/publications/the-nationalsecurity-strategy-a-strong-britain-in-anage-of-uncertainty (accessed 1st July,

2015).
Havelka D, Sutton SG, Arnold V. (1998) ‘A methodology for developing measurement criteria for assurance services: an application in information systems assurance’, Auditing: A J Pract Theory, 17, pp. 73–92.
Hansman, S. and R. Hunt, (2005) “A taxonomy of network and computer attacks”. Computer and Security.
Howard, J.D. and T. A. Longstaff, (1998) “A Common Language for Computer Security Incidents”. Technical report, Sandia National Laboratories.
Hutchins Eric M., Cloppert Michael J., Amin Rohan M,(2011) “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains” ICIW2011
IMG-S Integrated Mission Group for Security (2012), “IMG-S Position paper for Horiztion 2020,” IMGS, 2012
King, J., K. Lakkaraju, and A. Slagell, (2009) “A taxonomy and adversarial model for attacks against network log anonymization”. In ACM symposium on Applied Computing (SAC).
Kjaerland, M. (2005) “A taxonomy and comparison of computer security incidents from the commercial and government sectors”. Computers and Security, 25:522–538, October 2005.
National Institute of Standards and Technology (NIST) (2014) ‘PRISMA Review’, available at: http://csrc.nist.gov/groups/SMA/prisma/index.html (accessed 14th August 2015).
National Security Council, US Government (2009) 'The Comprehensive National Cybersecurity Initiative', available at: https://www.whitehouse.gov/issues/foreign-policy/cyber security/national-initiative (accessed 12th June, 2015).Meyers, C., S. Powers, and D. Faissol, (2009) “Taxonomies of cyber adversaries and attacks: a survey of incidents and approaches,” Technical Report, Lawrence Livermore National Laboratory.
Mirkovic, J. and P. Reiher, (2004) “A Taxonomy of DDoS Attack and DDoS Defense Mechanisms”, in ACM CCR, April.
Oborne, M.W. (2004) The Security Economy. OECD, Paris:ISBN 92-64-10772-X.
Peng, T., Leckie, C., and K. Ramamohanarao (2007) ‘Survey of network-based defense mechanisms countering the DoS and DDoS problems’, ACM CS 39(1), pp. 1-42
Pfleeger, S.L., Rue, R., Horwitz, J., and A. Balakrishnan, (2006) “Investing in cyber security: The path to good practice," The RAND Journal, 19(1).
Porras, P., Saidi, H., and V. Yegneswara, “An Analysis of Conficker's Logic and Rendezvous Points”. Malware Threat Center. SRI International Technical Report, February 2009
Schneier B. (2006) ‘Beyond fear’, New York: Springer

Shameli-Sendi, A., Pourzandi, M., Fekih-Ahmed, M. and M. Cheriet (2015) 'Taxonomy of Distributed Denial of Service mitigation approaches for cloud computing', Journal of Network and Computer Applications, 58, pp. 165–179

Simmons, C., Shiva, S., Bedi, H., and V. Shandilya (2013) "ADAPT: A Game Inspired Attack-Defense And Performance Metric Taxonomy", IFIP Advances in Information and Communication Technology.

Simmons, C., Shiva, S., Bedi, H., and D. Dasgupta, (2014) "AVOIDIT: A Cyber Attack Taxonomy", Symposium On Information Assurance, 2014

Vries, J.D. and Hoogstraaten H. and Berg, J.V.D. and Daskapan S,. (2012) 'Systems for Detecting Advanced Persistent Threats'. CyberSecurity. 54-61, IEEE Computer Society

Wang, J.W. and Ronga, L.L. (2009) "Cascade-based attack vulnerability on the US power grid", Safety Science, 47(10), pp. 1332–1336.

William, S. (1994) 'A Taxonomy of Computer Program Security Flaws, with Examples', ACM Computing Surveys, 26,3.

Wood, A.D., and Stankovic, J.A. (2004) 'A taxonomy for denial-of-service attacks in wireless sensor networks', in the Handbook of sensor networks: compact wireless and wired sensing systems, pp.739–63.

World Economic Forum and Deloitte (2015). Partnering for Cyber Resilience Towards the Quantification of Cyber Threats. Cologny/Geneva: World Economic Forum.

Wu, Z., Ou, Y. and Y. Liu (2011) "A Taxonomy of Network and Computer Attacks Based on Responses", International Conference of Information Technology, Computer Engineering and Management Sciences.

Zargar, S.T., Joshi, J., and D. Tipper (2013) 'A survey of defense mechanisms against distributed denial of service (DDOS) flooding attacks', IEEE CST, 99, pp. 1–24.