Human-Centered Specification Exemplars for Critical Infrastructure Environments

Shamal Faily
Bournemouth University
Poole, UK
sfaily@bournemouth.ac.uk

Georgia Lykou
Athens University of Economics & Business
Athens, Greece
glykou@gmail.com

Anton Partridge
Bournemouth University
Poole, UK
anton.partridge@me.com

Dimitris Gritzalis
Athens University of Economics & Business
Athens, Greece
dgrit@aueb.gr

Alexios Mylonas
Bournemouth University
Poole, UK
amylonas@bournemouth.ac.uk

Vasilis Katos
Bournemouth University
Poole, UK
vkatos@bournemouth.ac.uk

Specification models of critical infrastructure focus on parts of a larger environment. However, to consider the security of critical infrastructure systems, we need approaches for modelling the sum of these parts; these include people and activities, as well as technology. This paper present human-centered specification exemplars that capture the nuances associated with interactions between people, technology, and critical infrastructure environments. We describe requirements each exemplar needs to satisfy, and present preliminary results developing and evaluating them.

Critical Infrastructure; Specification Exemplars; Personas; Tasks; Risks; CAIRIS; Water; Rail

1. INTRODUCTION

Critical Infrastructure (CI), such as the water and rail sectors, are essential for day-to-day life. However, despite the attention given to parts of CI systems – such as water purification for water infrastructure, or train signalling in rail – there has been little work modelling the operating environments within which these parts are situated. Given the unforeseen circumstances that might arise due to complex interactions between people, technology, and the general environment, a security solution mitigating a risk in one type of CI system, may be inappropriate for addressing the same risk in another.

Specification exemplars are self-contained, informal descriptions of a problem in some application domain, and are designed to capture the harshness of reality (Feather et al. 1997). They can be used to promote research and teaching by introducing interesting and challenging problems, and provide a common model for evaluating solutions for the domain associated with the exemplar. Creating exemplars that address both of these needs can be difficult. For an exemplar to be useful, it needs to model different aspects of a problem, model a problem from different and potentially conflicting viewpoints, and deal with multiple sources of information.

Previous work by the authors (Faily et al. 2015) note that while specification exemplars focus primarily on modelling functional concerns, the nuances related to human issues are less easily modelled. By failing to model such nuances, exemplar users risk trivialising people and their work. In this paper, we present work designing and developing human-centered specification exemplars of nuanced CI environments. We describe five requirements for the exemplars before presenting preliminary results developing and evaluating them.

2. EXEMPLAR DESIGN PRINCIPLES

To address the issues in Section 1, we encapsulated five requirements into the design of each specification exemplar.

First, rather than being a textual description of a specific setting, each exemplar models the operating environment of a fictional CI company. Each model contains a goal model (van Lamsweerde 2009) representing the company's security policy and organisational constraints, asset models (Fléchais et al. 2003) describing the security properties associated with each asset, and floor plans of selected physical locations to provide context to how people and assets interact.

Proceedings of British HCI 2016 - Fusion, Bournemouth, UK

Figure 1: *CAIRIS generated model of the relationship between personas, tasks, assets, risks, and attackers*

Second, exemplars contain personas (Cooper 1999) of users in each environment, and tasks describing their typical work. Each task also contains information about how long it takes a persona to complete a task, how frequently the task occurs, how demanding the task is, and what conflicts may occur between different goals a persona might have. This information makes it possible to determine the impact that changing the environment might have on a persona's propensity for violating the security policy.

Third, each exemplar contains a selection of realistic vulnerabilities, attackers, threats, and risks specific to the type of CI system being modelled. By embedding these elements into the exemplar, one can see how vulnerabilities expose assets, and attackers realise certain threats to target assets by exploiting vulnerabilities.

Fourth, although exemplar models are static, model elements can be varied based on working contexts. As such, a threat with a high likelihood in one context may be insignificant or non-existent in another. Similarly, the same task carried out during the day might be more or less usable to a persona at night because the task might be truncated, or more stressful due to limited support in the event of problems.

Finally, exemplars are machine readable. For the purpose of our evaluation, exemplars were modelled as XML, to be compatible with the CAIRIS security design tool (Faily 2015). CAIRIS conforms to a meta-model for usable security (Faily and Fléchais 2010), enabling it to automatically generate visual models of how security, usability, and system elements interact with each other. For example, Figure 1 shows how tasks (blue ellipses) make use of certain assets (blue boxes) threatened or exploited by risks (red ellipses) within a given context of use. The model also shows the attackers and personas associated with each task and risk, and how usable personas find each task (different shades of blue). Although this model illustrates the complexity resulting from these different elements, CAIRIS provides facilities for filtering models, and generating documentation for some or all of the exemplar model. This allows exemplar users to focus on some aspect of a larger problem.

3. PRELIMINARY RESULTS

We developed two specification exemplars conforming to the requirements in Section 2; these are based on a fictional UK water company (ACME Water), and a fictional rail company in Southeast Europe (Balkan Rail). Each exemplar is grounded in empirical data from real CI companies. The data from ACME Water is drawn from two previous studies designing security for the water industry (Faily and Fléchais 2010; Faily and Fléchais 2011), and the data for Balkan Rail was collected specifically for the purpose of creating the exemplar. Both specification exemplars are publicly available (BANCIS Project Team 2016a,b).

The Balkan Rail exemplar is still under development, but the ACME Water exemplar provided context when evaluating a social engineering serious game (Beckers and Pape 2016). Personas, assets, and floor layouts were used as part of this game, where players were expected to devise social engineering attacks of people at ACME Water. Although successful, adoption of the exemplar was initially difficult due to the vast amount of information in the model. Consequently, when using the exemplar, one person was designated as a 'plausibility oracle', and consulted the exemplar in CAIRIS to determine the impact of proposed attacks on ACME Water. Future work will present a detailed evaluation and critical reflection of both exemplars.

4. ACKNOWLEDGEMENTS

REFERENCES

BANCIS Project Team (2016a). ACME Water Specification Exemplar. `https://github.com/failys/cairis/tree/master/examples/exemplars/ACME_Water`.

BANCIS Project Team (2016b). Balkan Rail Specification Exemplar. `https://github.com/failys/cairis/tree/master/examples/exemplars/Balkan_Rail`.

Beckers, K. and Pape, S. (2016). A serious game for eliciting social engineering security requirements. In *Proceedings of the 24th IEEE International Conference on Requirements Engineering*, RE '16. IEEE Computer Society. To Appear.

Cooper, A. (1999). *The Inmates Are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity (2nd Edition)*. Pearson Higher Education.

Faily, S. (2015). CAIRIS web site. `http://cairis.org`.

Faily, S. and Fléchais, I. (2010). A Meta-Model for Usable Secure Requirements Engineering. In *Proceedings of the 6th International Workshop on Software Engineering for Secure Systems*, pages 126–135. IEEE Computer Society.

Faily, S. and Fléchais, I. (2010). Barry is not the weakest link: eliciting secure system requirements with personas. In *Proceedings of the 24th BCS Interaction Specialist Group Conference*, pages 124–132. British Computer Society.

Faily, S. and Fléchais, I. (2011). User-centered information security policy development in a post-stuxnet world. In *Proceedings of the 6th International Conference on Availability, Reliability and Security*, pages 716–721.

Faily, S., Stergiopoulos, G., Katos, V., and Gritzalis, D. (2015). "Water, Water, Every Where": Nuances for a Water Industry Critical Infrastructure Specification Exemplar. In *Proceedings of the 10th International Conference on Critical Information Infrastructures Security*. Springer. To Appear.

Feather, M. S., Fickas, S., Finkelstein, A., and van Lamsweerde, A. (1997). Requirements and specification exemplars. *Automated Software Engineering*, 4(4):419–438.

Fléchais, I., Sasse, M. A., and Hailes, S. M. V. (2003). Bringing security home: a process for developing secure and usable systems. In *Proceedings of the 2003 New Security Paradigms Workshop*, pages 49–57. ACM.

van Lamsweerde, A. (2009). *Requirements Engineering: from system goals to UML models to software specifications*. John Wiley & Sons.