



Cybersecurity Culture Guidelines: Behavioural Aspects of Cybersecurity

DECEMBER 2018



About ENISA

The European Union Agency for Network and Information Security (ENISA) is a centre of network and information security expertise for the EU, its member states, the private sector and EU citizens. ENISA works with these groups to develop advice and recommendations on good practice in information security. It assists member states in implementing relevant EU legislation and works to improve the resilience of **Europe's critical** information infrastructure and networks. ENISA seeks to enhance existing expertise in member states by supporting the development of cross-border communities committed to improving network and information security throughout the EU. More information about ENISA and its work can be found at www.enisa.europa.eu.

Contact

For media enquires about this paper, please use press@enisa.europa.eu.

Legal notice

Notice must be taken that this publication represents the views and interpretations of ENISA, unless stated otherwise. This publication should not be construed to be a legal action of ENISA or the ENISA bodies unless adopted pursuant to the Regulation (EU) No 526/2013. This publication does not necessarily represent state-of-the-art and ENISA may update it from time to time.

Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication.

This publication is intended for information purposes only. It must be accessible free of charge. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication.

Copyright Notice

© European Union Agency for Network and Information Security (ENISA), 2018
Reproduction is authorised provided the source is acknowledged.

ISBN: 978-92-9204-267-7

DOI: 10.2824/324042

Table of Contents

Executive Summary	4
1. Introduction	5
1.1 Scope of the review	5
1.2 Behavioural sciences and cybersecurity - setting the context	5
2. Summary of evidence reviews	8
2.1 Evidence review of survey studies using social science constructs	8
2.2 Evidence review of models of cyber-security attitudes and behaviour	10
2.3 Evidence review of qualitative and mixed-method studies	12
2.4 Evidence review of current practices	15
2.5 Conclusions from evidence reviews	16
3. Guidelines for practitioners	19
3.1 Awareness	19
3.2 Analyse	21
3.3 Plan	24
3.4 Implementation	25
3.5 Evaluation and iteration	25
4. Recommendations for specific groups	26
5. References cited in the main report	30

Executive Summary

There is a growing recognition that technical cyber security measures do not exist in a vacuum, and need to operate in harmony with people. This has led to a plethora of academic research that seeks to address the role of the human in cybersecurity. It is against this backdrop that ENISA has conducted four evidence-based reviews of human aspects of cybersecurity: two based on the use (and effectiveness) of models from social science; one on qualitative studies; and one on current practise within organisations. These reviews are published online as a technical annex.

Across all four reviews, ENISA found a relatively small number of models, none of which were a particularly good fit for understanding, predicting or changing cyber-security behaviour. Many ignored the context in which much cybersecurity behaviour occurs (i.e. the workplace), and the constraints and other demands **on people's time and resources that it causes**. At the same time, there was evidence that models that stressed ways to *enable* appropriate cybersecurity behaviour were more effective and useful than those that sought to use *threat awareness* or *punishment* to urge users towards more secure behaviour.

There was little evidence that there are specific links between types of people (e.g. gender, personality) and security behaviours. However, by systematically approaching and analysing the current cybersecurity stance of the organisation, and carrying out an in-depth analysis of the causes of any problem(s), ENISA proposes that practitioners can take significant steps towards helping employees to act in a more secure way. This may involve skills-based training and support but may also require the restructuring of security **practises and policies to better align with people's workplace goals and/or capabilities**.

ENISA proposes a model of awareness, analysis and intervention for organisations to systematically plan and implement changes to address human aspects of cybersecurity.

The role of metrics is discussed, in particular the importance of using multiple measures in order to triangulate findings, and the avoidance of over-reliance on self-report measures and simple behavioural metrics.

Organisations should strive for adherence (active participation) rather than compliance - rapidly emerging threats require employees who are engaged and willing to step up. Organisational leadership has a key role in developing effective and workable security - by helping security specialists to fit security into the business, breaking down silos and leveraging other organisational capabilities (safety, HR, communications) - but not least by setting the tone and leading by example. Measures to improve security behaviour should be an ongoing, iterative process - the human factor in cyber-**security is never 'solved', and there is no simple 'solution', but human skills and knowledge, rather than vulnerabilities, can be made to work in favour of an organisation's defensive cybersecurity**.

The report concludes with recommendations for specific groups such as policy makers, management and organizational leaders, CISO and security specialists, CSIRT / CERT community, software developers and awareness raising managers.

1. Introduction

As recognized in the 2016 ENISA report “Definition of Cybersecurity: Gaps and overlaps in standardization”¹, there is no universally accepted definition of cybersecurity, with descriptions of the term varying across authors on multiple dimensions, including the nature of what is protected, from whom, and whether or not unintentional actions are included. Dictionary definitions tend to overly focus on the protection from external attack – for instance, the Oxford English dictionary defines cybersecurity as:

“The state of being protected against the criminal or unauthorized use of electronic data, or the measures taken to achieve this.”²

Alternatively, ISO/IEC JTC1 defines cybersecurity as “preservation of confidentiality, integrity and availability of information in the Cyberspace”³. The advantage of definitions based around confidentiality, integrity and availability (CIA) is that they keep open the potential role of humans in both cybersecurity risks, and protection from threats. However, CIA based definitions tend to underplay the potential for physical risks as a component part of cybersecurity (e.g. deliberate sabotage of a key component in a power grid), or for the use of cyberspace to conduct open source intelligence gathering (either as an end in itself, or in order to provide support to other actions).

1.1 Scope of the review

The present review does not seek to address the definition of cybersecurity, so will adopt a relatively wide definition as its terms of reference that includes external unauthorized access as well as unintentional (and intentional) end user actions that compromise or support the CIA of both information and systems.

The present report is concerned with human aspects of cybersecurity. Research on human behaviour (often also termed ‘behavioural science’) encompasses a wide range of disciplines, with the sole unifying aspect being that the subject of enquiry is the human actor. It therefore includes not only psychology and sociology, but also ethnography, anthropology, human biology, behavioural economics and any other subject that takes humans as its main focal point. The insight that humans are an integral part of delivering cybersecurity is not new, but only over the past 20 years has there been a significant body of social science research that looks at cybersecurity as a socio-technical problem and develops guidance on how to manage that problem effectively. The socio-technical perspective includes the actions (and decisions) of policy makers and security professionals; systems designers, developers and requirements engineers; and end users. Although a part of the cybersecurity ecosystem, ENISA does not in this report consider the human aspects of attackers and adversaries, since the focus of this report is managing the behaviour within defending organisations.

1.2 Behavioural sciences and cybersecurity - setting the context

In their foundational 1975 paper The Protection of Information in Computer Systems (Jerome Saltzer and Michael Schroeder, 1975) established ten principles for designing security. Three of those principles are

¹ <https://www.enisa.europa.eu/publications/definition-of-cybersecurity>

² <https://en.oxforddictionaries.com/definition/cybersecurity>

³ <https://www.iso.org/standard/44375.html>

rooted in knowledge from behavioural sciences: psychology (the security mechanism must be **'psychologically acceptable' to humans who have to apply it**), **human factors and economics (that each individual user, and the organisation as a whole, should have to deal with as few distinct security mechanisms as possible)** and crime science and economics (that the effort required to beat a security measure should exceed resources of/potential rewards for the attacker). Nearly 100 years before Schroeder and Saltzer, Auguste Kerckhoffs (1899) formulated six principles for operating a secure communication system, and three of those principles focus on human factors (easy to communicate and **remember the keys without requiring written notes, portability, and "it must be easy to use and must neither require stress of mind nor the knowledge of a long series of rules"**). **So those who laid the foundations of cybersecurity were very mindful that security involves humans, and how important it is to make security work for them.**

Yet, for nearly 25 years traditional computer security focused on securing technology - the systems and communication infrastructure that hold data and programs. Humans were treated as components whose behaviour can be specified through security policies, and controlled through security mechanisms and sanctions. That security mechanisms are largely ineffective as a result was shown in 2 studies published in 1999: Whitten & Tygar found that even well-motivated and -trained people could not use email encryption correctly, while Adams & Sasse showed that password policies and mechanisms were routinely bypassed by employees. Since then, further studies have demonstrated and explained the ineffectiveness security warnings (Herley 2009) and security awareness and education (Bada, Sasse & Nurse, 2016). **Cybersecurity has identified the human factor as "the weakest link", and sought to make security more 'usable' and 'acceptable' and looked to behavioural sciences to inform the design of policies and mechanisms.**

Many of the knowledge and methods harnessed to improve security come from psychology, human factors, economics and crime science. More recently, anthropological methods have been used to study individual cybersecurity professionals (such as analysts working in CERTs and SOCs - Sundaramurthy et al. 2014, 2015, 2016) and of cyber incident teams (Chen et al. 2014, Bada et al. 2014).

But arguably, there is a fundamental disconnect between what security professionals seek from behavioural sciences, and the guidance they offer. Cybersecurity professionals often attribute non-compliance to employees who are somehow defective, and are looking for interventions to make them behave as specified by the security policies. Thus, many look to psychology or behavioural economics for ways of motivating humans to take cybersecurity seriously (Protection Motivation Theory and Risk **Communication theories are particular favourites**), or how to **'nudge' people towards compliant behaviour**, using mechanisms such as framing (Caputo et al. 2016) or herding (Das et al. 2013). Another set of efforts seeks to reduce human error (such as ignoring warnings or clicking on links) through application of human factors and usability principles. A telling observation is the majority of these studies have been carried out by researchers from engineering disciplines, rather than social scientists. Social scientists would most likely **caution that 'it's complicated' - that the perspective of the human as 'a component whose behaviour we can specify' and control via a small set of 'levers' is misguided. Employee and consumer behaviour are driven by capabilities and limitations (e.g. 'what type of password can I remember?') and they are autonomous 'principal' agents whose behaviour is driven by their own goals, norms and values.** Pallas (2009) presented a security management framework based on new institutional economics, where principal agents are presented with a security value proposition that makes sense in the context of their **own goal and the organisation's wider goals** - most important of all, productivity goals. Recent research by Heath et al. (2018), Ashenden & Lawrence (2016), Coles-Kemp et al. (2018) has demonstrated how goals, values and norms drive behaviour of non-security experts. These studies show that security is ultimately a social construct and as such needs to be negotiated between the different stakeholders in the ecosystem.

The benefits of a collaborative stance are also suggested by economics: the time and effort individuals and organisations can expend is a limited resource - and economics tells us that we ignore such constraints at peril to our security goals - instead, we must apply economic principles to manage those resources effectively (Pallas 2009).

2. Summary of evidence reviews

2.1 Evidence review of survey studies using social science constructs

ENISA incorporated the results from a review carried out earlier this year with funding from the UK National Cyber Security Centre (NCSC)⁴. The review analysed 688 publications that claimed to use behavioural science constructs - variables that are not directly observable, such as attitudes or personality traits, and are assumed to influence human behaviour in cybersecurity - often towards compliance or non-compliance with security policies (see technical annex). Examples of how the results from such surveys in organisations might then be used include:

- The organisation screens prospective employees to identify those who score highly on constructs associated with compliance.
- The organisation screens existing employees to identify those who score highly on constructs associated with non-compliance, and targeting them with security awareness and/or behaviour modification activities.
- The organisation assesses the effectiveness of security awareness and/or behaviour modification activities by how selected or all employees score on such constructs.

Most of these publications claim to use well-established constructs and associated instruments from social sciences, they looked at whether a) the original (validated) constructs and instruments had been used, and b) the studies and information provided met scientific quality standards. The review revealed that there were 92 categories and 984 constructs that have been investigated in relation to cybersecurity behaviour.

Most studies claim to have found a link between constructs and behaviour - and generally assert that some factor (or pattern of factors) within the employees correlates with undesirable security behaviour. For instance, Safa, Von Solms and Furnell (2016) measured responses to the constructs information security knowledge sharing, collaboration, intervention and experience, plus attachment, commitment, personal norms and attitude to information security policy compliance and intention to comply with information security policies with 462 employees in four companies, and conclude that **'the lack of information security awareness, ignorance, negligence, apathy, mischief, and resistance are the root of users' mistakes'** - i.e. attributing undesirable security behaviour to failures by employees. From a scientific point of view, this and conclusions from similar studies is not tenable without some form of validation - triangulation and/or repeated measurements (see Section 2.4).

- It ignores the difference between correlation and causality - the possibility that other underlying factors influence both the constructs as measured by the instrument and the security behaviour - for instance, that daily experience of unworkable security policies shapes the attitude to security, as well as driving non-compliance behaviour. Adams & Sasse (1999), for instance, observed how unworkable password policies had led employees to conclude that this cybersecurity measure was put in place to make their life difficult, rather than offer protection.
- The conclusion assumes that compliance with security policies is sensible - i.e. that the security policies and measures that employees are supposed to comply with are assumed to be good, and that following them improves security - something that Herley (2009) demonstrated is not the case for

⁴ <https://verdi.cs.ucl.ac.uk/constructDB>

many common security measures, and our review of studies grounded in organisational contexts in (Section 2.4) confirms this.

- The 9 constructs used by Safa et al. are among 789 unique constructs identified in the review that have been used to try and explain security behaviour (<https://verdi.cs.ucl.ac.uk/constructDB/constructs/>) and range from personality traits measured through the widely used Big Five (Openness to Experience, Neuroticism, Conscientiousness, Extraversion, Agreeableness) over ethical stances (Utilitarianism vs. Formalism) to high security-specific intention of comply with information security policies. The top **investigated behaviour (60 studies) is Ajzen's (1991) generic Theory of Planned Behaviour (TPB)** - discussed in more detail below. The other top concepts are Compliance (40) and Intention (29). The large number of constructs itself is an indication that there is no agreement in the research community - apart from TPB - on which theories are likely to be applicable. The picture that emerges is one of **security researchers with engineering backgrounds 'grasping' plausible constructs that can be measured and explain non-compliant behaviour.**
- This latter related to the last point - lack reliable results. The review found that most results are not reliable - only a quarter of the studies met basic criteria for scientific survey research. Even where previously validated constructs have been used, the security surveys **often made 'tweaks' to adjust the original, validated instruments** - and then used them without further validation. The conclusion from the review is that most of these surveys are an exercise in trying to find something in employees that can be blamed for their **non-compliant security behaviour, and used by organisations to 'fix' it** (see points 1-3 above). But the results of three quarters of the studies cannot be regarded as reliable - a conclusion that is reinforced by largely divergent results.

The top **investigated behaviour (60 studies) is Ajzen's (1991) generic Theory of Planned Behaviour (TPB)**, which posits that - if people evaluate the suggested behaviour as positive (attitude), and if they think their significant others want them to perform the behaviour (subjective norm), this results in a higher intention (motivations) and they are more likely to perform that behaviour - so far, so rational. TPB then adds the construct of self-efficacy - whether a person believes that she can successfully execute the behaviour required to produce the desired outcomes. This is a concept adapted from Bandura (1977), who stated that self-efficacy(SE) is the most important precondition for behavioural change because it determines the initiation of coping behaviour. This construct was investigated in 11 studies, and a further 6 specially created variants (Computer SE, Knowledge SE, Knowledge sharing SE, Role breadth SE, Security SE, and SE to comply) appear in other studies.

Case study: Installing updates

Installing updates is one of the most obvious security behaviours that serves to protect **both employees and citizens (Reeder, Ion & Consolvo, 2017). However, many users don't** install updates to their operating systems or applications. Many users do not realise updates are important for security (and until recently most companies providing them did not flag that it was for security). Many users reported experiences of updates slowing their systems down, introducing unexpected (and often unwanted) changes to the user interface, or in the worst case it would stop their system working, and automatic updates requiring reboots were experienced as an unwanted disruption for activities and productivity (Wash, Rader, Vaniea, & Rizor, 2014). This is a clear case where educating users about the importance of updates for the security of their device, of their data, and potentially of other home or work network and systems devices, to protect from attacks such as ransomware, would seem to be an easy solution. But

organisations and suppliers can do much to make updates easier and more acceptable - for instance, by distinguishing important security updates from the improvements to non-security elements of the software, by pushing updates during non-working hours or when the user logs off for the day, and by finding ways to identify legitimate updates **from attempts by adversaries to install malevolent software on users' devices.**

This concept also emerges from the next two reviews on influencing security behaviours (sections 2.2 and 2.3). A study by Karlsson et al. (2017), however, found that awareness of a security measure plus intention to comply plus self-efficacy leads to secure behaviour unless the employee finds that such behaviour conflicts with other organisational values: *“employees’ compliance was, to a great extent, a function of the occurrence of conflicts between information security and other organisational imperatives, indirect conflicts with other organisational values”* - such as productivity. Another important insight from Karlsson et al. is that this effect was not detected via standard survey questions used by information security surveys (which ask about security behaviour alone) - but only when put in the context of other organisational goals. Taken together with the results from the qualitative studies (see section 2.3) this highlights that the vast majority of the studies in this review - which focus on security behaviour alone - may produce results that will not apply in real-world circumstances because of their **‘tunnel vision’ on security ignores the other factors driving security behaviour.**

2.2 Evidence review of models of cyber-security attitudes and behaviour

In order to identify the reliable human factors that predict cyber-security behaviours, ENISA further conducted a systematic review of the existing literature connecting attitudes and beliefs to cyber-security behaviours. Searching four different databases of journal and conferences papers yielded an initial 478 articles. After screening for relevance and quality, a total of 47 studies were reviewed, plus one relatively recent review on a related topic (Mayer, Kunz and Volkamer, 2017). The key findings from this activity (see technical annex) were:

1) Reliance on self-report measures

The vast majority of studies relied on self-report measures of cyber-security behaviours (or intention to behave). This was true for both studies of consumer cyber-security behaviour and organisational security policy compliance studies. However, there was some evidence of more creative measures being used – including asking people to create passwords (which were then checked for strength, or tested against the instructions provided to test for compliance); phishing simulations (where simulated phishing attacks are targeted at users); adoption of secure practises (e.g. diary keeping of backups); and the use of passively collected data from users’ operating systems (e.g. acceptance of recent updates; presence of unsigned executable files). This is potentially problematic since self-reporting does not always correlate with actual behaviour (Wash, Rader & Fennel, 2017).

2) Paucity of models

Three models for understanding human aspects of cyber-security dominated the studies reviewed. Both *Protection Motivation Theory* (PMT, see Figure 1) and the *Theory of Planned Behaviour* (TPB) were well represented in the review (see also the first evidence review). In studies of compliance with organisational information security policies, some studies also included general deterrence theory, a model adapted from criminology. There was no evidence that the research literature had moved from these pathway models to consider wider contexts or to integrate insights from behaviour change or persuasive design.

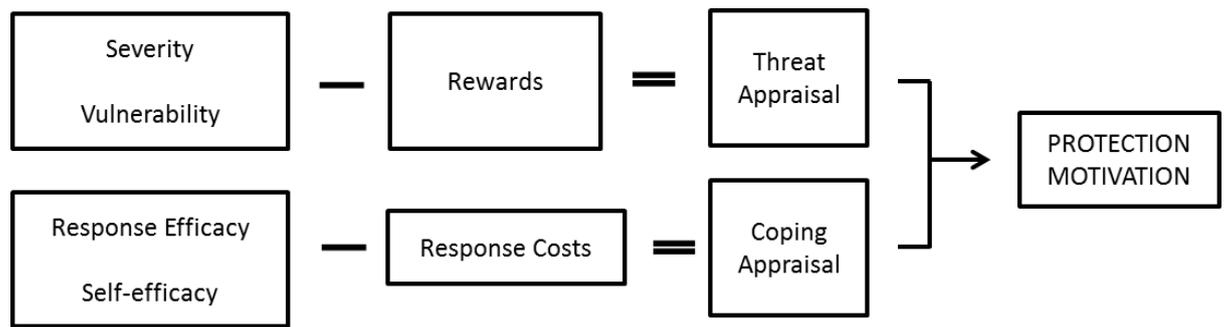


Figure 1: Protection motivation theory (from: https://commons.wikimedia.org/wiki/File:Protection_motivation_theory.png)

3) Threat models have little value predicting behaviour

Within *Protection Motivation Theory*, **people's motivation to protect themselves from potential threats** is determined by the relative balance of the severity and likelihood of the threat and the potential ways to cope with that threat. Studies that have studied the role of threat in predicting this motivation towards protection have reported either weak, neutral or even negative effects of increasing threat appraisal on motivation to take protective action. Similarly, studies of increased severity of punishment for violation of IS policies have found that they can backfire and lead to *less* compliance.

4) Coping models have more value predicting behaviour

On the other hand, studies that have included the resources people have to cope with a threat have produced more promising outcomes. The second element of protection motivation theory is the **individual's appraisal of their likely response to a threat, both in terms of the likely efficacy of the response and their own ability to complete the required response**. These two factors are commonly referred to as '*response efficacy*' and '*self-efficacy*'. In later PMT models, the cost of completing the response was also factored into the model. Within the theory of planned behaviour, '*self-efficacy*' or '*perceived behavioural control*' are used to signify **the users' belief in their ability to complete the desired behaviour**. Across studies of PMT and TPB, coping / self-efficacy was a reliable, moderately strong predictor of cyber-security intention and behaviour. This suggests that interventions that seek to **improve users' ability to respond** appropriately to cyber-threats (and belief that those responses will be effective) is more likely to yield positive results than campaigns based around stressing the threat.

5) Demographics and personality are not particularly useful

Relatively few studies in the review also studied personality or demographics (e.g. age, gender). Those that did found mixed results, with both older and young users often being found to be vulnerable, and gender only sometimes linking to security behaviour or attitudes. Personality rarely linked to security behaviour in a consistent way, although there was some evidence that models of general decision making might be more predictive.

Case study: Phishing guidance.

Most interventions **focus on training users ‘not to click on links’ or to undertake a series** of steps in order to verify the sender or link / attachment. At times, this message is amplified by conducting phishing simulations where employees are sent simulated phishing emails, clicking on which (sometimes) leads to training materials (but often does not).

Advice to ‘think before you click’ risks ignoring the importance of ‘clicking’ in most work environments. If we begin by assuming that clicking is a part of people’s jobs, and that ‘thinking’ before doing so adds a significant burden to people’s completion of their core work tasks, it becomes apparent that we are asking for a significant investment in employee time and energy across an entire organisation.

A more effective approach (e.g. <https://www.ncsc.gov.uk/phishing>) is to take systematic measures that reduce the number of phishing emails that reach users in the first place, and then to provide a reporting and feedback mechanism for employees to check if an email is genuine or not. It may also be useful to inform employees about the most common influence techniques used in phishing emails (e.g. authority, urgency) since this helps them identify multiple forms of malevolent influence attempts.

2.3 Evidence review of qualitative and mixed-method studies

Over the past decade, a few dozen studies have been conducted where the factors influencing behaviour **were first identified in a ‘grounded’ manner from interviews or observation. Rather than apply an existing theory** as to what is driving insecure behaviours, these studies identified and developed an understanding of the causes of non-compliant behaviours; in most of the studies this is the first step of a multi-phase investigation that either then tested hypotheses about causes and effects in a wider population, or made interventions to address those factors and then tested the effect. ENISA reviewed 12 such studies⁵ with general employees in organisations, consumers, and specific professional groups such as developers, CISOs and security analysts.

One of the multi-stage, mixed-methods studies (Ashenden 2015) investigated the links between attitudes and behaviour through attributions. The researcher first elicited attributions through a structured qualitative study in one organisation and tested the emerging constructs with a larger number of employees in the same company. She concluded there two segments of employees that need to be addressed with security awareness programmes and messages – the **‘I Can Handle It Group’** and the **‘It’s Out Of My Control Group’**. This was validated through an intervention study in a second organisation, where employees were profiled in a survey, followed by a targeted intervention (attributional framing was used to tailor persuasive messages to both groups). This construct is remarkably similar to the construct of **‘self-efficacy’**, part of the TPB (see above) which emerged as the most common and most likely framework for successful interventions in the first two reviews above.

From the studies that started by asking individuals to explain non-compliant behaviour, it emerged that the circumstances surrounding the security behaviour are driving it. The most common driver for non-

⁵ <https://www.enisa.europa.eu/publications/cyber-security-culture-in-organisations>

compliance were excessive workload and complexity of the security mechanism, which blocked or disrupted people's primary tasks and activities. Beautement et al. (2008) found that compliance consumes resources (employee time and attention), and when employees fear that the cumulative effect starts to interfere with their personal productivity (which is of course the root of organisational productivity), the tolerance of being compliant is reduced. The *Compliance Budget* is an instinctive response, rather than a consciously and explicitly calculated one. The lesson from this and a series of further studies using the same grounded approach (Kirlappos et al. 2015, Beautement et al. 2016) is that employees fear the consequences of not being productive enough more than they fear the consequences of being the cause of a cybersecurity incident. This judgement is reinforced by organisational leadership being "tacitly complicit" in employees 'productivity first' decisions (Kirlappos et al. 2015). An indicator of this is that non-compliance with security policy is generally bemoaned, but in most organisations not acted upon - until there is an incident. As Pfleeger et al. (2014) point out, safety research learned the importance of acting immediately on any unsafe behaviour, since these latent failures combine and lead to incidents. Safety research also found that treating the "human as hazard" and prescribing only one way to do things does not lead to compliance (passive following of security policies) or adherence (active engagement - including explaining the rules and reminding others) by employees. In the modern environment, organisations should aim for adherence rather than compliance - 1) because the old 'comply or die' does not apply with modern workers who are autonomous agents (Pallas 2009, Kirlappos et al. 2015), and 2) because organisations need employees who are empowered and can step up to counter rapidly evolving threats. Pfleeger et al. also note that cybersecurity needs what safety researcher Reason (2008) described as acts of 'heroism' - dealing successfully with novel threats that overcome the organisation's defences. Security heroism can only be performed by employees who are engaged (adherence rather than compliance) and skilled, trusted and supported by the organisation. Beris et al. (2016) found that how employees feel towards the organisation (positive or negative) is at least as important a driver for adherence as awareness of the risks posed by cyber threats.

Lack of skills and effective support currently affect even stakeholders who are seen as 'very technical' by most - software developers. Research with developers found that security primitives (such as crypto APIs) are difficult to implement (Fahl et al. 2013), and that the demand for productivity means developers cannot research and test the code they produce to the level that security specialists think they ought to (Acar et al. 2016). Productivity pressures also lead to widespread code reuse via repositories such as *Stack Overflow* and *GitHub*. The researchers demonstrate that the effective and efficient way to improve security - rather than tell developers they should 'do more security' and sacrifice productivity - is for security specialists to provide them with secure components and examples that are easy to use. There is also a hint that changes in education and professional certification courses would help to integrate security better in the development process. Poller et al. (2017) conducted an anthropological study in a software development organisation of developers reacted to the results of security reviews and audits carried out by security consultants, and found that developers would follow advice given, but were unable to embed 'security as a routine' in subsequent development processes. This shows how hard it is to embed security in established workflows that have become routine. Whilst organisations may struggle to shift current development practice towards including security, shifting it to include usable security is even harder. Caputo et al. (2016) conducted 3 case studies in organisations that claimed to develop usable security, but found that those working on the projects had little understanding of, and respect for usability, and were not measuring it, either. Most interviewees were convinced there is a trade-off between security and usability - and that security comes first. IEEE Security & Privacy Magazine dedicated a special issue to debunking that myth (see the Expert Round Table discussion - Sasse et al. 2016), but it is still pervasive among developer and security experts. The downstream consequences of this are that tools are difficult and cumbersome to use

- which in turn drives users to non-compliance and not valuing security, as the review of studies with employees shows.

For security specialists such as security analysts and Chief Information Security Officers, security is a primary rather than a secondary task and they work in an environment where everyone puts security first. Sundaramurthy et al. (2014, 2105, 2016) carried out a series of anthropological studies with security analysts working in Security Operations Centres (SOCs), Computer Emergency Response Teams (CERTs) / Computer Incident Response Teams (CSIRTs), and found that - somewhat surprisingly to non-security specialists - the work of these analysts is hampered by conflicts within the organisation. Management in security operation centres do not understand the importance of automation to handle routine cases, so analysts can focus on novel ones, and to ensure analysts have time to perform reflection and automation. As a result, analysts are stuck in operation mode, leading to low job satisfaction and burnout in the worst case. Since CERTs and CSIRTs are critical to the successful defence of many private and public sector organisations, and there is a shortage of staff with the required technical skills, it is essential for organisations to address these issues and skill and support staff. Chen et al. (2014) also point out that CSIRTs need to work in teams and multi-teams, and managers need to foster collaboration rather than competition for these structures to work.

CISOs are the key stakeholders in operational security in organisations. Ashenden & Sasse (2013) found that - despite claiming to want to get staff on board with security - CISOs struggle to engage with other members of the organisation, and rely on one-way communication and **'comply-or-die' approach**. This leads to other employees - **including 'highly technical' ones as developers** - to not engage and seek security advice and guidance. Instead they try to tell security practitioners as little as possible, and as late as **possible, for fear of those 'shooting their baby'**. However, **security issues that emerge late in projects are either expensive or impossible to fix**. Ashenden & Lawrence (2016) conducted a series of Action Research Studies to provide security practitioners with communication and negotiation skills, to enable them to be seen as approachable and supportive. A workshop conducted by the UK National cybersecurity Centre (NCSC) with security practitioners found that CISOs themselves feel they are not understood and **supported by business leaders, while business leaders fear being bamboozled by cyber experts who don't understand what their organisation's business needs, and will instead sell them technology they don't need** (<https://www.riscs.org.uk/2017/10/24/communicating-with-the-board-workshop-summary/>).⁶

What is emerging from these studies is the need to improve communication, collaboration and the working relationship between security specialists and other functions in the organisation. There have been a number of examples of creative security engagement techniques (first mentioned by Dunphy et al., 2014) with employees, consumers or citizens. They are encouraged to reflect on security in their environment, the emotions they feel, the constraints they experience, the pressures that they undergo as well as the actions and the tasks that they perform when generating and sharing information. All of this generates insights on what is needed to make security work for a particular group, in the context of their goals and daily activities, and the physical and social environment in which this takes place. The EU Trespass Project (<https://www.trespass-project.eu/>) has developed and pioneered a number of such techniques, model

⁶ In 2018 ENISA has carried out an additional study entitled CSIRT Law Enforcement cooperation: interaction with the judiciary, that seeks to analyse aspects of cooperation across these three communities (i.e. CSIRTs, LE and the judiciary), for the purpose of enhancing response to combating cybercrime. Behavioural and organisational differences across these three communities may have a role to play in the way cybercrime response is channelled; there is space for improvement along the lines of better understanding of organisational practices, awareness raising, training and tools used.

building with *Lego Serious Play* among them. This type of physical modelling bridges the space between the typical diagrams - flow-charts and UML (unified modeling language) diagrams for example - that security practitioners commonly work with, and the everyday practices of the consumers who are affected by security design. Heath et al. (2018) report a successful case study where this method was used to model security for a home banking application, which identified areas where human intervention and support needed to be provided to make security work overall.

These studies provide examples of different ways of engaging with employees, consumers and citizens on security. This type of engagement and negotiated security solutions is part of a growing trend. Coles-Kemp et al. (2018) conducted a study with citizens in a community in the UK and concluded **“Trust and collaboration ... are necessary for effective cybersecurity.”**

2.4 Evidence review of current practices

Formal evaluation and feedback mechanisms are critical components of any security awareness, training, and education program. But since ENISA (2007) concluded **“Organisations appear to find it very difficult to put effective quantitative metrics in place”** and **“There is little consensus on the most effective measures”** in their report about *Current practice and the measurement of success*, not much has changed in practice.

Still most companies consider collecting metrics in a constantly changing risk environment as challenging, especially given the lack of **universally accepted measurements. And many organizations don’t collect metrics at all.**

Organizations derive those security metrics usually from statistical numbers, performance metrics, tests/inspections or audit results, which can be categorized into the following three general types of security metrics defined by NIST (see National Institute of Standards and Technology, 2008):

- Implementation measures to measure execution of security policy (e.g. compliance with ISO/IEC 27001 or regulations);
- Effectiveness/efficiency measures to measure results of security services delivery (e.g. costs of single activities or whole programme, user satisfaction, change in risk exposure); and
- Impact measures to measure business or mission consequences of security events (e.g. costs of security incidents, cybersecurity budget vs. IT budget).

While those measures increase accountability and effectiveness, and demonstrate compliance of security controls, they do not provide good enough insights into organisational behaviour and the strategies to influence it (see Table 1).

SOURCE	EXAMPLE	ISSUE
Statistical numbers	No. of IT Service Desk tickets related with security	Statistical numbers are often hard to interpret. In the given example, an increase in tickets related with security could mean either that security awareness has dropped, and users behave more insecure, or that security awareness has increased, and users detect and report more incidents

Performance metrics	No. of staff trained No. of visits of Intranet security page	Performance numbers often look good at first sight, but do not help to understand the organisational performance in a way that informs future strategies. Such metrics are called vanity metrics.
Tests / Simulations	Phishing tests Cyber defence simulations Red team vs. blue team	Tests and simulation can give valuable insights into human behaviour patterns. But they are very limited to specific situations and do not provide information about strategies of how to influence the behaviour.
Audit results	ISO/IEC 27001 PCI/DSS	While audit results are the most complete metrics, today's standards and best practice catalogues do not cover the full spectrum of social and psychological items that influence human behavior.

Table 1 Source of metrics to measure cybersecurity awareness and their issues

The maturity of an organization’s cybersecurity awareness programme determines the type of measures that can be gathered successfully. But even the organizations with the most mature programme, rarely collect data which are based on behavioural theories and which allow to draw real conclusions on the behaviour.

It is therefore evident, that the industry must shift from the technology and process centric view to a human centric view and adopt the knowledge from behavioural theories to be successful in the digital age of cybersecurity.

2.5 Conclusions from evidence reviews

The four evidence reviews draw similar conclusions despite drawing on different methods and approaches. One is the importance of metrics in any assessment of the human factor in cybersecurity. While **quantifiable, behavioural measures are often seen as the ‘gold standard’ for the measurement and** evaluation of human aspects of cyber-security, it is important to recognise that deeper insights can be gained by combining these with other approaches, including those drawn from qualitative research and organisational studies. Most metrics in current practice are not suited to measure human behaviour or to provide information about strategies of how to influence the behaviour. It is required that the industry shifts from a technology and process centric view to a human centric view and adopt the knowledge from behavioural theories to be successful in the digital age of cybersecurity.

A second conclusion is that many of the models currently used to study human aspects of cybersecurity are a poor or moderate fit to actual behaviour. The first evidence review found that a model of attitude-behaviour (the theory of planned behaviour) was the most commonly used theoretical model to study cybersecurity. The second review found a combination of the same model and protection motivation theory were the most widely used to study changes in behaviour. Neither are an ideal fit. In the case of the theory of planned behaviour, it ignores wider contextual factors (e.g. organisational factors), and the models tend to assume that compliance (e.g. to a security policy) is a positive outcome. While attitudes, **social norms and perceived behavioural control do predict someone’s intention to undertake a particular**

cyber-behaviour behaviour, most of this work is based around compliance to security policies, that may not be the most appropriate or useful outcome measure.

Third, there is increasing evidence that increasing users' understanding of the threat posed by cybersecurity breaches, or fear of the consequences, is not an effective tool for changing behaviour. There are a number of possible reasons for this. One could be that we have reached 'peak threat appraisal', where the message that cybersecurity threats are pervasive and dangerous is well understood by a large proportion of the population. Another is that increasing fear without providing the tools or skills to address the threat may lead to a sense of helplessness and apathy in the face of impending doom. Alternatively, it may lead to a lack of trust in the message if the threat does not come to pass, or may lead to psychological defence mechanisms such as avoidance. **The lack of effectiveness of fear messaging shouldn't be a surprise** - almost 20 years ago a review concluded that fear appeals are not effective in changing behaviour (Witte & Allen, 2000), a finding replicated more recently in a study of climate change (Feinberg & Willer, 2011), and work on password strength (Weirich, 2006).

Finally, however, the reviews do conclude that there is moderately reliable link between people's ability to cope in the face of threats and their cybersecurity behaviour. Coping can be divided into the *effectiveness of the response* (i.e. does it work) and *ability to carry out the response* (i.e. can I make it work). In both cases, increasing users' coping skills and beliefs reliably leads to changes in their cybersecurity behaviours. **The latter, i.e. "can I make it work" increases the motivation of the individual actors within an organisation and it turns them into positive agents that want to contribute to the solution of the problem threatening the organisation they work for and they have reasons to be loyal to.**

Case study: Public Wi-Fi networks

When on travel, standard advice is to not use public Wi-Fi networks for sensitive work. However many studies have shown that this advice is being ignored.. Recent studies in the UK and Japan have shown that this is despite most people being very aware of the risks. Clearly, scaring them is not working. Many people use their mobile data plan, but when they have less than 50% left, they switch to public Wi-Fi instead (Sombatrung et al. 2016, 2018).

Considering what the research about the *Protection Motivation Theory* tells us, this behaviour must be expected: the severity of the risk of communication interception is very abstract and mostly not understood, and the likelihood of vulnerability rated as very low by the users, while the rewards of staying unsafe are high (comfort, performance and bandwidth, no extra costs).

The more effective approach is to increase the coping appraisal. Response costs should be reduced by giving the users tools that enable them to use public Wi-Fi securely, i.e. an easy-to-install, easy-to-use, secure VPN product on devices they use on the move. And response and self-efficacy should be increased by training them to use it all the time, so they acquire a secure habit. Additionally, it is important that the VPN solution is always on by default, this way the user does not has to switch between two scenarios (not sensitive vs. sensitive work). Having two ways of accessing public Wi-Fi increases **complexity, plus the "secure way" will not become a habit** - and if it is not a habit, people will struggle, make mistakes or give up, especially when under pressure.

This does not mean that organisations should simply invest in more training and skills as such. Trust and collaboration are the foundation of successful cybersecurity - to protect themselves organisations need staff who are loyal and engaged, especially given the rapidly evolving nature of threats. Blaming people or **their inability to comply with security policies (e.g. by referring to them the 'Weakest Link')** is counterproductive. Security must support and work with staff and business leaders to develop workable and effective security solutions, not fight them; intra-organisational efforts need to be targeted against fighting inertia and complacency to counter threats. We should **stop trying to 'fix the human' and fix the security** instead. Insecure behaviour is largely driven by security being too complex and/or effortful. Security needs to accept that human effort and attention is a precious resource primarily dedicated to productivity. Thus, security needs to fit into work processes and tasks rather than disrupt them; policies and tools need to be targeted and easy to follow. If it were possible to make coping with cybersecurity threats a simple, obvious action, we massively increase the likelihood that people will behave in a secure way and thus reduce the likelihood of threats ever materialising. However, making coping with cybersecurity threats an obvious action is a non-deterministic goal and targeted action needs to be taken for significant change to be observed.

3. Guidelines for practitioners

While the reviews that have informed this report identify specific recommendations for security practitioners, ENISA proposes that they be implemented as part of circular, ongoing process (Figure 2, below) that seeks to maintain awareness of an organisations’ security stance and the human aspects that contribute to that stance, followed by in-depth analysis of *why* vulnerabilities might exist (and re-visiting previous efforts to address them). This is followed by a strategic planning stage where available options are weighed up and interventions designed. Critical at this stage is to identify what the *goal* of any change is, and the success (or failure) will be measured. This naturally leads to a final part of the cycle where the effort is evaluated against the original goals, target end state, and the awareness process begins again.



Figure 2: Framework for designing interventions for human aspects of cyber-security

3.1 Awareness

The starting point for any organisation is to gain understanding of its current cybersecurity status, and the ways in which human factors might support or detract from that defensive stance. Our review identified a variety of statistical measures an organisation might use to gain awareness, as well as ways in which groups can gain an understanding of security culture through quantitative (e.g. surveys) and qualitative (e.g. interviews) methods.

The specific methodology adopted will in part depend on the organisation, and its maturity in terms of both cyber-security and data collection and analysis. The results of our reviews suggest that multiple methods are preferable to a ‘one-shot’ approach, and that reliance on easily collected metrics (e.g. phishing simulation click rates, IT support tickets) and self-report surveys are unlikely to reveal more ‘hidden’ human aspects of cyber-security. In the review of qualitative research, ENISA identified a number of studies that provided good blueprints for ways to begin investigation using multiple, open methods that

probe multiple levels in an organisation (e.g. Beutement et al. 2016, Ashenden & Lawrence 2016 which also includes interventions)

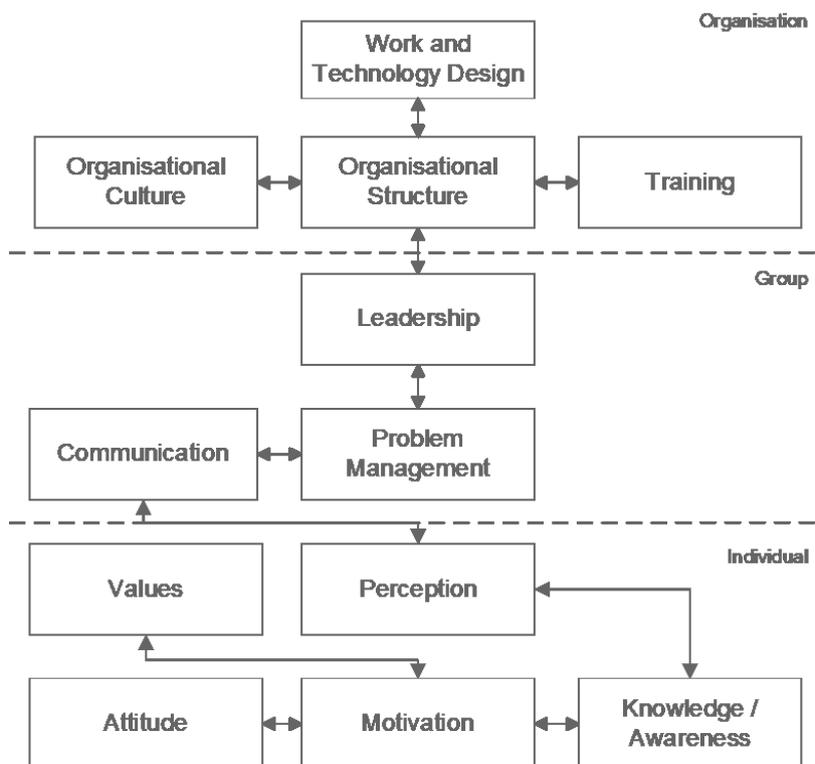


Figure 3: Organisational behaviour model to assess cybersecurity culture (arrows are only exemplary to demonstrate interconnections between domains)

The model of organizational behaviour in Figure 3 is used to study the security culture within organisations (Schlienger, 2006). It has been applied and validated in several studies over the last 15 years, and can be used to guide how an open approach to gaining awareness might look at different levels within an organisation. It is mainly used with quantitative surveys, but since it is an open framework, it can be used in qualitative approaches as well.

Since the framework is based on a model of organizational behaviour, the measurement results indicate concrete starting points for improving and changing the cybersecurity culture. The experience with this model has shown, that the main pain points most often are not in awareness and training, but in supporting domains that strongly influence the work environment of the users, like work and technology design, organisational structure, leadership and problem management. These domains strongly influence the perception and the motivation, which in fact finally have a strong impact on the behaviour.

When developing awareness, it is also important to recognise that many ‘problems’ that seem ‘human’ in cause may be due to a variety of causes, some of which are not within the control of the individual. For instance, some security practices may effectively ‘force’ people to behave insecurely by requiring unrealistic steps or cognitive skills in order to comply. Awareness should therefore include the organisation, work processes as well as local factors (see Figure 1).

Caveats:

- Behaviour change will only happen if the target (secure) behaviour is achievable in the context of an individual's everyday activities.
- Behaviour change takes time and is unlikely to be achieved through a single campaign or intervention. This limitation however, can be balanced with a dedicated and loyal workforce that accumulates secure behaviour messages to the point that response to a threat that brings about positive outcomes becomes a natural reaction.
- **Trying to 'fix the human' without fixing the system won't work alone.**

3.2 Analyse

Once a basic awareness of the organisation's current status is achieved (using whatever method is most appropriate), the next stage is to analyse what may be the root causes of any identified weaknesses or problems. Analysis can be divided into two core elements: analysis of the problem (and root causes), and choice of appropriate method to study the problem (and to measure success). Again, a multi-method approach is particularly suitable here, since some causes may not be readily identifiable using self-report surveys or statistical data. In some cases, the easily measurable (e.g. click rates on phishing simulations) might not be the most suitable measure (e.g. of security culture). Some suitable methods might be to use workgroups or focus groups to identify whether a behaviour is not being conducted due to ability, motivation or another factor before designing any intervention. While surveys can be valuable (see technical annex), organisations should also consider that collection of data via surveys consumes a precious resource – employee time and effort, and many organisations report that employees suffer from survey fatigue. Care needs to be taken to not over survey staff, or to rely solely on self-report survey data when possible.

Selecting the right metrics

It is important to select the right instrument for any analysis or evaluation. The instrument selected should give clear and relevant information. And the organisation should also choose the right number of metrics, not too few, but also not too many, so that they give valuable insight but still are manageable.

Choosing a metric

When choosing or designing a measurement instrument, it is advised to follow the common SMART criteria. SMART stands for Specific, Measurable, Actionable, Relevant and Time-related.

When looking at a proposed metric, make sure that it is:

Specific: Does it target a specific area for improvement?

Measurable: Is it quantifiable or does it at least suggest an indicator of progress?

Actionable: Can the results be used to define concrete improvement actions?

Relevant: Is it relevant for your organisation taking your context into consideration and does everybody understand the result?

Unpicking the causes of a behaviour

Practitioners should consider conducting a more detailed analysis of the causes or barriers to conducting the desired behaviour. In our review of intervention models, ENISA identified two approaches to unpicking the causes of (non) behaviour that are particularly suitable for cybersecurity: **COM-B** and **Fogg’s behaviour** model.

The ‘**COM-B**’ model (Michie et al., 2011; Figure 4 below) argues that whether or not a behaviour is enacted (e.g. locking a screen when leaving for lunch) is dependent upon three interrelated factors: 1) capability (can they do it? Do they know how to?); 2) opportunity (do they have the chance to do the action?); and 3) motivation (are they motivated to lock the screen?). The type of intervention is dependent upon the cause of the (non)behaviour - so for instance, if users are able to lock a screen, have the opportunity to do it, but are not motivated to, then interventions should be based around creating a motivation (e.g. education, awareness, reward/punishment). If initial analysis found that users were motivated, but did not know how to lock screens (capability), then an intervention should be based on training (skills-building). However, this example also shows the importance of identifying the real cause of underlying behaviour and responding with the appropriate intervention: in several of the qualitative studies reviewed (e.g. Kirlappos & Sasse, 2015), employees knew they were supposed to lock screens, and how to do this, but did not lock them because - in their specific work context - they felt it would signal a lack of trust in their work colleagues - and that is the inhibitor to changing to the secure habit.

ENISA has highlighted at several points in this report that trust between the organisation and staff, and between staff is important for cybersecurity, but here an education intervention is required: employees **need to understand that screen locking is a key security habit, and that ‘it’s business, not personal.’** Once employees know their colleagues will interpret the security behaviour as an organisational security habit, rather than personal mis-trust, the blocker is removed. Trust in the real world vs. trust online is an excellent example of where organisations should target education and skills-building (Kirlappos & Sasse, 2012): attacks like phishing and social engineering exploit the fact that most people rely on trust models from the physical world - understanding where trust in the online world is different (you cannot always be sure an email from who they claim to be) and how one responds appropriately is important to understand (e.g. consult security or colleagues, verify identity via different channels, refuse request politely). At the end of the day an employee would want to take credit and be proud of demonstrating a suitable response that “saves the day” for the employing organisation when an attack is deployed.

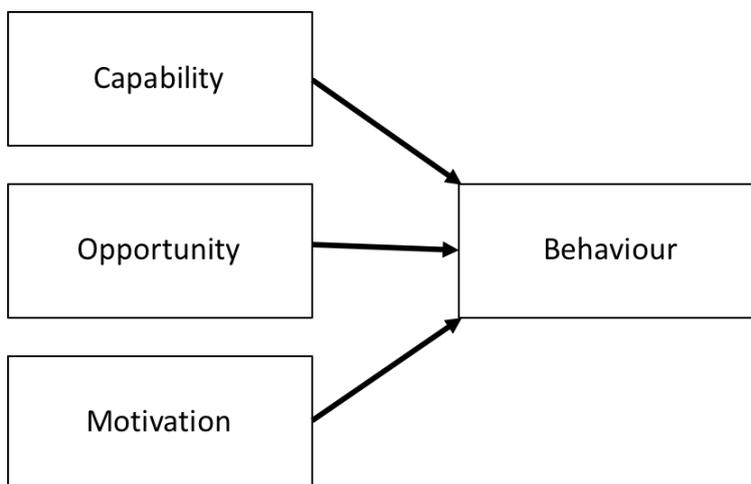


Figure 4: COM-B model (adapted from Michie et al., 2011)

A related model is that developed by Fogg (2009) that seeks to identify the type of cue needed to **encourage the appropriate action, dependent on an individual’s motivation and ability to perform the act** (see Figure 5). According to the B=MAT model, the likelihood of a behaviour occurring is a product of motivation (M), Ability (A), and the appropriate trigger (T).

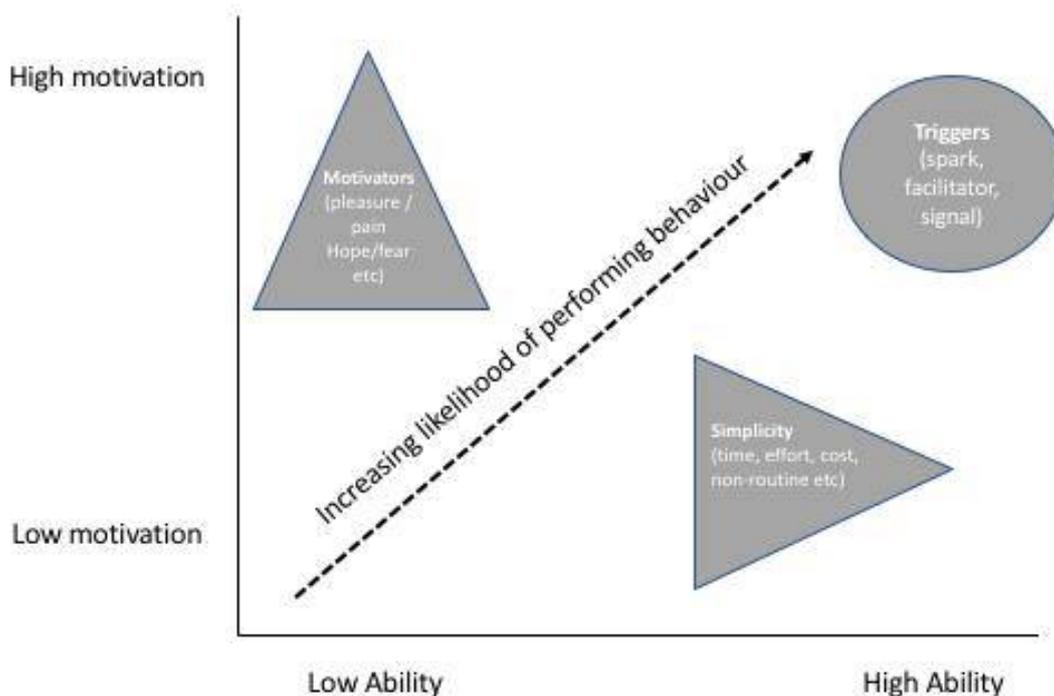


Figure 5: B=MAT model (adapted from Fogg, 2009)

According to the B=MAT model, the type of persuasion required to bring about a behaviour depends on where it lies in the motivation / ability dimensions, with different interventions needed to increase either motivation or ability. For instance, if people are motivated to undertake a task (e.g. updating software), then addressing their ability (e.g. by reducing the cost or effort) should increase the likelihood of carrying out the behaviour. Similarly, if an action is simple and the person is able to complete it, then addressing motivation (e.g. fear of outcome, hopes, pain) should also increase the likelihood. **Once motivation and ability are addressed, according to Fogg’s model, we should then look to triggers that signal to people that a behaviour is required.** These triggers can take the form of: 1) signals (e.g. a message saying that updates are ready to be installed), best used when someone has motivation and ability; 2) sparks that seek to motivate as well as trigger a behaviour (e.g. warning that the **computer will be at risk if the update isn’t installed**); or 3) **facilitators, that seek to both trigger a behaviour and make it easier (e.g. “just click here to download and install the update”)**.

As noted in the reviews that contributed to this paper, interventions designed to improve the ways in which users cope with cybersecurity threats have shown reliable associations with actual behaviour. ENISA

propose then that the COM-B model is used to identify *why* a desired behaviour may or may not be carried, and that **Fogg’s behaviour model guides thinking about possible interventions.**

It is also important to note that in some cases the *cause* of a non-compliance / non-protective action may well sit outside the user or employee. For instance, ENISA has documented studies of **‘shadow security’ where policies interfere excessively with employees’ work goals, and so are circumvented. In these cases, changing the policy is more suitable than changing employee behaviour.**

In our review of measures (see technical annex), ENISA make a series of recommendations around the selection of valid, reliable measures that would allow practitioners to track changes in security behaviour or culture. ENISA advise the use of multiple measures that can be triangulated, and for the collection pre- and post any changes aimed at improving the human aspects of cybersecurity.

3.3 Plan

Once awareness and analysis has been completed, the next stage is to plan based around that. The exact nature of that planning will be determined by the diagnosis of around the problem in the previous two stages. For instance, if the analysis stage identified a problem with understanding, a traditional awareness-raising or training campaign is appropriate. Alternatively, the analysis may identify specific issues around the design of policies and work practises. In this case, it would be appropriate to either consider the **goodness of fit of policies, and redesign then around workplace requirements. Similarly, ‘opportunity’ may be diagnosed due to a lack of usability of a particular security protocol or tool, in which case, rather than training users on a poorly designed or implemented tool, the tool should be redesigned. For instance, if users want to encrypt email, but the tool for doing so is cumbersome and difficult to use, redesigning the tool is more likely to yield success than attempting to redesign the human.**

In the case of training, much current training is actually awareness raising in some form or other. For instance, phishing simulations do little to provide real, implementable skills for employees to use in their everyday work. ENISA therefore proposes **that practitioners consider training to be ‘skills building’, with identifiable outcomes that help employees cope with cybersecurity threats of various kinds. However, if this training is still to support policies and practises that interfere with employees’ core work goals, it will have little impact unless those same policies and practises are also changed to align with work and business processes.**

PROBABLE CAUSE	EXAMPLE ORGANISATIONAL RESPONSES	COMMON ACTIVITIES
Capability	Redesign policies & tools Education Skill-building Restrict	‘Fix security’ - review & change policies & tools Build employee security skills (‘proper’ training) Remove admin rights
Motivation	Awareness campaigns Incentives Organisational response	Security culture programme ‘Good security’ awards, security performance as Key Performance Indicator Visible organisational reaction to all policy breaches & errors

Opportunity	<p>Engage employees in security review/design</p> <p>Security champions</p> <p>Nudge / prompt</p>	<p>Identify policies & tools that cause friction</p> <p>Identify & support employees who want to build security skills</p> <p>Support transition to secure habits through alerts & reminders</p>
Audit results	ISO/IEC 27001 PCI/DSS	While audit results are the most complete metrics, today's standards and best practice catalogues do not cover the full spectrum of social and psychological items that influence human behavior.

Table 2: Example links between analysis and potential organisational responses

3.4 Implementation

Changing behaviour (or indeed organisational culture) is likely to be a long-term project. Some techniques and approaches may yield immediate results. In other cases, it may take considerably longer, not only because habits take time to form, but also because restructuring security policies and practises to align with organisational goals and work tasks is likely to be a considerable undertaking.

Implementation should also be monitored as part of the process - ideally with interim measurement and analysis as the programme is on-going.

3.5 Evaluation and iteration

Evaluation of any intervention can take two main forms - process and outcome. Process evaluation seeks to identify how the attempted change or intervention ran - was it implemented correctly? How did the different stakeholders interact? What process elements could be improved? Outcome evaluation seeks to identify whether or not the change achieved its stated goals. As mentioned earlier, this relates closely to the issues around measurement and metrics. It is critical that practitioners know what success looks like - be able to identify what metrics or measures will change in response to any intervention ahead of time. Ideally, the same measures and metrics are used before and after any change in order to quantify the **effort. If this isn't possible, a control group can be used for comparison purposes (e.g. if the organisation has multiple sites, and one site can be kept outside of the intervention).**

4. Recommendations for specific groups

Policy makers

For policy makers ENISA identified a clear lesson from the reviews - increasing cybersecurity literacy and skills is an evidenced method to support citizens to protect their cybersecurity. There are further ways in which policy makers can support cybersecurity:

- Ensure responsibilities are assigned to those who have the capability to discharge those responsibilities (key principle Risk Management ISO 27001).
- Signpost trustworthy competence - e.g. where to look for guidance on specific issues.
- Promoting collaboration and trust-building within and across organisation; promote respect and ban disrespectful language about stakeholders
- Support the development and use of evidence-based metrics and measures to assess cybersecurity skills, knowledge and organisational culture.
- Encourage and support collaboration between technical and social / behavioural domain experts in tackling cybersecurity behaviours.
- **Don't assume that awareness of a threat will lead to protective action!** Citizens (and workers) also need the skills to avoid the threat, and an understanding that doing so will be effective.

Management and organisational leadership

The outcome of the qualitative and practice reviews is that organisational leaders need to shift their perspective on what their role and responsibilities in managing cybersecurity in their organisation are. Cyber threats are existential to all organisations whose business relies on IT and networked communications, but the management of most organisations has to date seen the problem as a technical one, and handed it to technical specialists (usually the CISO) to manage. That means decisions about security policies and how to implement them and are made from a security perspective, with little to no consideration of the impact on individual employees going about their daily work tasks. The result is that many security policies cause friction with business, and when that happens most employees put **'productivity first'**. Organisational leadership - while happy to keep saying how important security is to the **organisation, are "tacitly complicit" in this behaviour (Kirlappos et al. 2015). Organisational leaders need to take an active role to ensure that the security behaviour staff are asked for is doable in the context of the business.** Organisational leaders should:

- Take responsibility: Decide which security risks they want to manage, and commit the resources required. This includes the cost of buying security equipment and services, but more important the *total cost of operating those security measures*:
 - the time staff have to spend on security
 - the time and effort required to change insecure behaviours into secure ones
 - business that may be lost as a result of staff following the policies
 - building the skills different parts of the organisation need
- Lead by example: always follow security policies (you are a high-value target) and commit a portion of your time to work with security specialists and staff to finding workable solutions in your areas of expertise (e.g. operations, finance, marketing)
- Manage your organisation to help security: The report has found that effective management of security can draw knowledge and tools from safety. Similarly, Human Resources can assist with the

design of incentives and KPIs, and helping to identify staff who are interested in acquiring new skills and take on extra responsibilities. Communications and marketing can help to improve effectiveness of awareness materials. Leveraging your own resources and breaking down silos to solve problems is a management task.

CISO and security specialists

Studies with CISOs show a need for a shift in what the job entails, and how to work to be effective.

- Calculate the impact of your policies. Security policies and measures can only be effective if they are adopted and used correctly. To ensure this, CISOs need to know the impact it has on staff in daily business operations - so they should know the time and effort compliance requires before putting a **security measure in place**. **General MacArthur's dictum 'Never give an order that can't be obeyed'** should be framed and hung over their desk.
- Be visible and approachable. CISOs should be security cheerleaders, not policemen - to engage staff, they need to be visible and approachable at all times. And to build trust, they need to listen and **negotiate, rather than try to 'stamp out' non-compliance** and throttle innovation and experimentation.
- You need soft skills. CISOs and other security specialists **need to acquire the 'soft skills' to do this** effectively. The programme developed by Ashenden & Lawrence (2016) with funding by the UK NCSC is a pioneer effort that should be replicated, and become incorporated into academic and professional security courses.
- Stop verbally bashing people. Effective engagement and trust require respect for others. The security profession needs to revise its perspective of non-specialists, and accept they are the only stakeholders in the ecosystem for whom security is the primary roles. And they must change the language they use to reflect this - **no more referring to humans as 'the problem', 'stupid users' and the 'Weakest Link.'**

CSIRT's /CERT's & SOC's

Incident response teams and security operations centre staff are among the most important assets in the fight against cyber threats. Enabling them to perform well is clearly important - and so is having sufficient capacity to fill positions. Given that there is a growing shortage of cyber skills professionals (Oltzig, 2017), the finding from our review that burnout is a serious problem should give rise for concern. Organisations need to take steps to manage this precious resource effectively.

- Look after your staff. Burnout is largely driven by overload, but also by boredom through repetitive tasks. Organisations need to ensure sufficient staffing levels - this can be a challenge in CSIRTs where demand is high during incidents, but lower during other times. More flexible task allocations - e.g. having staff work on tool development, skills resources and team building and knowledge-sharing between teams - would be a way forward. Managing incidents affecting safety and cyber physical systems can be stressful - organisations need appropriate support to help staff deal with the aftermath.
- Invest in training and personal growth. Having the skills, and ability to grow are key for effective performance, confidence and job satisfaction of security staff. In a rapidly evolving threat environment, funding for research and specialist skills is an essential, not a luxury. Some skill-building can happen through online courses, but hands-on case analysis, master classes and wargaming are more engaging and effective.
- Support team and multi-team approaches. Effective cybersecurity defence is a team sport - building trust among analyst teams, and between them and management is essential. This is a new problem for

many organisations, so some investment in external expertise (research or consultancy) may be required.

Software Developers - and those who manage and educate them

The review of literature on software developers revealed that they - like other staff - are currently caught between producing code and delivering products, and make sure what they deliver is secure. The studies show that there are high-impact first steps that can be taken to make it easy for developers to produce more secure code - by ensuring that platforms, tools and APIs they use have secure defaults and settings, and that code that is frequently copied is vetted and made safe. One major study found that incorporating security in the development process through expert security guidance was not successful, but another showed that enabling security experts to become more approachable and supportive can lead to more - and most critically earlier - engagement between security and staff.

The lessons for developers are:

- You cannot be a security expert, but you need to think about security from the start (security by design) and throughout the whole lifecycle (secure software development lifecycle). The longer you leave thinking about security, the more expensive it will be to make it secure.
- Security experts can advise and support you - work with them, and tell them everything.
- If your code includes a security mechanism, you need to make sure it is usable: how much time will it take? Will they be able to understand the decisions you are asking them to take? Work with usability experts to help you design and test for usable security.

Those who manage and educate developers

One of the studies reviewed found that even organisations who claimed to have processes to ensure their software was secure (and usable) had no criteria or metrics by which staff could determine it was so. Also, problems arising with security, and usability of security, tend to end up with help desks and support staff, rather than the developers. If they did, they would not only have an incentive to reduce those problems, but learn over time how to avoid them in the first place.

- Like other staff, developers need time to keep up-to-date with threats and update their skills.
- Assign developers to work on help desks and support, so they experience first-hand the consequences of failing security.
- Include the number of support calls and other costs associated with insecure or unusable code in performance evaluation

Those who educate developers - computer science and software engineering courses - do currently not teach security-by-design. Some offer specialist modules in security, but this is mostly optional.

- Security courses should be compulsory in academic computer science and engineering courses.
- It should be taught as an integral part and from year one programming should teach students to program securely, rather than just how to program.
- Material in all other modules should be reviewed to remove or annotate examples that would lead to insecure code.

Awareness raising managers

Our reviews clearly noted that awareness based around threat is not effective. And yet, many awareness campaigns still spend considerable time and energy repeating the scale and vulnerability of cybersecurity **threats. Whether we have reached 'peak awareness' may be unknown, but the evidence does support that** we should be aiming to provide users with the skills in order to cope with threats, and the knowledge that a simple act can be effective protection (e.g. accepting updates immediately). While efforts to further understand the attitudes and beliefs of a population might be worthwhile, they should be linked closely to an analysis that leads to tailored campaigns based on identified issues. The COM-B and Fogg models outlined in the report give a structure to begin strategically tailoring awareness campaigns towards specific causes of a (non) behaviour.

5. References cited in the main report

Note: see technical annex for full references and evidence reports

Acar, Y., Backes, M. Fahl, S., Kim, D., Mazurek, M., Stransky, S. (2016): You Get Where You're Looking For - The Impact of Information Sources on Code Security. Proceedings of the 2016 IEEE Symposium on Security and Privacy - Oakland'16.

Adams, A. & Sasse, M. A. (1999): Users are not the enemy: Why users compromise computer security mechanisms and how to take remedial measures. *Communications of the ACM*, 42(12):40-46, December 1999.

Ajzen, I. (1991): The theory of planned behaviour. *Organizational Behavior and Human Decision Processes*. Volume 50, Issue 2, December 1991, Pages 179-211.

Ashenden, A.: Information Security Awareness: Improving Current Research & Practice. PhD thesis, UCL Department of Computer Science. 2015

Ashenden, D. & Lawrence, D. (2016): Security Dialogues: Building Better Relationships between Security and Business. *IEEE Security & Privacy Magazine*, May/June 2016.

Ashenden, D. & Sasse, M. A. (2013): CISOs and organisational culture: Their own worst enemy? *Computers & Security*, Volume 39, Part B, November 2013, Pages 396-405.

Bada, M., Sasse, M. A. & Nurse, J. R. C. (2015): Cyber Security Awareness Campaigns: Why do they fail to change behaviour? International Conference on Cyber Security for Sustainable Society, 118–131.

Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2), 191-215.

Beautement, A., Becker, I., Parkin, S., Krol, K., & Sasse, M. A. (2016). Productive Security: A Scalable Methodology for Analysing Employee Security Behaviours. Procs. SPOUPS 2106 USENIX Association.

Beautement, A. Sasse, M. A., Wonham, M. (2009): The compliance budget: managing security behaviour in organisations. In *New Security Paradigms Workshop (NSPW)*, 2008, pages 47-58.

I. Becker & M. A. Sasse (2018): Separating Security Science from Pseudo-Science:

A Systematisation of Knowledge (SoK) review of survey constructs measuring security behaviour. Report and analysis available at <https://verdi.cs.ucl.ac.uk/constructDB/>

Beris, O., Beautement, A., & Sasse, M. A. (2015). Employee Rule Breakers, Excuse Makers and Security Champions:: Mapping the Risk Perceptions and Emotions That Drive Security Behaviors. Proceedings of the 2015 New Security Paradigms Workshop, 73-84.

Caputo, D.D., Pfleeger, S.L., Sasse, M.A., Ammann, P., Offutt, J., Deng, L. (2016). Barriers to Usable Security? Three Organizational Case Studies. *IEEE Security and Privacy*, 14 (5), 22-32.
doi:10.1109/MSP.2016.95

Chen, T. R., Shore, D. B. Zaccaro, S. J., Dalal, R. S. Tetrick, L. E. & Gorab, A. K. (2016): An organizational psychology perspective to examining computer security incident response teams. *IEEE Security & Privacy*, (5):61–67, 2014.

Coles-Kemp, L., Ashenden, D., O'Hara, K (2018): **Why Should I? Cybersecurity, the Security of the State and the Insecurity of the Citizen.** *Politics and Governance* (ISSN: 2183–2463) 2018, Volume 6, Issue 2, Pages 41–48.

Dunphy, P., Vines, J., Coles-Kemp, L. Clarke, R., Vlachokyriakos, V. Wright, P., McCarthy, J. & Olivier, P. **“Understanding the Experience-Centeredness of Privacy and Security Technologies,”** in *Proceedings of the 2014 Workshop on New Security Paradigms Workshop NSPW 2014*, pp. 83–94

ENISA (2007). Information security awareness initiatives: Current practice and the measurement of success. In ENISA, *Raising Information Security Awareness across Europe*, Office for Official Publications of the European Communities (ISBN 978-92-9204-002-4).

Fahl, A., Harbach, M. Perl, H. Kötter, M., Smith, M. (2013): Rethinking SSL Development in an Appified World. *Proceedings of the 2013 ACM conference on Computer and Communications Security - CCS'13*.

Feinberg, M., & Willer, R. (2011). Apocalypse soon? Dire messages reduce belief in global warming by contradicting just-world beliefs. *Psychological science*, 22(1), 34-38.

Fogg, B. J. (2009, April). A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology* (p. 40). ACM.

Heath, C., Hall, P. & Coles-Kemp, L. (2018): Holding on to dissensus: Participatory interactions in security design. *Strategic Design Research Journal*, 11(2): 65-78 May-August 2018.

Herley, C. (2009): So long, and no thanks for the externalities: the rational rejection of security advice by users. *Proceedings of the New Security Paradigms Workshop (NSPW) 2009*. ACM.

Karlsson, F., Karlsson, M., Åström J. (2017): Measuring employees' compliance – the importance of value pluralism. *Emerald Insight* 25 (3) 279 – 299. 133-144.

Kerckhoffs, A. (1883) 'La cryptographie militaire', *Journal des sciences militaires*, vol. IX, pp. 5–38, Jan. 1883, pp. 161–191, Feb. 1883.

Kirlappos, I., Parkin, S., Sasse, M.A. (2015). "Shadow security" as a tool for the learning organization. *ACM SIGCAS Computers and Society*, 45 (1), 29-37. doi:10.1145/2738210.2738216

Mayer, P., Kunz, A., & Volkamer, M. (2017). Reliable Behavioural Factors in the Information Security Context. Paper presented at the *Proceedings of the 12th International Conference on Availability, Reliability and Security*, Reggio Calabria, Italy.

Michie, S., Van Stralen, M. M., & West, R. (2011). The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation science*, 6(1), 42.

National Institute of Standards and Technology, 2008. *Performance Measurement Guide for Information Security*, Gaithersburg, USA: NIST Special Publication 800-55 Revision 1.

Oltsig, J. (2017): Research suggests cybersecurity skills shortage is getting worse. CSO Online Jan 11, 2018 <https://www.csoonline.com/article/3247708/security/research-suggests-cybersecurity-skills-shortage-is-getting-worse.html>.

Pallas, Frank, Information Security Inside Organizations - A Positive Model and Some Normative Arguments Based on New Institutional Economics (August 11, 2009). Available at SSRN: <https://ssrn.com/abstract=1471801> or <http://dx.doi.org/10.2139/ssrn.1471801>

Pfleeger, S. L., Sasse, M. A., & Furnham, A. (2014). From Weakest Link to Security Hero: Transforming Staff Security Behavior. *Journal of Homeland Security and Emergency Management*, 11 (4), 489-510. doi:10.1515/jhsem-2014-0035.

Poller, Andreas, Laura Kocksch, Sven Türpe, Felix Anand Epp, and Katharina Kinder-Kurlanda. "Can Security Become a Routine?: A Study of Organizational Change in an Agile Software Development Group." In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 2489-2503. ACM, 2017.

Reason, J. (2008): *The Human Contribution: Unsafe Acts, Accidents and Heroic Recoveries*. 2nd Edition. Routledge.

Reeder, R., Ion, I. & Consolvo, S. (2017): 152 Simple Steps to Stay Safe Online: Security Advice for Non-Tech-Savvy Users. *IEEE Security & Privacy Magazine*.

Safa N. S., Von Solms, R. & Furnell, S. Information security policy compliance model in organizations. *Computers & Security*, Volume 56, February 2016, Pages 70-82.

Sasse, M.A., Smith, M., Herley, C., Lipford, H., Vaniea, K. (2016). Debunking Security-Usability Tradeoff Myths. *IEEE SECURITY & PRIVACY*, 14 (5), 33-39.

Saltzer J. H., Schroeder, M. D. (1975): The protection of Information in Computer Systems. *Proceedings of the IEEE* Vol. 63, Issue: 9, Sept. 1975.

Schlienger, T., 2006. *Informationssicherheitskultur in Theorie und Praxis*. Doctoral thesis. Fribourg, Switzerland: iimt University Press.

N. Sombatruang, M. A. Sasse, and M. Baddeley, "Why do people use unsecure public wi-fi?: an investigation of behaviour and factors driving decisions," in *Proceedings of the 6th Workshop on Socio-Technical Aspects in Security and Trust*. ACM, 2016, pp. 61–72.

N. Sombatruang, Y. Kadobayashi, M. A. Sasse, M. Baddeley, and D. Miyamoto, "The continued risks of public wi-fi and why users keep using it: Evidence from Japan," In *Press*, 16th Annual Conference on Privacy, Security and Trust. IEEE, 2018.

SC Sundaramurthy, J Case, T Truong, L Zomlot, M Hoffmann (2014):

A tale of three security operation centers. *Proceedings of the 2014 ACM workshop on security information workers*, 43-50.

Sathya Chandran Sundaramurthy, Alexandru G Bardas, Jacob Case, Xinming Ou, Michael Wesch, John McHugh, S Raj Rajagopalan (2015) A Human Capital Model for Mitigating Security Analyst Burnout. *Procs SOUPS 2015* pp. 347-359.

Sathya Chandran Sundaramurthy, Michael Wesch, Xinming Ou, John McHugh, S Raj Rajagopalan, Alexandru Bardas. (2017): Humans are dynamic. Our tools should be too. Innovations from the Anthropological Study of Security Operations Centers. IEEE Internet Computing.

Vania, K. E., Rader, E., & Wash, R. (2014, April). Betrayed by updates: how negative experiences affect future security. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 2671-2674). ACM.

Wash, Rader, Vania, Rizor (2014): Out of the Loop: How Automated Software Updates Cause Unintended Security Consequences. Procs SOUPS 2014

Wash, R., Rader, E., & Fennell, C. (2017, May). Can People Self-Report Security Accurately?: Agreement Between Self-Report and Behavioral Measures. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (pp. 2228-2232). ACM.

Weirich, D. (2006). Persuasive Password Security. PhD Thesis, University College London 2006. http://discovery.ucl.ac.uk/1446157/1/Weirich_thesis.pdf

Whitten, A. & Tygar, D. (1999): Why johnny can't encrypt: A usability evaluation of pgp 5.0. In Proceedings of the 8th USENIX Security Symposium, August 1999.

Witte, K. & Allen, M. (2000). A meta-analysis of fear appeals: Implications for Effective public health programs. Health Education and Behavior 27(5), 591-615.



ENISA

European Union Agency for Network
and Information Security
1 Vasilissis Sofias
Marousi 151 24, Attiki, Greece

Heraklion Office

Science and Technology Park of Crete (ITE)
Vassilika Vouton, 700 13, Heraklion, Greece



TP-04-18-964-EN-N



1 Vasilissis Sofias Str, Maroussi 151 24, Attiki, Greece
Tel: +30 28 14 40 9710
info@enisa.europa.eu
www.enisa.europa.eu

ISBN: 978-92-9204-267-7
DOI: 10.2824/324042

