# A Model of Security Culture for e-Science

S. Faily and I. Fléchais

Computing Laboratory, University of Oxford, Oxford, United Kingdom
e-mail: shamal.faily@comlab.ox.ac.uk

## Abstract

There is a need to understand the cultural issues affecting security in large, distributed and heterogeneous systems; such systems are typified by e-Science projects. We present a model of security culture for e-Science, grounded both in the security literature and in empirical data from an e-Science project. From this model, we present five concepts, which have differing effects on security culture. Each concept is discussed in terms of how the literature treats it, and how it impacts security culture in practice. This discussion highlights differences and similarities between the two domains.

## Keywords

Security Culture, e-Science, Roles, Responsibility, Context, Sub-Cultures, Requirements

## 1. Introduction

Designing for a multi-organisational context, with no coherent organisational control, raises social, as well as technical, challenges. Software engineers are ill-equipped to reconcile the myriad of values people hold about assets, controls, risks and usability in these environments. Consequently, accidental complexity arises due to conflicting values; security mechanisms usable to members of one sub-culture may be unusable to members of another.

e-Science is concerned with global collaboration in key areas of science and the next generation of infrastructure that will enable it (Taylor, 2001). As e-Science grows to encompass the needs of culturally disparate stakeholders, so too will the impact of security on them. The impact of security artifacts on end-user communities in different national cultures has been explored (Singh et al., 2007), as have conflicting values between end-users and system administrators (Adams & Sasse, 1999)(Kraemer & Carayon, 2007), and end-users and security developers (Zurko & Simon, 1996)(Whitten & Tygar, 1999). Yet, the distinction between roles such as user, developer and administrator begins to blur in e-Science.

Understanding different Security Cultures within the e-Science community may support information systems design by reducing accidental complexity. However, Security Culture remains a hackneyed term in case studies, and little evidence exists for the applicability of insights in single organisations scaling up to the dynamic, multi-organisational contexts found in e-Science. Moreover, much of the previous work on Security Culture does not appeal to the multifaceted nature of culture.

Instead, security culture is described as a concept influenced by security awareness (Helokunnas & Kuusisto, 2003)(Da Veiga & Eloff, 2007) or obedient behaviour (Thomson et al., 2006)(Thomson & Solms, 2005). Intrinsic case studies have analysed Safety Culture within particular contexts (Cooper, 2000)(Haukelid, 2008), but these findings may not be universally applicable; safety engineering is primarily concerned with unintentional, rather than intentional failure.

This paper presents a model of Security Culture for e-Science, grounded in the literature and validated through empirical research. In section 2, we detail the method used to build a model of Security Culture from the relevant, peer reviewed literature, and describe how this model was empirically validated. In section 3, we analyse this model before comparing and contrasting the roles played by five concepts shaping this model, both within the literature and in practice.

## 2. Method

To derive meaning about Security Culture in e-Science, an analytical induction approach was taken. Grounded Theory (Corbin & Strauss, 2008) was selected as the methodology for this approach, as this prescribes procedures for generating theory for observed real-world phenomena. We analysed two data samples: one based on a selection of the research literature on security culture in general, the other based on interviews with stakeholders in an e-Science project. Our methodology consisted of building one model for each sample, and then comparing and consolidating both into a single, unified model of security culture.

A literature-based model of Security Culture was grounded in a sample of 21 peer-reviewed papers from the safety and security culture literature. Analysis was initially performed on a sample of 17 papers. Based on emergent concepts from the axial coding, a further 4 papers, considering these concepts, were added to the sample before the theory was considered saturated -- the point of analysis when further data gathering and analysis added little to the model.

The corpus of empirical data used to ground an empirical model of Security Culture was collected from participants of the NeuroGrid project. NeuroGrid was a UK Medical Research Council funded project to develop a Grid-based collaborative research environment for different clinical researcher communities (Geddes et al, 2005). NeuroGrid was used by three different clinical exemplars: Stroke, Dementia, and Psychosis. The sensitivity and distributed nature of the clinical data drove the need to find secure and effective ways of accessing and managing it.

Qualitative interviews were held with various participants of the NeuroGrid project, amounting to approximately 500 minutes of transcribed data. To ensure a balanced coverage, two interviews were held with members of each clinical exemplar. Interviews were also held with developers, managers, and researchers involved with NeuroGrid, but not associated with any particular exemplar.

The method for building the comparative model of Security Culture from the empirical data was identical to that used to derive a security culture model from the

literature. The models were compared to validate whether the tenets of the Security Culture theory held in both cases. The empirical model was initially induced independent of its theoretical counterpart. As an axial theory began to emerge from the empirical data, both models were refined. The empirical model was refined where identical concepts were synonymous with more grounded concepts in the literature. The theoretical model was refined when the empirical data provided insight to concepts in the literature, which were previously latent within the model. Variances were found in both models; nevertheless a single, unified theory of Security Culture was evident in both cases.

## 3. Results

The method described above produced the model of Security Culture for e-Science shown in Figure 1. In the following sections, we highlight the characteristics of intangible and tangible factors, consider key concepts influencing this model, and compare the literature's perception of these concepts with how these are realised in practice.

### 3.1. Resultant Model of Security Culture?

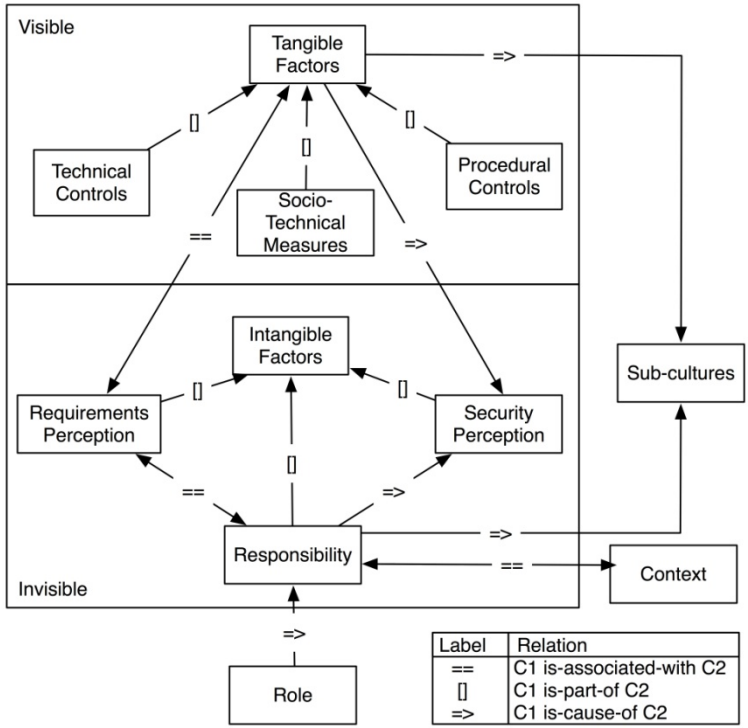


**Figure 1: Security Culture network diagram**

Many authors describe Security Culture as a mentally perceived concept, such as security awareness or obedient behaviour, but we believe such a definition fails to explain a larger picture. Like other writers on Security Culture, we take inspiration from Schein and his layered model of organisational culture (Schein, 1992), but we also consider the work of those who conceive culture as an ordered system of symbols where meaning is based on individual participants, rather than the organisation as a whole, e.g. (Hirschheim & Newman, 1991).

We define Security Culture as a combination of Tangible Factors and Intangible Factors within both an organisation's culture and its subcultures.

Intangible Factors are invisible assumptions, norms, and values of a culture's participants. Tangible Factors are visible artefacts of a culture or subculture. These artefacts are represented by technical controls, procedural controls or socio-technical measures. Technical controls represent the mechanisms controlling security; these include passwords and digital certificates. Procedural controls are organisational policies and procedures reflecting the presence of security; these include security policies, instructions for using technical controls, and guidelines for secure data handling. Socio-technical measures augment both technical and procedural controls, and are designed to increase the potency of intangible factors. Examples of these measures include security awareness programmes, and guidelines for ethical conduct.

3.1.1. Intangible Factors

Although we identified several classes of intangible factor, space restrictions only allow us to describe a small subset of them. Requirements Perception is the perception of how security requirements should be designed into the system. Responsibility is a consequence of a person's organisational role, several of which were identified as playing a part in discharging organisational security obligations. Security perception is the mental perception of how security is managed, allowing members of an organisation to be sensitive to threats and vulnerabilities. Several other concepts were found to have moderating influences on security culture, the most prominent of which were roles, context, and sub-culture norms.

Although the models emerging from both the literature and empirical data were similar, two major differences were identified. First, although the relationship between tangible and intangible factors is bi-directional in theory, it only appeared to be unidirectional in practice. While perceptions of security inform socio-technical measures and security policies, there was no evidence this occurred in NeuroGrid. Second, the theoretical model indicates the emergence of different sub-cultures follows from intangible factors. There is, however, no evidence of these sub-cultures influencing any intangible factors. In the empirical model, this relationship is unidirectional but in a different direction, as there was evidence of sub-cultures influencing the perception of security.

### 3.1.2. Tangible Factors

The most visible tangible factor in NeuroGrid was its technical controls, specifically access-control policies and digital certificates. Digital certificates controlled access to the NeuroGrid web service interfaces, although data access was additionally facilitated by XACML access control policies. These policies were specified by the data owners within each clinical exemplar, but were manually handcrafted by the NeuroGrid Security Team.

The most visible control cited by all of the participants was the [X.509] digital certificate. In most cases, perceptions of the control were influenced by usability problems, not in their day-to-day usage but in their initial set-up. One administrator described how, despite the presence of step-by-step instructions, the process of certificate installation was considered onerous enough to dissuade many potential NeuroGrid users.

Access control policies were prominent to the NeuroGrid system administrators, but largely invisible to end-users and application developers. These policies were recognisable only as an artefact under the control of the infrastructure team, and as something to initially write, based on the requirements of data owners, or as something to be breached. The evidence supporting this perception may be biased by the tedium associated with manually authoring XML policy documents, although it did reinforce the idea of technical controls being considered only within a limited context.

The mixed-visibility of controls in different contexts led to a formulation of security perception based on incomplete information. This is illustrated by differing descriptions of how security is mandated in NeuroGrid by different users. Some users believed that access control was based on the issue of passwords to the NeuroGrid portal, while in reality it was based on digital certificates and access control policies.

The literature suggests that procedural controls indirectly document the values held by an organisation's Security Culture. In many ways, this proposition remains valid based on the empirical data. Information security policies or guidance did not exist for NeuroGrid. Such guidance could have been written based on the documentation that did exist, but no participant was aware of procedural controls beyond those encountered on a day-to-day basis. This may be partially explained by a dichotomy between control and data responsibility, such that guidance focuses on one or the other, but not both. An example of a procedural control focusing on technical controls was the comprehensive on-line guidance for installing certificates.

Although not as prominent as technical controls, procedural controls influenced the norms of different sub-cultures. For example, participants who were sensitive to control usability problems described useful security guidance as that which focused on technical controls. Similarly, participants, sensitive to the value of data assets they worked with, cited guidelines and procedures for secure data handling; these

included Medical Research Council policies, institutional procedures for handling patient data, and DICOM standards for anonymisation.

### 3.2. Key Concepts

The following sections describe key factors influencing the e-Science Security Culture model. For each factor, we examine how this is reflected in the literature, before considering how these factors are evident in practice.

#### 3.2.1. Role

Responsibility in the Security Culture literature stems from two roles. The first of these, management, is responsible for imbuing information security into the organisational culture. The most cited means of doing this is policy communication. The second, organisational peers, are members of an organisation responsible for understanding and discharging their security obligations. Although such obligations are ever-present to organisational peers, they cannot be completely responsible for their security, due to many issues beyond their control (Furnell, 2008). Users are, however, expected to understand their role and how to fulfil it in a secure manner. The literature also suggests that, in addition to responsibilities, sub-cultural norms and values can also evolve based on roles.

Analysing the different roles within the empirical data points to a more elaborate taxonomy, with different classes of software developer and end-user. For example, rather than a single organisational collective, more refined classes of users were present. Some were user proxies, testing NeuroGrid applications on behalf of end-users. The end-users, who were usually clinical researchers, were either data providers, data consumers or a hybrid of both. Data providers, several of which had been delegated responsibility for data by their managers, were responsible for access control policy decisions. Data consumers were researchers who had been granted access to subsets of NeuroGrid data in order to carry out clinical research. Their primary focus was not on the sensitivity of the data they were handling, but on the research they were carrying out. In many cases, security was only visible to these different roles by the technical controls constraining access to data. In contrast, much of the data cited in the literature is based on small, less distributed environments; these do not yield different classes of developer due to their comparative scale. The management role was much less grounded than suggested by the literature because, as a science project, more authority and responsibility was delegated than commonly found in industry.

#### 3.2.2. Responsibility

The literature indicates that clear definitions of responsibility lead to increased security perception. The security responsibilities explicit from the literature were organisational and moral responsibility. Organisational responsibility represents the accountability mechanisms in place for justifying and managing security management decisions. Moral responsibility represents instilled norms, which allow security concerns, affecting the wellbeing of an organisation, to take priority over

other internalised norms in the organisational context. This sense of moral responsibility is only evident once users are made aware of their organisational responsibility, and sensitised to information security issues.

(Fléchais et al., 2005) report that a shared sense of moral responsibility can lead to benevolence, the property allowing an agent to gain gratification from the well being of another agent. This property can engender Security Culture, but may also lead to attacks when a security policy is inadvertently subverted. An example of such an attack is presented by (Furnell, 2008), where benevolence, and the desire to increase one's social capital, allowed users to reveal personal information to a potentially malicious stranger. While not part of the sample, benevolent actions are also described by (Miller, 2002) as costly to the performer, contextual, and based both on the perception of need and a moral judgement of the needy's predicament.

Like roles, the concept of responsibility is more elaborate in practice than in theory. Rather than a singular concept of organisational responsibility, responsibility, and thereby accountability, is split between technical controls and assets. Additional forms of responsibility evident in the empirical data, as perceived by the participants included:

- The responsibility of applications to safeguard delegated security controls.
- The responsibility of participants to safeguard sensitive data, curatorship of which has been delegated to them.
- Legal responsibility of data as described by the UK Data Protection Act.
- The responsibility of participating institutes with regard to line management of participants and physical control of data.
- Moral or ethical responsibility to safeguard the privacy and anonymity of confidential data.

Adding to this myriad of responsibilities were related concepts, which tempered or weakened levels of responsibility. All of the participants interviewed were aware of the sensitivity of data within NeuroGrid, this tempered both the moral responsibility of individuals, and the perception of management. Conversely, management perceptions regarding the low take-up of NeuroGrid -- a possible conflict with their institutional responsibility -- led to an evident security and usability trade-off.

### 3.2.3. Sub Culture Norms

Security Culture is ideally considered in the singular, but the literature suggests the reality is more complex. Security Culture may nest other sub-cultures, which vary between organisational units. Moreover, not only can security perception vary between these sub-cultures, but members of the sub-cultures can affect security controls based on their perception of other sub-cultures' security perceptions. For example, (Kraemer & Carayon, 2007) describe how network administrators “walled around” a particular network because they perceived its users to be uninterested in security. Network administrators viewed end user errors as intentional, yet considered their own errors as unintentional. Security sub-cultures may also be

microcosmic of the management of organisational culture, and cultural stereotypes can arise not from organisational units, but from the style of management adopted.

Each NeuroGrid clinical exemplar appeared to constitute a sub-culture. The security norms of each sub-culture were most evident when participants described their handling of data. Sub-culture values were also evident from participants' descriptions of controls. For example, participants in one exemplar espoused strong obligations for anonymising data, to make it safe for release to NeuroGrid; after anonymisation, the controls were considered superfluous. Another exemplar anonymised data, although not to the extent that their own applications would be rendered useless. This led to a sub-culture with a strong reliance on security afforded by technical controls. In both cases, different perceptions of security were held by each sub-culture, based on norms and values associated with data and controls.

Indifference to security issues perceived to be beyond their sphere of control was observed within the sub-cultures. Sub-cultures close to the technical controls perceived them as a means of rendering NeuroGrid secure, irrespective of the context. Sub-cultures close to the data perceived the process of obtaining ethical approval, to use the data, sensitised them to the principles of information security. This phenomenon is evidence of diffusion of responsibility, as no one individual or party wishes to take holistic responsibility of security. This notion has been examined in depth by (Darley & Latané, 1970), who concluded that rather than benevolence being dispensed universally, its propensity depends on prejudice, or details of a particular situation (Keltner & Marsh, 2006).

3.2.4. Context

The literature talks about operational and cultural contexts. Operational contexts are governed by procedural controls to determine acceptable conduct, yet they are also organic in that individual values and management influence help shape them. The environment and the passage of time also influence these contexts. Underpinning these operational contexts are cultural contexts, which are also organic. The potential for cultural conflict arising from conflicting organisational and cultural norms is ever-present. (Singh et al., 2007) describe how sharing passwords and credentials is necessary when a visitor to a commercial hub needs to shop or carry out business on behalf of an entire outlying community, even though such practices are disallowed by many banks and financial organisations.

These insights into operational and cultural contexts remain valid in the empirical data. Participants described NeuroGrid contexts of use in terms of workflows, and it was these that were built into the design of NeuroGrid. However, several other operational contexts of use were also identified from the empirical data; these include set-up of technical controls, the process of data anonymisation, and the development and maintenance of NeuroGrid end-user applications.

Cultural contexts within NeuroGrid were based on sub-cultures within the project. Artefacts familiar to some sub-cultures but not others influenced perceptions of

usability when their affordances were not explained or illustrated. For example, data sharing on NeuroGrid relied on WebDav technology, but the term WebDav folder persisted as a source of frequent confusion until either explained, or illustrated using the NeuroGrid portal. In another example, two participants with similar roles and levels of expertise, but working within different sub-cultures, were asked to describe the process of installing the same security control. To one participant, the operation was described as a trivial exercise. To the other, the description made explicit several assumptions not apparent by the first participant. Within the cultural context of the first participant, the handling of security controls was frequent, but for the second the control was only a link within a larger application chain.

3.2.5. Requirements Perception

The term requirement is prevalent in the literature; requirements are described as a panacea for resolving security issues. Security requirements should factor education and training, conform to external audit and governmental requirements, and help align employee behaviour towards compliance with organisational security goals. These definitions lead us to ask what the literature means when referring to requirements?

Despite the value of requirements, their use on NeuroGrid was limited to informing the detailed design of the NeuroGrid infrastructure. While the team developing the NeuroGrid infrastructure believed it had agreed security requirements, these appeared to conflict with performance and usability requirements assumed by application developers.

The initial indifference towards requirements and the lack of consistent policies for compliant behaviour led to confusion about how security should best be achieved. The NeuroGrid requirements were written before the exemplar-specific workflows were developed. In some cases, these requirements were written on behalf of the clinical exemplars, who simply signed off the documents once they were complete. Inevitably, conflicting security perceptions did not become apparent until the workflows were documented.

Even though everyone interviewed held a high regard for security, the lack of agreed consensus about how to achieve this did little more than cement pre-existing security values held by individual sub-cultures. In one case, an application developer proposed modifying his NeuroGrid application to store user-supplied certificates, thereby allowing user access to be controlled via passwords; these were thought to be less cumbersome than the controls already in place. This subversion of the security requirements was considered necessary as many end-users refused to use NeuroGrid due to perceived security hurdles.

## 4. Conclusion

Factoring security culture into the design of information systems is an important part of securing e-Science systems for their different contexts of operation. This paper has made a contribution towards a better cultural understanding of secure systems

design. We have developed a model of security culture for e-Science, grounded in the literature, and supported by empirical data from the NeuroGrid project, and highlighted key concepts influencing this model.

## 5. Acknowledgements

The research described in this paper was funded by EPSRC CASE Studentship R07437/CN001. We are very grateful to Qinetiq Ltd for their sponsorship of this work.

## 6. References

Adams, A. and Sasse, M. (1999). Users are not the enemy. Communications of the ACM, 42, 41-46.

Cooper, M. D. (2000). Towards a model of safety culture. Safety Science, 36(2), 111-136.

Corbin, J. M. and Strauss, A. L. (2008). Basics of qualitative research : techniques and procedures for developing grounded theory. Sage Publications, Inc.

Da Veiga, A. and Eloff, J. H. P. (2007). An information security governance framework. Information Systems Management, 24(4), 361-372.

Darley, J. M. and Latané, B. (1970). Norms and normative behaviour: field studies of social interdependence.

Fléchais, I., Riegelsberger, J., and Sasse, M. A. (2005). Divide and conquer: the role of trust and assurance in the design of secure socio-technical systems. In NSPW '05: Proceedings of the 2005 workshop on New security paradigms 33–41.

Furnell, S. (2008). End-user security culture: A lesson that will never be learnt? Computer Fraud and Security, 2008(4), 6-9.

Geddes et al (2005, June). NeuroGrid: using grid technology to advance neuroscience. Computer-Based Medical Systems, 2005. Proceedings. 18th IEEE Symposium on, 570-572.

Haukelid, K. (2008). Theories of (safety) culture revisited-An anthropological approach. Safety Science, 46(3), 413-426.

Helokunnas, T. and Kuusisto, R. (2003). Information security culture in a value net. Engineering Management Conference, 2003. IEMC '03. Managing Technologically Driven Organizations: The Human Side of Innovation and Change, 190-194.

Hirschheim, R. and Newman, M. (1991). Symbolism and information systems development: Myth, metaphor and magic. Information Systems Research, 2(1), 29-62.

Keltner, D. and Marsh, J. (2006). We Are All Bystanders. Greater Good, 3(2).

Kraemer, S. and Carayon, P. (2007). Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. Applied Ergonomics, 38(2), 143-154.

Miller, D. (2002). 'Are They My Poor?': The Problem of Altruism in a World of Strangers. Critical Review of International Social and Political Philosophy (CRISPP), 5(4), 106-127.

Schein, E. H. (1992). Organizational Culture and Leadership. Jossey-Bass.

Singh, S., Cabraal, A., Demosthenous, C., Astbrink, G., and Furlong, M. (2007). Security design based on social and cultural practice: Sharing of passwords. In (476-485). Springer-Verlag.

Taylor, J. (2001). Presentation at e-Science Meeting by the Director of the Research Councils, Office of Science and Technology, UK. Available at: http://www.e-science.clrc.ac.uk, .

Thomson, K.-L. and von Solms, R. (2005). Information security obedience: A definition. Computers and Security, 24(1), 69-75.

Thomson, K.-L., von Solms, R., and Louw, L. (2006). Cultivating an organizational information security culture. Computer Fraud and Security, 2006(10), 7-11.

Whitten, A. and Tygar, J. D. (1999). Why Johnny can't encrypt: a usability evaluation of PGP 5.0. In SSYM'99: Proceedings of the 8th conference on USENIX Security Symposium 14–14.

Zurko, M. E. and Simon, R. T. (1996). User-centered security. In NSPW '96: Proceedings of the 1996 workshop on New security paradigms 27–33.