

Practitioners' Views on Cybersecurity Control Adoption and Effectiveness

Louise Axon
louise.axon@cs.ox.ac.uk
University of Oxford
UK

Arnau Erola
arnau.erola@cs.ox.ac.uk
University of Oxford
UK

Alastair Janse van Rensburg
alastair.jansevanrensburg@cs.ox.ac.uk
University of Oxford
UK

Jason R. C. Nurse
J.R.C.Nurse@kent.ac.uk
University of Kent
UK

Michael Goldsmith
michael.goldsmith@cs.ox.ac.uk
University of Oxford
UK

Sadie Creese
sadie.creese@cs.ox.ac.uk
University of Oxford
UK

ABSTRACT

Cybersecurity practitioners working in organisations implement risk controls aiming to improve the security of their systems. Determining prioritisation of the deployment of controls and understanding their likely impact on overall cybersecurity posture is challenging, yet without this understanding there is a risk of implementing inefficient or even harmful security practices. There is a critical need to comprehend the value of controls in reducing cyber-risk exposure in various organisational contexts, and the factors affecting their usage. Such information is important for research into cybersecurity risk and defences, for supporting cybersecurity decisions within organisations, and for external parties guiding cybersecurity practice such as standards bodies and cyber-insurance companies.

Cybersecurity practitioners possess a wealth of field knowledge in this area, yet there has been little academic work collecting and synthesising their views. In an attempt to highlight trends and a range of wider organisational factors that impact on a control's effectiveness and deployment, we conduct a set of interviews exploring practitioners' perceptions. We compare alignment with the recommendations of security standards and requirements of cyber-insurance policies to validate findings. Although still exploratory, we believe this methodology would help in identifying points of improvement in cybersecurity investment, describing specific potential benefits.

CCS CONCEPTS

• Security and privacy → Systems security; Network security; Human and societal aspects of security and privacy.

KEYWORDS

Cybersecurity Risk, Control Effectiveness

ACM Reference Format:

Louise Axon, Arnau Erola, Alastair Janse van Rensburg, Jason R. C. Nurse, Michael Goldsmith, and Sadie Creese. 2021. Practitioners' Views on Cybersecurity Control Adoption and Effectiveness. In *The 16th International Conference on Availability, Reliability and Security (ARES 2021)*, August 17–20, 2021, Vienna, Austria. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3465481.3470038>

1 INTRODUCTION

In order to mitigate the harms resulting from cyber-attacks, security practitioners working in organisations seek to strengthen defences and improve resilience, yet there is not necessarily a clear path to achieving this. In 2019, organisations were purchasing an average of three new cybersecurity products annually, mainly chosen to address threat trends, despite 70% of organisations performing one or fewer revisions of their cybersecurity posture in the same time frame [6]. This suggests that most organisations spend little time assessing their cybersecurity exposure; and for those that do, it is difficult to make decisions on which controls to adopt, or to assess whether different controls could provide better risk reduction.

There is a need to better understand the real value of risk controls: for organisations looking to prioritise security deployments and strengthen defences, and to support research in the area of cybersecurity risk and defence. Insurance and audit are key drivers of good cybersecurity practice, and a strong understanding of the value of risk controls is also vital to supporting external parties such as cyber-insurance providers, who insure businesses against cyber-risk, in accurately assessing and making recommendations for reducing cyber-risk [20]. There are numerous standards, such as the Center for Internet Security's 20 Critical Security Controls (CIS CSCs) [11], NIST Cybersecurity Framework (CSF) [21], ISO 27002 [1], and Cyber Essentials [5], aimed at facilitating information-security management by organisations. The value of deploying controls and complying with standards is not easy to quantify. Furthermore industry, for example cyber-insurers, increasingly suggest products and approaches for reducing cyber-risk [19]. This can influence take-up, and it is important that such guidance is supported by up-to-date evidence on control effectiveness.

A wealth of experiential knowledge is possessed by the cybersecurity practitioner community on the effectiveness of cybersecurity controls, and on the frequency with which and ways in which cybersecurity controls are deployed in organisations. The aim of our research is to synthesise this knowledge by identifying and



This work is licensed under a Creative Commons Attribution International 4.0 License.

ARES 2021, August 17–20, 2021, Vienna, Austria
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9051-4/21/08.
<https://doi.org/10.1145/3465481.3470038>

analysing the views of security practitioners on the impact of using security controls on reducing the risk exposure of an organisation. We also aim to explore the ways in which organisations actually deploy risk controls, and investigate the factors that impact on effectiveness and deployment.

Prior research has aimed to determine control effectiveness in specific organisational contexts [15, 17, 18]. Technical approaches, such as automated modelling of risk-control properties and optimisation methods, have been applied to reasoning about control effectiveness [2, 3, 8, 9]. Some qualitative studies have provided insights into the value of a limited number of specific types of control and some of the factors that are important in their deployment. The work of *Such et al.* [24] explores the views of industry stakeholders on the characteristics of 20 assurance techniques, such as vulnerability scans, red-team exercises, and configuration review. A survey was conducted to review them in terms of cost, effectiveness, personnel, experience and time required, and to determine the three most complementary controls. This information is then used to provide a cost-effectiveness measure, and the sets of controls that are more effective and more cost-effective. Similarly, *Dietrich et al.* [7] explored the opinions of system operators on system misconfiguration. They conducted interviews and a survey, aiming at determining human elements (causes) of misconfiguration, and provide recommendations to reduce their frequency and impact.

None of this prior work fully achieves our aim of understanding practitioners' perspectives of the relative effectiveness of security controls in reducing the risk exposure of an organisation. Furthermore, this work does not provide a comprehensive account of the ways in which organisations deploy risk controls. In this paper we aim to progress understanding of these topics by identifying practitioners' perceptions of the effectiveness of various risk-control setups and gathering their experiences of the ways risk controls are currently being deployed by organisations. We explore two main research questions (RQs):

RQ1 How effective do security practitioners perceive different cybersecurity controls to be in addressing organisational cyber-risk? There is a need to comprehend the real impact of security controls on the risk run by organisations of being harmed by cyber-attacks; in particular, to understand how the controls deployed, and the ways in which these controls are positioned and configured, mitigate cyber-attacks on networks and their assets, and/or reduce the harms that may ensue. We aim to collect the perceptions of security practitioners on this question.

RQ2 How are different cybersecurity controls deployed in practical environments? Identifying how organisations deploy controls, the challenges faced in doing so, and the factors deployment depends on, benefits the development of improved practical risk-control approaches. It could also aid estimation of which controls an organisation with particular characteristics is likely to have deployed. This reduces the organisation-specific data required in the estimation of cyber-risk, improving the ability to obtain accurate results from sparse data. This benefits both research into the wider picture of the cyber-risk landscape, and into the risk assessment of individual organisations, as the data gathered directly

from organisations on the controls they use and the way in which they are configured may be incomplete or difficult to obtain [27].

We examined these research questions by soliciting the perspectives of security practitioners through an online survey and interviews. While the question of control effectiveness is clearly complex and, like the deployment of controls, will vary between organisational contexts, we aimed to ascertain whether any consensus could be reached, and also to establish a stronger understanding of these concepts: the factors that contribute to them, and the ways in which they can be reasoned about, measured and studied.

In this article we present the studies and reflect critically on the findings, providing recommendations for furthering knowledge. Our contributions are as follows:

- We present the qualitative perspectives of security practitioners on the effectiveness of controls and the deployment of controls by organisations.
- We identify trends in the effectiveness and deployment of controls according to the practitioner community and present these findings in the context of organisational risk.
- We highlight the factors that impact on the effectiveness of the risk controls used by organisations, and that drive their deployment.

Our methodology is presented in Section 2, and the quantitative and qualitative results in Sections 3 and 4 respectively. We discuss the implications in Section 5, and conclude in Section 6.

2 METHODOLOGY

We conducted two studies seeking to elicit the perspectives of security practitioners on our research questions:

- (1) **Online survey.** We aimed to obtain a quantitative account of security practitioners' perceptions of the relative effectiveness of the CIS 20 CSCs (v7), and of the frequency with which organisations deploy these controls. This standard was chosen since it is widely recognised and used [10, 22], with clearly defined categories suitable for the survey. The online survey was completed by 30 participants between June and August 2019.
- (2) **Semi-structured interview.** Interviews were conducted to gather qualitative evidence of the views of security practitioners on our research questions. The aim was to delve deeper into perceptions of control effectiveness and the challenges of deployment that had been possible in the survey, through open-ended questions [12]. Seven security practitioners participated in interviews, three of whom had worked in cybersecurity at a senior level within one or more organisations, and four of whom had experience working as penetration testers.

2.1 Recruitment

We recruited participants by emailing our existing contacts in organisations that employ security practitioners. The researchers advertised the survey on their LinkedIn pages, targeting security-practitioner contacts, and on relevant mailing lists. We also attended events for executive-level security-practitioner groups to

promote the studies. By recruiting through this range of channels, through each of which a range of different security practitioners were reached, we aimed to reduce any possibilities of bias in our participant selection. Our anonymisation of the online-survey data prevents us from determining the overlap between participants in the survey and the interviews. Ethical approval for the studies was granted by our Research Ethics Committee.

2.2 Online survey

Participants' demographic data was collected, including job role, years of experience in that role, and years of experience in cybersecurity. Participants were then asked to record their views on the effectiveness of each of the CIS 20 CSCs in securing the networks of organisations by:

- Rating the effectiveness (in their opinion) of each control using a five-point Likert scale with response categories *very ineffective, ineffective, neutral, effective, very effective* (or *don't know*).
- Selecting up to three *most effective* and three *least effective* controls using drag-and-drop to the relevant boxes.

Participants were also asked how commonly each of these 20 controls is implemented by organisations by:

- Rating the frequency of deployment of each control using a five-point Likert scale with response categories: *almost never, rarely, neutral, often, almost always* (or *don't know*).

Finally, participants were offered the opportunity to record any additional comments at the end of the survey. This allowed participants to add additional information they considered relevant, such as more detailed perceptions on the control-effectiveness concept, or further insights into the factors that drive the deployment of controls by organisations.

The controls were listed in a randomised order to each participant, with the aim of reducing the likelihood of biasing responses through the consistent placement of these elements in the same area and, in particular, next to the same *most* and *least* ranking boxes. We ran a pilot of the survey with three people with relevant experience: two had worked as penetration testers, and one had performed security-compliance audits for organisations. This enabled us to refine the phrasing and structure of the survey, and the clarity of the interface.

2.3 Interviews

Interviews took place face-to-face or over the phone, to suit the needs of participants. Each was audio-recorded and lasted approximately 30 minutes.

The interviews were guided by a set of questions through which we aimed to explore the concepts of the online survey in greater depth. Participants were guided towards a discussion of real-world security practice and their perceptions of risk-control deployment in organisations, and the effectiveness of these controls; the full list of questions is presented in Appendix A. At the end of each session, participants were given the opportunity to make any additional comments.

The first set of guiding questions focused on security-control deployment: the ways in which decisions are made on control deployment and configuration, resources used to guide these decisions, the impact of broader considerations such as the interconnection of assets and aggregation of harms on these decisions, and perceptions of the most important security controls. This was followed by questions on security-control effectiveness: perceptions of the concept, methods of determining control effectiveness and data required for this, and methods of determining the residual risk of an organisation. We finished with a set of questions on perceptions of and the treatment of interdependencies between controls, and the ways in which these may impact on usage and effectiveness.

2.4 Analysis

Given discrepancy in the community as to how to treat Likert-scale data [14, 23], we analysed the survey responses by calculating the mode and median rating, and a comparison of non-neutral scores (CNNS), which takes the ratio of scores less than, versus greater than, "*neutral*". We present all three measures, which support the same conclusions.

We manually transcribed our interview recordings, producing transcripts for each discussion. After becoming familiarised with the data, the researchers coded it using Template Analysis [16]. This is a qualitative data-analysis technique that begins with coding according to a-priori themes (in this case, the interview questions). The codebook then evolves iteratively, with relevant sections of data that do not fit these existing themes being assigned new codes. We thus produced a template of codes, developed through iterative application to the dataset. We interpreted and documented the findings within the themes of the resulting template, engaging in frequent reflections to avoid bias and the influence of personal beliefs.

3 SURVEY RESULTS

Table 1 shows the demographics of participants in the online survey: twelve executive-level cybersecurity professionals, four IT Directors, and a range of security professionals working internally (security analysts, engineers and officers) and externally (security consultants and penetration testers). All participants had at least three years' experience (and a majority more than ten) working in the field of cybersecurity.

Figure 1(a) shows the responses made by participants when they were asked to rate the effectiveness of the 20 CIS CSCs. A majority of participants viewed all controls as being either "effective" or "very effective". The mode and median averages, and the CNNS are presented in Table 2, and it is clear from the strongly positive CNNS ratios that there was a high level of agreement between participants in rating all controls either "effective" or "very effective".

Figure 2 shows the CSCs that were considered "most" and "least" effective, respectively, by participants. The controls perceived to be "most effective" by the greatest number of participants were CSC3 "*Vulnerability assessment*", perceived to be "most effective" by seven participants, CSC4 "*Admin privileges*", perceived to be "most effective" by seven participants, and CSC19 "*Incident response*", perceived to be "most effective" by six participants. While this provides some indication of agreement, none of these numbers

Job Role	Count	Experience (years)	Count
CISO/ Head of Cybersecurity	12	4 or fewer	2
IT/ Technical Director	4	5–9	6
Information Security Officer	3	10–14	15
Security Analyst/ Engineer	5	15–19	3
Cybersecurity Consultant	4	20 or more	4
Penetration Tester/ Auditor	2		

(a) Job Role

(b) Years of Cybersecurity Experience

Table 1: Online survey participant demographics

Control	Description	Control Effectiveness			Control Deployment		
		Mode	Median	CNNS	Mode	Median	CNNS
CSC1	“Device inventories”	effective	effective	2:18	neutral	neutral	12:8
CSC2	“Software inventories”	effective	effective	2:17	almost never	neutral	12:7
CSC3	“Vulnerability assessment”	effective	effective	2:24	rarely	rarely	15:5
CSC4	“Admin privileges”	v. effective	v. effective	0:28	often	neutral	8:12
CSC5	“Secure hosts”	effective	effective	0:27	often	often	8:14
CSC6	“Log monitoring”	effective	effective	1:20	rarely	neutral	13:6
CSC7	“Web and email defence”	effective	effective	1:21	often	often	5:14
CSC8	“Malware defences”	effective	effective	2:18	almost always	almost always	0:24
CSC9	“Protocol controls”	effective	effective	1:21	neutral, often	neutral	10:9
CSC10	“Data recovery”	effective	effective	1:21	often	neutral	9:11
CSC11	“Secure network devices”	effective	effective	1:26	often	often	5:18
CSC12	“Boundary defences”	effective	effective	5:19	almost always	often	2:20
CSC13	“Data protection”	effective	effective	2:21	rarely, neutral	neutral	11:7
CSC14	“Access control”	effective	effective	2:23	rarely	neutral	14:8
CSC15	“Wireless access control”	effective	effective	3:18	neutral	neutral	4:13
CSC16	“Account monitoring”	effective	effective	0:20	neutral	neutral	10:7
CSC17	“Skills and training”	effective	effective	1:19	often	neutral	7:12
CSC18	“Application security”	effective	effective	3:19	rarely	rarely	15:7
CSC19	“Incident response”	effective	effective	0:26	rarely	rarely	15:7
CSC20	“Penetration testing”	effective	effective	3:22	often	neutral	13:8

Table 2: Likert Scale Results: Control Effectiveness and Deployment. CNNS denotes *Comparison of Non-Neutral Scores*. Provided descriptions are shorthands to enable ease of reference. For full descriptions, see SANS 20 Critical Security Controls [11].

represent a particularly large proportion of the 30 respondents, and it is apparent that the question divided opinion.

The controls perceived to be “least effective” were not clearly differentiated: although CSC12 “Boundary defences” was rated “least effective” by the greatest number of participants (five), many others were selected three or four times, and the margin is small. Opinions on this question were too divided to allow us to draw reasonable conclusions.

Figure 1(b) shows the responses made by participants when they were asked to rate how commonly each of the 20 CIS CSCs is deployed by organisations. It is clear that there was greater difference in opinion between participants on this question than on the previously-presented “effectiveness” question: most of the CSCs obtained a set of different responses covering the full range of the Likert scale. As is made clear by the mode average ratings given to each CSC, presented in Figure 3, there was also greater variation

in the average perception of how commonly each control is deployed by organisations (with modes ranging from “almost never” for some controls to “almost always” for others) than in the average perception of how “effective” the control is (as shown, the average was “effective” or higher across controls).

The mode, median and CNNS of the Likert-scale ratings recorded by participants are presented in Table 2. Despite the variation in opinion between participants as a whole, we can observe some areas of agreement using these statistics. Some controls were considered to be deployed “often” or “almost always” on average (mode and median) with highly positive CNNS: CSC8 “Malware defences”; CSC12 “Boundary defences”; CSC7 “Web and email defence”; and CSC11 “Secure network devices”. CSC5 “Secure hosts” was also considered to be deployed “often” on average (mode and median), although the fact that the CNNS is closer to an equilibrium indicates that this view was held less consistently. It is worth noting that the

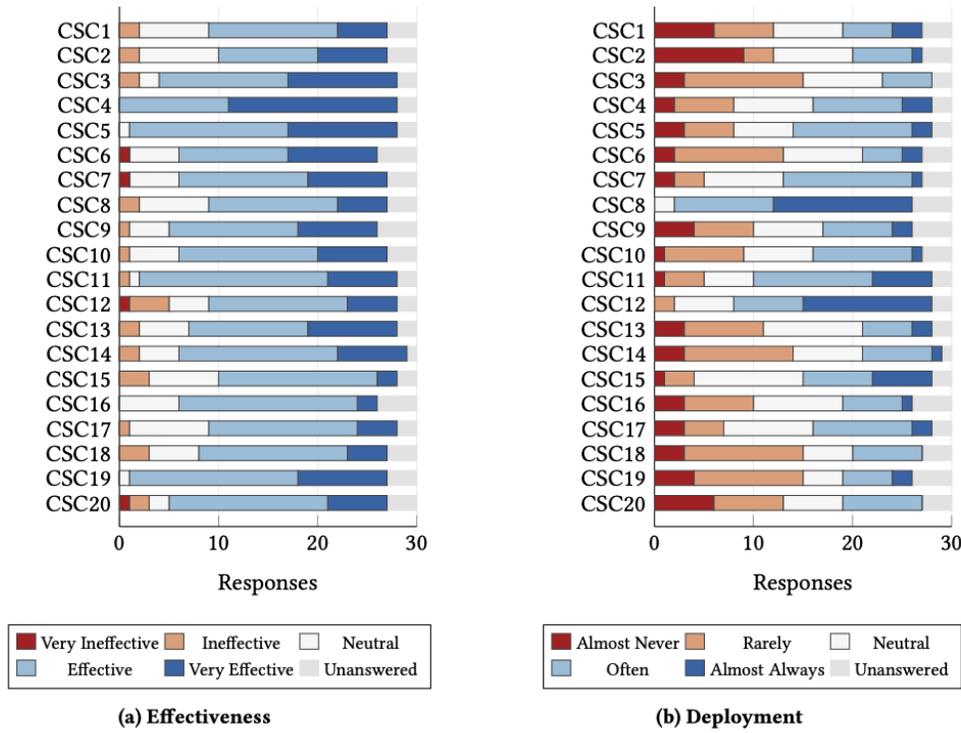


Figure 1: Likert responses regarding the effectiveness and deployment of controls

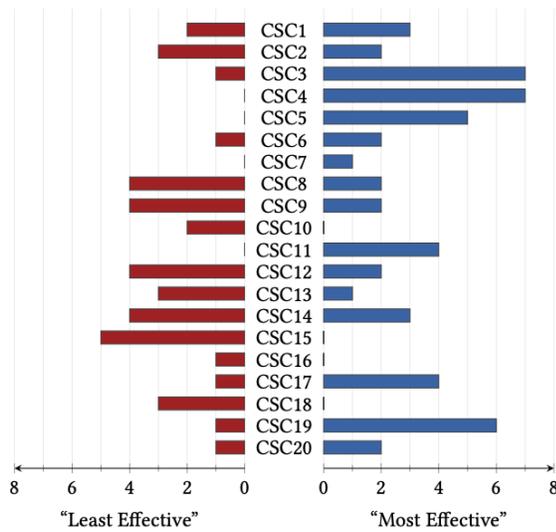


Figure 2: Count of "Most Effective" and "Least Effective" bucket placements

only control for which the CNNS indicates that the view that it is deployed "often" or "almost always" was expressed unanimously is CSC8 "Malware defences".

Some controls were considered to be deployed "rarely" or "almost never" on average (mode and median), with highly negative CNNS: CSC3 "Vulnerability assessment"; CSC18 "Application security"; and CSC19 "Incident response".

When the CSC effectiveness responses of each participant were correlated against that participant's deployment response for the same control, a weak negative correlation was found. This could indicate that participants felt that more effective controls were less likely to be deployed. Alternatively, there may be a "holy grail" effect in which practitioners expect controls to be more effective when they have less experience of them (and therefore possibly of their failings). While it is not possible to present confident conclusions from this correlation, the lack of a strong correlation gives at least some indication that practitioners' choices of controls to deploy may not be directly tied to their perception of their effectiveness. This correlation was reinforced by analysis of participants' choices of "most effective" and "least effective" controls, which indicated a similar weak negative correlation between perceived control effectiveness and perceived deployment frequency.

4 QUALITATIVE PERSPECTIVES

We present participants' views on control effectiveness and deployment: their additional comments in the online survey (labelled PS1–PS30); and responses during the interviews (PI1–PI7).

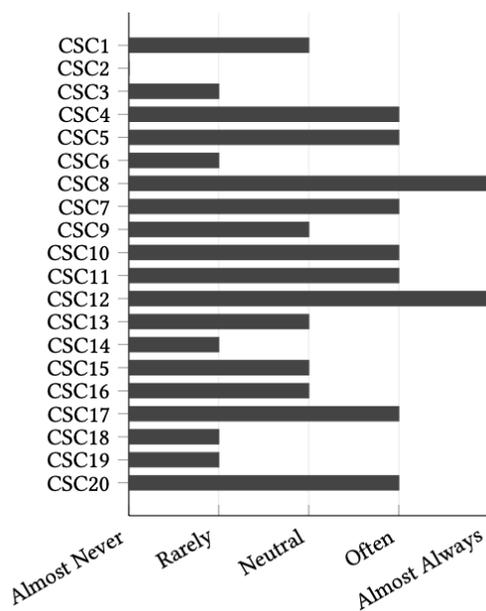


Figure 3: Modal responses to CSC deployment

4.1 Control effectiveness

Participants communicated their understanding of the concept of control effectiveness. Controls may be considered effective if they reduce an organisation’s vulnerability to being “as low as practical” (scepticism was expressed at the idea of the complete effectiveness of any control). This may mean assessing vulnerability levels relative to other organisations: “if you’re more attractive than the person next door, you’re more likely to have a burglary” (PI1).

A number of technical controls (including firewalls, encryption of data, and up-to-date anti-malware systems) and procedural controls (training and awareness of users, data-backup procedures, network- and endpoint-monitoring capability and patching processes) were widely perceived to be fundamental to security in all organisations. Control effectiveness is context-specific: a control’s effectiveness depends on numerous factors and has a propensity to vary over time: “you may have 100% control when you put it in, but you almost certainly won’t in 1 month’s/6 months’ time” (PI1).

Controls can vary in their effectiveness at protecting different assets (network-boundary defences may protect servers more effectively than clients, for example). In assessing the effectiveness of controls, potential attacker objectives are an important consideration: an attacker aiming to exfiltrate data and an attacker wanting to “cause the biggest disruption possible” might face very different challenges in trying to circumvent security controls. The threat actor being defended against also has an impact: “anti-malware is effective against ‘mass-market’ attacks but won’t help you much if you’re defending against nation states” (PS10).

The effectiveness of a control depends on the make and model used, since weaknesses may exist in controls from certain vendors. Considering the quality of a control’s implementation is critical to assessing its effectiveness: *A measure can be effective, and in theory*

it gets implemented, but the way it gets implemented... is ineffective. Examples: logs get logged but not looked into; admin access restrictions applied but access rights ‘creep up’ as the employee moves across the organisation” (PS17).

A participant noted potential influences on their perceptions: “I seem to have listed the easier to implement controls as most effective, and the harder ones as least effective... this is probably insightful to the effectiveness of these controls in a real-world environment” (PS19).

This variety of influencing factors means measuring the effectiveness of controls is complex. A possible approach is to use threat-detection platforms to measure the amount of malicious activity penetrating an organisation’s systems (and the time taken to detection and mitigation) and thus draw conclusions on the effectiveness of the control setup. Formal third-party assessments (e.g., penetration tests) are another approach. Assessments of control effectiveness should account for the importance of assets to the organisation: “has anything ever been able to get at the things that you most value in your organisation?” (PI1).

The effectiveness of controls can also be considered in terms of an organisation’s confidentiality, integrity and availability aims. Integrity was considered to be particularly difficult to measure, especially if insiders to the organisation are involved in an attack: “if they have the authorisation to alter it... it’s virtually impossible to know whether they’ve particularly done something to the integrity of the system” (PI1). Availability, on the other hand, can be measured using metrics such as the amount of network downtime.

Even after the deployment of controls, residual risk remains. This can be measured using risk-assessment templates, which may include assessments of the likelihood and impact of a harmful event occurring. It is the likelihood that an organisation has “the most chance of influencing” through the deployment of security controls, while “the impact probably stays much the same, unless your business changes or the world changes in some way” (PI1).

There is “nuance and complex interplay between controls” (PS10). This poses challenges to directly comparing their individual effectiveness, and interdependency must be taken into account. Asset inventories, for example, “only yield benefits for an already mature organisation” (PS10), who can act on the information collected; while “security awareness is only really effective if you have an incident response capability” to act on reports. In assessing controls, there is a need to obtain a holistic understanding of the security of the entire network and the way in which controls work collectively to protect it:

Very few controls in isolation prevent a particular kind of attack... they are usually reinforced by other controls. My background is in defensive security; generally speaking when I see controls failing, either in a simulated exercise or real life, it is due to a failure of multiple controls (PI3).

Controls can also compensate for each other: boundary defence, for example, is “one of the quickest controls to deploy and would offer partial mitigation against a lack of endpoint security (hardening, anti-malware)” (PS10).

4.2 Control deployment

Participants commented on the ways in which organisations, in their experience, deploy security controls: selecting the controls to

be used (often guided by standards) and implementing (configuring and placing) these controls. The segmentation and blocking of sensitive assets is handled well by some organisations possessing highly sensitive data: *“for companies that deal with security clearance work, or government work, that stuff’s locked away really well and you’d be hard-pushed to get to it”* (PI5). In other industries, the network lacks segmentation, creating security problems: *“you can just gain access anywhere”* (PI5).

The decisions made by organisations on the risk controls to be deployed, and the implementation of them, can be driven by standards and regulations, risk-assessment procedures, professional judgement, and guidance from technical vendors and partners. Standards may be chosen by organisations based on their priorities, types of system, and need for tailoring: *“a lot of them are rather generic and you really need to understand what you’re doing. I mean, ISO 27001 isn’t really going to help you unless you know what it is”* (PI1).

Regulators mandate the adoption of standards and controls by organisations in certain sectors: finance and healthcare/pharmacy were cited as two industries within which compliance requirements drive companies to have a strong adoption of security controls. The view was expressed that outside such heavily regulated industries control adoption is less advanced, and networks tend to lack to necessary segmentation: *“if they are not forced, if they have to do it by themselves, most of the time they will not put in the effort”*.

A number of factors influence the selection of controls for deployment. Companies tend to follow a cost-benefit analysis approach, investing *“to protect the crown jewels, not haphazardly but to the minimum level required”* (PI2). The need to prioritise control setups that protect the assets they *“really do not want people to get at”* (PI1) leads to differences in risk-control usage between industries. One participant gave examples of observed priorities in the defence sector:

In the defence and security sector... you’re really most concerned about the confidentiality of your data, and its integrity... they probably have been much less concerned than other companies have usually been about things that would now come under GDPR... they haven’t got much personal data, but they’ve got a lot of intellectual property and that’s what they’re concerned with, and they’ve also got a lot of state secrets (PI1).

In some sectors, the prevalence of legacy systems designed with little security in mind and with backdoors built in as a *“norm”* makes producing effective cybersecurity difficult: *“if you’re trying to put a firewall between a distributed system and an old legacy system... you practically always come a cropper”* (PI1). Finance was cited as a sector in which this problem has been commonplace, with transactions systems run on Windows XP, and which has seen improvements in recent years. The prevalence of legacy systems in the healthcare sector was also highlighted, particularly in the context of their vulnerability to ransomware.

The type of systems (e.g., information versus operational technologies) being run by an organisation impacts hugely on the deployment of controls: *“whether it’s potentially a standalone system, or a safety-critical system, or an aircraft... it’s going to really dramatically change what controls we pick”* (PI2). The security of an organisation’s network is impacted by its size: *“the larger the network the harder it becomes to apply your controls”* (PI5). On the other

hand, large organisations in certain sectors may be more advanced in their adoption of controls than SMEs due to the capability of their dedicated security teams:

The answer depends largely on the sector you are talking about and the size of the company. SMEs seldom deploy many of the controls other than having some anti-virus software, probably patching and maybe a firewall, but large companies in sectors like finance or defence and aerospace deploy almost all of them and it depends how robust and good their implementation is and the products they select (PS9).

Smaller companies often outsource the security of their systems to the cloud, which can shift responsibility for deploying controls to service providers. In the private sector, financial risk is an important driver for control deployment: *“a lot of it comes down to how they convince their shareholders that they’re making an adequate risk cost-benefit analysis”* (PI2).

5 DISCUSSION AND IMPLICATIONS

5.1 Factors affecting control effectiveness and adoption

We have obtained an account of the range of factors that can impact on the effectiveness of a control, and established characteristics of an organisation that can affect the way in which they deploy controls. Both are summarised below.

Control effectiveness.

- A control’s **implementation** – the quality of the product used, its initial configuration, and its treatment over time (updating and adaptation in line with changing networks and threats) – impact on its effectiveness.
- The **capabilities and motivations of the threat actors** a control defends against impact on its effectiveness.
- The **assets** a control defends must be considered when reasoning about its effectiveness.
- The **wider organisational context** also has an impact: the business use-cases of the network, and the size of the network and organisation.
- The **interactions and interdependencies** between controls impact on their effectiveness. There is a need to take a holistic view of controls, as controls depend on others, compensate for others, and work in combination.

Control deployment.

- Organisations prioritise the **protection of their “crown jewels”**, and controls are often deployed with a focus on this protection.
- The **sensitivity** of the assets handled by an organisation can impact on its deployment of controls: companies handling highly sensitive data may be more likely to have segmented their networks effectively, for example.
- **Sector** can impact on priorities and therefore on control deployment:
 - Organisations in the defence and security sector are likely to be most concerned with preserving the confidentiality of data.

- In the private sector reliance on the approval of stakeholders makes financial risk a key driver in the adoption of controls.
- Regulation drives the deployment of controls in certain sectors, and in organisations that carry out certain types of business.
- Finance and defence were cited multiple times as sectors in which control deployment was more advanced in general.
- The **existing infrastructure** of an organisation may impact on its control deployment: securing legacy versus more recent systems, for example.
- The existing **security maturity** of an organisation may influence its selection of cybersecurity standards, as more mature organisations are more likely to adopt more advanced standards (ISO 27001, for example).
- The **size** of an organisation may impact on their control deployment. SMEs are less likely to have deployed controls as thoroughly as larger enterprises in general, and would begin with basics such as anti-virus software, patching and firewalls (and may outsource systems and security to cloud providers).

5.2 Trends in control effectiveness and adoption

Despite the range of factors impacting on control effectiveness, our results suggest that some controls are perceived to be more effective than others in general. For instance, the survey showed that CSC3 “*Vulnerability assessment*”, CSC4 “*Admin privileges*”, and CSC19 “*Incident response*” were perceived to fall into the top 3 “most effective” controls by the most participants (although it is important to note that the margin was not large, and other controls were considered “most effective” by other participants).

We were unable to draw reasonable conclusions on the “least effective” controls due to a lack of consistency in responses across participants; this may suggest that participants found difficulty in reasoning about “least effective” controls due to a belief that all are “effective” (indeed, all controls were rated as being “effective” or higher). These facts appear to be in contradiction with the multitude of descriptions throughout the studies of control failures that had led to security events; we posit that perhaps the “ideal” version of all of these controls is considered “effective”; however in reality this ideal version is rarely deployed.

Trends emerged from the online survey suggesting that certain controls are deployed more commonly across organisations. The controls considered to be commonly deployed were: CSC8 “*Malware defences*”, CSC12 “*Boundary defences*”, CSC7 “*Web and email defence*” and CSC11 “*Secure network devices*”. Other controls were considered to be deployed infrequently in general: CSC3 “*Vulnerability assessment*”, CSC18 “*Application security*” and CSC19 “*Incident response*”.

5.3 Interpreting the findings in light of related work

Considering these findings alongside the results of various related work (as described in Section 1) that has explored the questions of control effectiveness and deployment is important to support

drawing reliable conclusions and understanding additional relevant context.

Such *et. al.* [25], for example, presented a set of cost-effectiveness measurements for a range of assurance techniques. Where there is overlap between the control sets explored (which is the case in many of the controls for security-testing of systems, configuration and code, although Such *et. al.* did not consider controls for incident response, operational monitoring or training), we can observe similarities between these results and ours – for example, in their finding that vulnerability scans were perceived to be one of the most cost-effective controls. This work also helps to provide additional context for some of the results of our study – e.g., considering whether and how perceptions of the cost of controls, and the time and expertise required to deploy them, might (even subconsciously) have influenced responses to our study. Other work such as [15], which explored the prioritisation of the ISO 27001 controls in a specific organisation scenario using fuzzy Analytic Hierarchy Process, and [9], which used model optimisation to identify cost-effective CIS controls for specific scenarios and threats), can provide similar support and context.

The work of Dietrich *et. al.* [7], who investigated system operators’ perspectives on common security misconfigurations, also provides valuable support for the interpretation of our results on control effectiveness and deployment. Dietrich *et. al.* presented findings on misconfigurations that commonly occur in the deployment of security controls (for example, faulty firewall settings, use of bad or default passwords, and faulty assignment of access privileges) – which provides additional context relating to the deployment of these controls (and indeed, perceptions of such common misconfigurations may have influenced respondents’ perceptions of a control’s overall effectiveness or deployment in our study).

5.4 Relevance of findings to stakeholders

We anticipate that these results will be beneficial to a range of parties from both academia and industry. For researchers seeking to reason about and measure organisational cybersecurity risk, insight into the collective knowledge and perceptions of the security practitioners who observe and address these risks on a day-to-day basis will help to refine understanding of the use of cybersecurity controls and their real value in mitigating risk. This in turn is critical for supporting research in a range of areas including the refinement of cyber-insurance models and the development of approaches to improving organisational cybersecurity posture, and will help to ensure that such research is aligned with the current industry practice.

As we have noted, determining prioritisation of the deployment of risk controls and determining their likely impact on overall cybersecurity posture is a challenging task, and the collective views of the expert community are a valuable resource on this topic. We envisage that our results on overall effectiveness will be useful to organisations making decisions on the implementation of controls, and that the identification of the factors impacting effectiveness and deployment can highlight the considerations that may influence these decisions.

The perspectives of security practitioners on this topic also have implications for the providers of cybersecurity standards and guidance. We might assume that these sources recommend those controls considered most important to maintaining security, and therefore provide an indication of effectiveness. We might also assume that organisations follow key standards, implementing those controls they recommend, and that these sources therefore provide an indication of deployment.

In an examination of six key industry standards (the CIS 20 CSCs [11], NIST Cybersecurity Framework [21], ISO 27002 [1], GCHQ 10 Steps to Cybersecurity [4], Cyber Essentials [5], and COBIT 5 [13]), we found frequent reference to our “most effective” controls (CSCs 3 and 4 are referred to in all, and CSC 19 in five of the six, standards), and all four controls that obtained highest deployment frequency are referenced across all six standards. This implies that these standards are recommending those controls perceived to be most effective by practitioners, and that they play a key role in guiding the adoption of controls by organisations. It is important to note that the recommendations made in these standards are usually updated over time, to reflect perceived changes in the importance of various risk controls and practices.

Finally, we suggest that these results are valuable to the cyber-insurance community. A refined understanding of control effectiveness could inform pre-screening forms, enable more accurate calculation of insurance provisions for an organisation based on the residual risk they face having deployed controls, and inform guidance given by the insurance industry on risk-control products [19]. In calculations of residual risk for organisations in which data on the deployment of controls is lacking, assumptions might be made about the likelihood that controls are in place based on results such as these on control-deployment frequency. Risk modelling must take into account the identified influencing factors, for example the implementation of a control and the organisational context. Understanding which controls organisations deploy, and the ways in which they deploy them, allows us to understand the likely effectiveness of these controls more accurately (and this argument is bolstered by the emphasis placed on control implementation as a factor affecting control effectiveness by participants in the studies).

A comparison of the controls most frequently referenced in 24 insurance forms offered by insurers based in the UK and US [26, 28] with our results showed that our “most effective” controls were referenced relatively little (CSC3 especially). Of the controls that we found to be deployed most frequently, CSC8 and CSC12 are referenced across a large number of these forms, while CSC7 and CSC11 are not. It is important to note that the insurance forms studied are from 2017. Cyber-insurance has of course evolved since, and positive changes have been made; for example, supplemental questionnaires focused on specific topics like ransomware are now quite common. It would nonetheless be valuable to further explore potential discrepancies between data collection by the insurance community, key standards and guidelines, and the perceptions of the practitioner community, and this may highlight key controls that are currently not adequately considered.

5.5 Limitations

It is important that the limitations of the methodology are taken into account in the interpretation of the results. Owing to the nature of our qualitative data collection there was variation in the level of detail in which different participants discussed each question. Furthermore, this paper can report only those factors that emerged through this data collection; it is possible that others would emerge in surveying and conversing with other participants. While our results represent the views of a range of experts in the field, therefore, they are not necessarily exhaustive, and further studies would be valuable in verifying and consolidating these findings. Nevertheless, this is only the first attempt in using this methodology to assess security controls effectiveness and deployment.

6 CONCLUSION AND FUTURE WORK

The effectiveness and deployment of controls varies across different organisational contexts, affected by the characteristics and priorities of organisations, the way in which they configure controls and the threats they face, amongst other factors. We propose a new methodology to explore practitioners' perspectives on the effectiveness of controls in reducing organisational risk (*RQ1*) and on the deployment of these controls by organisations (*RQ2*). We conduct the first attempt to use this methodology with a set of research studies.

Our results showed that certain controls (including vulnerability remediation, administrative-privilege control, and incident-response capability) were considered more effective than others, and that there is variation in the frequency with which controls are deployed (with malware defences and boundary defences being deployed often by organisations). Some controls are reliant on others, and interdependencies between controls impact on their effectiveness but are not necessarily considered systematically during deployment.

Since cyber-risk and cybersecurity practice are fast-changing, findings on control effectiveness and adoption may vary over time. If methodologies such as these are used to collect data (that may influence practice and guidance), it is important that the lifetime for which the data remains relevant is considered. It would be valuable to repeat follow-up studies at regular time intervals, to capture changes in perceptions over time.

In order to understand the validity of both commonly held and differing beliefs, it is important to further comprehend their basis. Therefore, we would recommend that future research seeks to understand the preconceptions that underpin these views, including the possible influence of widely-held misconceptions and past experience. Moreover, our research has showed that practitioners are aware of information that is not reflected in the standards. The community may, therefore, be benefitting from shared and tacit knowledge that helps to improve cybersecurity, but is not formalised or documented fully. Codifying this knowledge will not only help promote best practice, but could also lead to more accurate quantification of risk exposure.

ACKNOWLEDGMENTS

This research was sponsored by AXIS Insurance Company, whose support is gratefully acknowledged.

7 APPENDICES

A INTERVIEW QUESTIONS

Security controls

- (1) How would you make decisions on which security controls to apply to assets?
- (2) How would you make decisions on how to apply these security controls? I.e., their placement and configuration.
- (3) How would you decide what security standards/framework to adopt in an organisation?
- (4) To what extent are connected assets and aggregate (secondary, tertiary) harms considered in control selection?
- (5) What would you say are the three most important security controls today? Why? What do they protect against?

Security control effectiveness

- (1) What would you consider the “effectiveness” of security controls to mean?
- (2) How would you determine the effectiveness of security controls? What factors contribute to a control’s effectiveness?
- (3) What are the types of ways in which organizations can measure a control’s effectiveness? What metrics do you consider relevant?
- (4) What are some of the types of data that would need to be collected to measure control effectiveness?
- (5) Residual risk is the risk remaining after a risk mitigation measure has been applied – do you agree with this description?
- (6) Do you seek to measure or determine the levels of residual risk? If so, how? What metrics do you consider relevant?

Control dependencies

- (1) How do you regard and treat risk controls as dependent upon each other?
- (2) To what extent do organisations have a good understanding about the interdependencies between risk controls?
- (3) How would you go about identifying and understanding the interdependencies between risk controls?
- (4) Are there any interdependencies that you consider risky and monitor accordingly?
- (5) How does the effectiveness of controls impact the interdependencies between controls?

REFERENCES

- [1] [n.d.]. ISO/IEC 27002 Code of practice for information security controls. <https://www.iso27001security.com/html/27002.html> [accessed on 07/02/2020].
- [2] Mohiuddin Ahmed and Ehab Al-Shaer. 2019. Measures and metrics for the enforcement of critical security controls: a case study of boundary defense. In *Proceedings of the 6th Annual Symposium on Hot Topics in the Science of Security*. 1–3.
- [3] Nadher Al-Safwani, Yousef Fazea, and Huda Ibrahim. 2018. ISCP: In-depth model for selecting critical security controls. *Computers & Security* 77 (2018), 565–577.
- [4] National Cyber Security Centre. [n.d.]. 10 Steps to Cyber Security. <https://www.ncsc.gov.uk/collection/10-steps-to-cyber-security> [accessed on 07/02/2020].
- [5] National Cyber Security Centre. [n.d.]. Cyber Essentials. <https://www.cyberessentials.ncsc.gov.uk/> [accessed on 07/02/2020].
- [6] Cylon. 2019. Signal from noise: how to win customers and influence cisos. <https://blog.cylonlab.com/signal-from-noise-how-to-win-customers-and-influence-cisos> [accessed on 07/02/2020].
- [7] Constanze Dietrich, Katharina Krombholz, Kevin Borgolte, and Tobias Fiebig. 2018. Investigating System Operators’ Perspective on Security Misconfigurations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1272–1289.
- [8] Ashutosh Dutta and Ehab Al-Shaer. 2019. Cyber defense matrix: a new model for optimal composition of cybersecurity controls to construct resilient risk mitigation. In *Proceedings of the 6th Annual Symposium on Hot Topics in the Science of Security*. 1–2.
- [9] Ashutosh Dutta and Ehab Al-Shaer. 2019. “What,” “Where,” and “Why” Cybersecurity Controls to Enforce for Optimal Risk Mitigation. In *2019 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 160–168.
- [10] Center for Internet Security. [n.d.]. Cybersecurity Best Practices. <https://www.cisecurity.org/cybersecurity-best-practices/> [accessed on 01/12/2020].
- [11] SANS/Center for Internet Security. [n.d.]. 20 Critical security controls. <https://www.cisecurity.org/controls/> [accessed on 07/02/2020].
- [12] Julie M Haney, Simson L Garfinkel, and Mary F Theofanos. 2017. Organizational practices in cryptographic development and testing. In *2017 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 1–9.
- [13] ISACA. [n.d.]. COBIT 5. <https://www.isaca.org/cobit/> [accessed on 07/02/2020].
- [14] Susan Jamieson et al. 2004. Likert scales: how to (ab) use them. *Medical education* 38, 12 (2004), 1217–1218.
- [15] Hamid Khajouei, Mehdi Kazemi, and Seyed Hamed Moosavirad. 2017. Ranking information security controls by using fuzzy analytic hierarchy process. *Information Systems and E-Business Management* 15, 1 (2017), 1–19.
- [16] Nigel King. 1998. Template analysis. *Qualitative Methods and Analysis in Organizational Research: A Practical Guide* (1998).
- [17] Philip Kobezak, Randy Marchany, David Raymond, and Joseph Tront. 2018. Host Inventory Controls and Systems Survey: Evaluating the CIS Critical Security Control One in Higher Education Networks. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.
- [18] Sangjae Lee, Seongil Jeon, and ByungWon Lee. 2019. Security Controls for Employees’ Satisfaction: Perspective of Controls Framework. *SAGE Open* 9, 2 (2019), 2158244019853908.
- [19] Marsh. [n.d.]. 2019 Cyber Catalyst Designations: 17 Cybersecurity Solutions Chosen. <https://www.marsh.com/us/services/cyber-risk/cyber-catalyst.html> [accessed on 07/02/2020].
- [20] Jason RC Nurse, Louise Axon, Arnau Erola, Ioannis Agrafiotis, Michael Goldsmith, and Sadie Creese. 2020. The Data that Drives Cyber Insurance: A Study into the Underwriting and Claims Processes. In *International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*. IEEE.
- [21] National Institute of Standards and Technology. [n.d.]. Cybersecurity Framework. <https://www.nist.gov/cyberframework> [accessed on 07/02/2020].
- [22] Dimensional Research. 2016. *Trends in Security Framework Adoption: A Survey of IT and Security Professionals*. Technical Report.
- [23] Judy Robertson. 2012. Likert-type scales, statistical methods, and effect sizes. *Commun. ACM* 55, 5 (2012), 6–7.
- [24] Jose M. Such, Antonios Gouglidis, William Knowles, Gaurav Misra, and Awais Rashid. 2016. Information assurance techniques: Perceived cost effectiveness. *Computers & Security* 60 (2016), 117–133. <https://doi.org/10.1016/j.cose.2016.03.009>
- [25] Jose M Such, John Vidler, Timothy Seabrook, and Awais Rashid. 2015. Cyber security controls effectiveness: a qualitative assessment of cyber essentials. Lancaster University.
- [26] Daniel Woods, Ioannis Agrafiotis, Jason RC Nurse, and Sadie Creese. 2017. Mapping the coverage of security controls in cyber insurance proposal forms. *Journal of Internet Services and Applications* 8, 1 (2017), 8.
- [27] Daniel Woods and Tyler Moore. 2019. Does insurance have a future in governing cybersecurity? *IEEE Security and Privacy Magazine* (2019).
- [28] Daniel Woods and Andrew Simpson. 2017. Policy measures and cyber insurance: a framework. *Journal of Cyber Policy* 2, 2 (2017), 209–226.