Cyber Security

Knowledge Transfer Network

Knowledge Transfer Networks
A DTI business support solution
Delivered through the Technology Programme

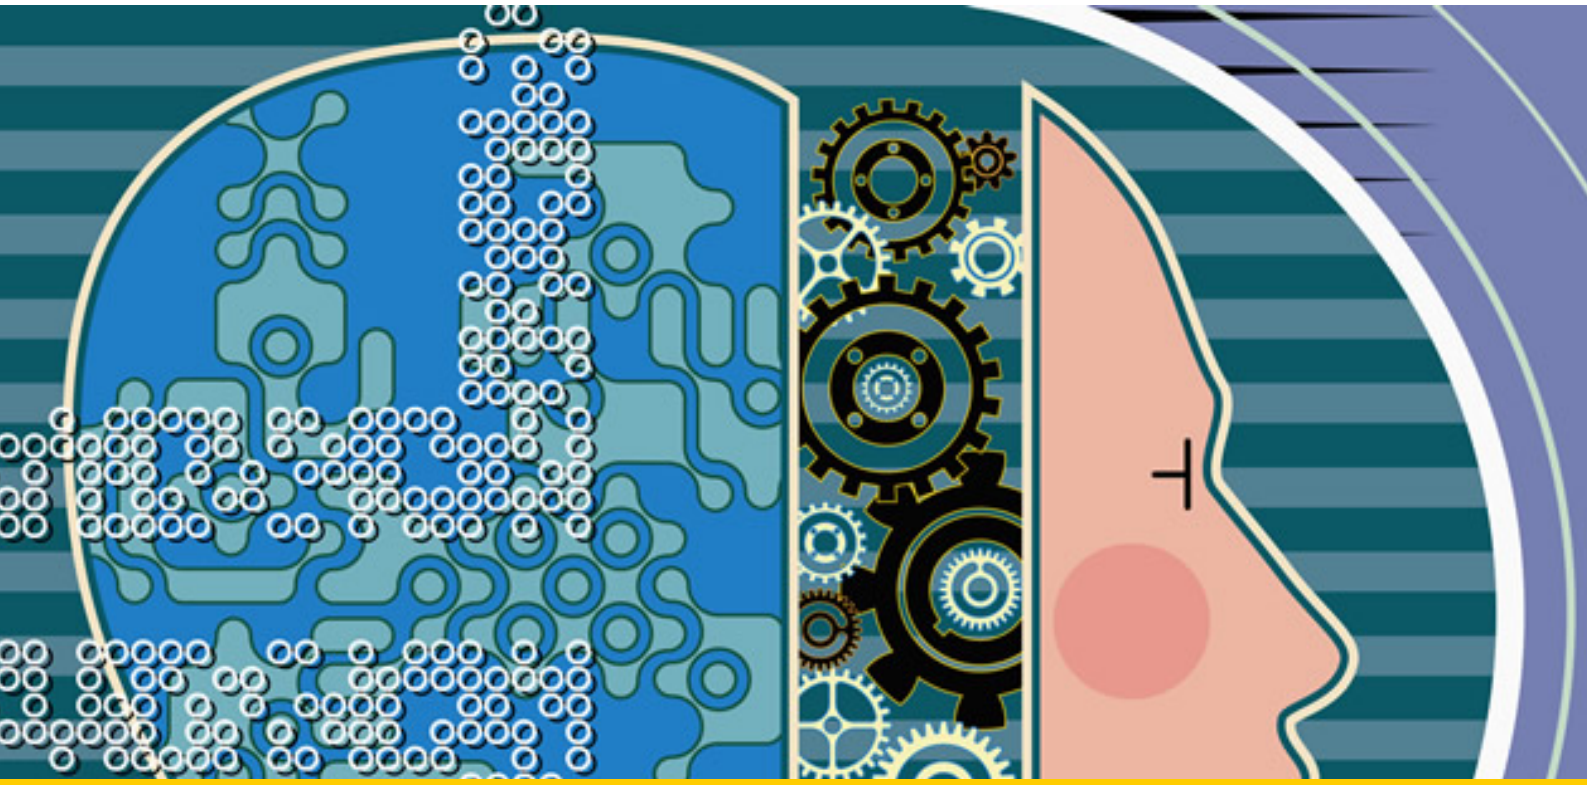

HUMAN FACTORS WORKING GROUP WHITE PAPER

Human Vulnerabilities in Security Systems

M. Angela Sasse (UCL), Debi Ashenden & Darren Lawrence (Cranfield University), Lizzie Coles-Kemp (RHUL), Ivan Fléchais (Oxford University), Paul Kearney (BT)

Introduction

This White Paper addresses a specific problem in security - human vulnerabilities that lead to security being breached. The aims of the White Paper are:

1. **To provide a description of the problem: what human vulnerabilities are, why they occur, and how their occurrence puts systems at risk.**
2. **To identify knowledge and techniques from a range of research areas that can be applied to remove or mitigate human vulnerabilities.**
3. **To provide an indication of the effectiveness of available measures, and under which circumstances they are most likely to succeed.**
4. **To identify a set of challenges for managing human vulnerabilities, arising from technical, economic and societal developments.**
5. **To present promising research directions for addressing them.**

The primary target audiences of the White Paper are

1. **Those in charge of securing the assets of their organisations (company owners, managers, CEOs).**
2. **Those charged with delivering security (security officers, system administrators, CSOs, security service providers and security researchers).**

But a key message of the paper is that security is everybody's business. To manage human vulnerabilities effectively, we need to involve all stakeholders in the design and operation of secure systems, and communicate effectively about risks, and everybody's roles and responsibilities in managing it.

Definition

Security, in its simplest sense, can be described as *"things that should happen, do, and things that shouldn't happen, don't"*. Organisations are socio-technical systems, consisting of human and technical elements. A vulnerability is weakness in the system that makes it more likely that an undesirable event happens, or a desired one fails to happen. Technical systems have vulnerabilities, and so do humans. We define human vulnerabilities as *"human characteristics and behaviour that create, exploit, or contribute to the exploitation of, a system vulnerability."*

1. **The definition covers unintentional behaviour (mistakes that create vulnerabilities, or facilitate their exploitation), as well as intentional malicious behaviour (exploiting a vulnerability). Given recent developments (see *Problem Space*), effective security management**

must acknowledge that every human within an organisation has the capability to become an attacker.

2. **Security breaches caused by unintentional behaviour cannot simply be blamed on individuals. The underlying causes for human vulnerabilities are systemic failures, caused by errors in the design of systems and management of people.**
3. **Human vulnerabilities need to be identified and managed before they lead to an actual breach of security. Currently, there is a tendency to ignore undesirable characteristics and behaviour until a breach of security occurs, but unacknowledged vulnerabilities cause systemic failures (see point 2).**
4. **It would be wrong simply to equate human vulnerabilities with "undesirable behaviour." A characteristic or behaviour that leads to a security breach in a specific context may be highly desirable in another (for instance, being helpful to customers, trusting a colleague).**

Problem Space

Why have human vulnerabilities have become a key concern? Firstly, they are a side effect of the success story of information technology: 20 years ago, interaction with information and communication technology was the preserve of a few professionals. Today, most businesses in the UK make extensive use of ICT, and nearly 60% of households have PCs and broadband. With government services moving online, citizens who have no technology skills are in danger of being disenfranchised. Given the speed with which technology has been adopted, it is not surprising that there is a knowledge and skills gap when it comes to technology in general, and security in particular. The gap affects not only individuals, but also businesses and other organisations. Even professionals who design, build, and supply technology do not always understand the security implications of their decisions.

The number and variety of threats for networked systems is increasing rapidly, and with it the number of security measures. Most security mechanisms are currently chosen to protect the technology, with little or no consideration of the impact on individuals. This compounds the effect of increasing system complexity. Many existing mechanisms create an high workload for individual users. The cumulative effect of the demands of several such security mechanisms (e.g. passwords and PINs) at work and at home means that many individuals simply cannot cope and make mistakes. This negative experience of security antagonises people, reduces their motivation to follow security policies, and creates negative attitudes towards security. Security mechanisms are also chosen without considering the impact on the business processes they are meant to protect - they get in the way of "getting the job done" and induce people to bypass them. High workload, complexity and habitual bypassing of security mechanisms not only increase the

likelihood of mistakes, but also create many opportunities for new types of attacks.

Few individuals are currently self-motivated to overcome the security knowledge gap. A key factor is the traditional "command-and-control" approach to security, which casts humans in the role of a component that needs to be managed through a set of policies that are constructed top-down, and enforced through sanctions. This approach can work, but is expensive and does not fit well with many modern organisations with flat management hierarchies, which encourage their knowledge-worker employees to be flexible, manage themselves, and take initiative.

The second key factor underlying current developments is that organisations are changing in order to be competitive in a global economy. Increasingly, business is conducted at a distance, with people we never meet, and who operate in a different social, cultural and legal context. Growing integration of supply chains, outsourcing, off-shoring and virtual organisations means established, perimeter-based security concepts no longer apply: it is harder to draw the system boundary and to determine who should have access to which systems or data. The deperimeterisation of security has changed the technical mechanisms used to protect systems, but thinking on how to manage the human element in security has not moved on from the command-and-control approach. Globalisation has also led to changes in the relationships between organisations and employees. Whilst knowledge workers in some organisations have become more empowered and participate actively in management, in other sectors, many employees feel less loyalty towards their employers. Information technology and flatter management hierarchies have reduced the number of administrative and lower management jobs. Organisations are able to interact with their customers at a distance, reducing the need for local branches. The creation of call centres has provided many low-skilled clerical jobs, but in the global economy, such jobs are constantly under threat from outsourcing and off-shoring. This is also true for some types of knowledge work, notably software development. Employees who have no expectation of job security feel less loyalty towards their organisation, which means they are likely to be tempted when they come across an opportunity to defraud the organisation, or when offered money in exchange for data or information, and is now easy to copy and remove large amounts of data. Many organisations respond to these human threats with increased monitoring and surveillance. The downside of these countermeasures is that they can be expensive, threaten employees' privacy, reduce trust and loyalty, and hence contribute to employee disaffection. Monitoring and surveillance tends to be aimed at lower-level employees. Insider attacks with the most serious consequences, however, tend to be perpetrated by executives, who often enjoy unlimited trust and access to systems. The role of trust between humans and its

economic implications are currently not well understood in the security community. The technology platform that mediates business interactions and human interactions today often fails to support people in making good trust decisions.

Understanding risk is a key part of effective security. It is a fundamental aspect of ISO 27001 (the international standard for information security) and it underpins an organisations' ability to prove compliance with a raft of corporate governance requirements. However, in many organisations, information security risk assessment takes the form of "dashboards" or checklists to satisfy legal or regulatory requirements. Compliance does not lead to effective security if it becomes a box-ticking exercise. There is a tendency to focus on achieving the best measurement of risk, whilst the need to manage the perception of risk and associated behaviours within organisations is overlooked. Owners of business processes often fail to engage in risk analysis, seeing it as task for specialists. Whilst expert advice and support is undoubtedly needed, the wholesale abdication of responsibility to IT/security experts has led to security measures that are disconnected from the business processes they are meant to protect, and leads to a choice of security measures that interfere with business processes and threaten profitability.

What Can Be Done?

In this section, we examine how techniques that promote learning and participation can be applied to manage human vulnerabilities in organisations.

Managing Risk

An information security management system (such as that embodied in ISO 27001) offers a systemic approach to security management, with security objectives set by risk assessment and supported by policies. Policies are an important communication tool - they should signal that senior management supports security goals, and provide an explanation of security decisions to the whole organisation. Policies are also one of the key interfaces through which humans interact with security mechanisms, and it is in the interpretation of policies that human vulnerabilities often occur. In our view, the best way to manage human vulnerabilities in a modern organisation is through communication and co-operation between management, security practitioners and employees. Correctly implemented risk management systems deliver continuous feedback on the effectiveness of security measure through audit, incident management and risk assessment. To obtain a complete and honest assessment, blame-free incident reporting must be put in place: security incidents should be seen as opportunities for learning, where the learning takes the form of corrective action review (single-loop learning), and in addition may take the form of preventive action review where the underlying causes of the incident are

challenged (double-loop learning - see Organisational Behaviour).

Trust

In information security, it is a commonly held belief that people need to be protected because they are not capable of making good security decisions. People are commonly blamed for security problems (such as disclosing their passwords or downloading malware). This leads to the widespread perception that trusting people to behave securely is something to be avoided. However, trust is a pervasive phenomenon and economically powerful: in successful trust relationships, both trustor and trustee derive economic benefits. Organisations who can trust their employees save money: expensive assurance procedures can be avoided, and employees are more productive, cooperative and creative. In the context of human vulnerabilities in security, it is useful to distinguish between three classes of relationship involving trust:

1. **Trust between a human and the technical infrastructure.**
2. **Trust between humans, especially where the relationship is a mediated by technology.**
3. **Trust between a human and the organisation within which he or she functions.**

In all three cases, a system designer must aim to support people in making correct decisions about the trustworthiness of the other entity, and to foster (well-placed) trust between the entities.

Design

One of the key aspects of designing and developing secure technical systems revolves around correctly identifying the requirements for those systems. All stakeholders (owners, employees, customers) have some security requirements, but may not associate these requirements as “security”, or be able to express them in security terminology. As Dieter Gollmann states: *“security-unaware users have specific security requirements but usually no security expertise”*. One way of addressing this issue is to adopt a participative approach to security analysis and design - involving the stakeholders in the technical discussions and decision-making surrounding security design. Through participation, stakeholders can gain a better understanding of security issues able to communicate their own security needs. Security design also needs to take into account the wider business perspective. The cost of a security mechanism (e.g. purchase cost, reduced productivity, user workload) needs to be balanced against its effectiveness and usefulness. Security needs to be a productive addition to the business. Security mechanisms often are unusable because they are a technical “one fits all” solutions added onto systems, irrespective of their context of operation. Usable technology is fitted to the requirements of its particular users, into the tasks or

business process they are working on, and into the physical and cultural context of use. Human factors researchers distinguish between production tasks and supporting tasks. Production tasks are activities that are directly linked to the main goal (e.g. getting the product out of the door). Security supports the long-term survival of the business process, but it does not contribute to production in the short term. When there is competition for resources (a person's time, attention, memory), production tasks are prioritised – especially when remuneration is linked to productivity. Security mechanisms that create an unreasonable physical or mental workload are bypassed altogether (e.g. for a door protected by a keypad, used by people carrying loads, will end up permanently propped open). Security must be integrated into people's tasks and business processes, rather than interfere with them. Key design principles are:

1. **Identify the performance requirements of the production task, and make sure the security task does not significantly reduce productivity.**
2. **Minimise the physical and mental workload of the security task; use a mode of interaction that fits with the production task activity (e.g. use voice-based mechanisms telephone-based interactions, a hands-free mechanism for tasks where both hands are occupied).**
3. **For frequently executed security mechanisms, design for speed; for infrequently use mechanisms, design for memorability (step-by-step user guidance, recognition-based interfaces).**
4. **Minimize the scope for error. Human Factors research, especially on Human Error, provides ample guidance on how to design systems that minimise the likelihood of error, and the impact of errors. Systems must be designed such that a single error by an individual does not lead to serious security incidents.**
5. **Incentivise secure behaviour (as well as - or as an aspect of -productivity goals).**

In addition to error prevention, design should aim to reduce the opportunities and motivation for intentional attacks. Research in criminology has produced frameworks for situational crime prevention, which offer a range of options for reducing the likelihood of attacks. Finally, design can incorporate persuasive techniques to influence risk perception and reinforce appropriate behaviours.

Awareness, Education, Training

The terms "security awareness", "security education" and "security training" are often used interchangeably, but ought to be thought of as 3 steps on the route to changing people's behaviour. The role of security awareness is to get people interested in security – attract their attention, make them realise affects them and others they care about, that it is something

worth paying attention to. Attention can be attracted through striking design, and surprising or humorous messages. Good awareness-raising material is like good advertising, and should be designed in collaboration with experts. Once aware, people are more likely to respond to security education – materials or courses that provide information about the threats and vulnerabilities, and what actions they should take to protect themselves or the organisation. Education can be delivered through tutorials on websites, for instance, but the material must be designed so people can find an answer quickly when they need it (e.g. *"How do I construct a password that is secure and memorable?"*), and provide sufficient depth of understanding to equip people for dealing with uncertainty and complexity in security decision-making. Awareness and education may prepare the ground, but actually changing people's behaviour – breaking old habits, establishing new ones – requires training. Training is a programme in which new behaviours are not only presented, but tested and corrected. To be effective, security training must be based in the work context and address specific security needs. Collaboration with training professionals is advisable. Once the 3 steps have been effective, regular reminders of the key messages are needed. There are commercial providers of screen savers, posters and trinkets that can fulfil this role, but organisations opting to buy these should bear in mind that (1) reminders can only reinforce messages delivered through an AET approach, not replace it, and (2) that materials perceived as "silly" or "cheesy" can decrease respect for security, and thus be counterproductive.

Organisational Behaviour

There is a need to educate and persuade people to think and act in a security-conscious way. Human behaviour in the workplace is the focus of a wide-ranging body of literature that could illuminate and support initiatives aimed at securing competitive advantage and resilience to organisational threats. Whilst there is currently no specific literature on security behaviours, there is ample literature on other indicators that most likely have common roots with issues associated with security behaviour, e.g. productivity, job satisfaction, staff turnover, absenteeism and sickness rates. There are formal and informal aspects that govern the behaviour of people in organisations. Formal security policies, like job descriptions, can only specify a fraction of the positive security behaviours needed from employees within their everyday work roles. To meet this shortfall between what is prescribed and what is actually required, organisations rely on organisational citizenship behaviour. This is related to the jobs people do, but not necessarily tied to formal reward systems. These informal aspects of work roles and relationships can be managed though psychological contracts, which most organisations negotiate with their employees. Security behaviours should become part of such contracts, which have to be based on concordance (mutual understanding and agreement). Literature on human behaviour has shown that concordance is a more successful

route to behaviour change than the command-and control approach currently prevalent (people should do what they are told, and carry the blame if they fail to do so).

Research needed

Having reviewed the problem space, and what existing knowledge from a range of disciplines can contribute to the management of human vulnerabilities, we have identified a number of areas that are potentially promising, but need further research to develop knowledge, methods or tools that can be used in practice.

1. **Crafting cultural change: How can an organisation develop and maintain a healthy security culture? Which psychological contracts promote good security behaviour? Can we link psychological contracts and security policies?**
2. **Creating Consent: How do we move from command-and-control, expert-led security to participative security models? How do we manage stakeholder participation in the design and operation of secure systems?**
3. **Persuasive technologies: Can we harness technology to promote good security behaviour, and dismantle unhelpful perceptions and attitudes?**
4. **Opportunity reduction: Why do some employees exploit opportunities in systems, and how do they do it? Can we adapt organisational design approaches for managing safety (such as Human Error) to security? How can we identify and remove opportunities that tempt people?**
5. **Balancing vetting/monitoring against privacy and trust: How can organisations balance the need for (and potential benefits of) trusting employees with the need to protect their assets? How do we balance the need to vet employees and monitor their behaviour with their privacy needs and data protection rights?**
6. **Distributed Security Management: How do we manage human vulnerabilities in a deperimeterised security environment? How should roles and responsibilities be assigned? What tools can support effective communication about risks and incidents in a distributed environment?**
7. **Design: There is a need for tools and toolkits for managing security design, including representations of human vulnerabilities, and for security design patterns that encapsulate usability principles.**
8. **Modelling a human-technical system for security decision-making: Can we establish a catalogue of undesirable behaviours, and integrate them within the frameworks provided by established (mathematical) systems modelling techniques? How can we model the economic value of the security measures for human vulnerabilities, and the effort needed to operate them?**

9. **Next-generation risk assessment: How can we integrate risk assessment with system and policy design? Can we develop mechanisms for evidence-based threat assessment?**
10. **Risk communication: How can we change public discourse about security, communicating that security-conscious behaviour is desirable and attractive, rather than nerdy and paranoid?**

Conclusions

It is important to understand that human vulnerabilities in current systems are a consequence of the success story of ICT uptake, and global economic trends. And while there can be no doubt that the number of security breaches caused by human vulnerabilities has been increasing rapidly, the problem is a tractable one. Together, engineering and social science research can explain why the vulnerabilities occur, and this understanding provides the first step towards addressing them. We have identified a number of actions that organisations can take today, and promising ideas that need to be explored and developed through further research.