

RESEARCH DIRECTIONS FOR TRUST AND SECURITY IN HUMAN-CENTRIC COMPUTING *

Sadie Creese

*Systems Assurance Group,
QinetiQ, Malvern Technology Centre, UK.*

S.Creese@eris.QinetiQ.com

Michael Goldsmith

Formal Systems (Europe) Ltd.

www.fsel.com

michael@fsel.com

Bill Roscoe

Oxford University Computing Laboratory.

Bill.Roscoe@comlab.ox.ac.uk

Irfan Zakiuddin

*Distributed Technology Group,
QinetiQ, Malvern Technology Centre, UK.*

I.Zakiuddin@signal.QinetiQ.com

Abstract Pervasive networks foresee communicating and computing devices embedded throughout our environment. This will cause huge increases in the complexity of network infrastructures and the information services available over them. The challenge of managing information services, while maintaining security and privacy will be great. It is not clear that current security paradigms will map readily into such future environments. This paper outlines the authors' current position regarding the technical challenges which will need to be addressed in order to make secure pervasive computing environments a reality.

*This research is being conducted as part of the FORWARD project which is supported by the U.K. Department of Trade and Industry via the Next Wave Technologies and Markets programme. www.forward-project.org.uk

1. Introduction

The ubiquitous paradigm foresees devices capable of communication and computation embedded in every aspect of our lives and throughout our environment. This will increase both the complexity of information infrastructures and the networks which support them. New forms of interaction are envisaged, which will aim to push the technology into the background making the information services human-centric in delivery. Computing devices will be less and less noticeable, creating a feeling of being surrounded by “ambient intelligence”.

As these pervasive computing technologies become deeply intertwined in our lives we will become increasingly dependent on them, implicitly trusting them to offer their services without necessarily understanding their trustworthiness. Undoubtedly the timely provision of bespoke services will require personal or valuable data to be digitally stored and made available. The increased digitisation of our assets, coupled with the increasingly intangible way that networks use information, will make it difficult to ensure that trusted services are indeed trustworthy. Will users have to decide how to interact with systems without understanding the associated risks?

In this paper we organise and present our thoughts on a particularly important, often critical, property that will be required of such systems, namely *Information Security*. In addition to technical concerns we also devote attention to the role of human users and to the challenges of achieving trust and security in *human-centric* computing. The aim of the paper is to stimulate debate and to highlight and clarify the issues and problems that need to be addressed by the research community. The thoughts that we present in this short article are substantially influenced by our previous work on pervasive computing security: Creese et al., 2003a and Creese et al., 2003b.

2. Challenges to Information Security

The concept of *authorised access* is enormously important to security, underpinning most principal security properties:

- **Confidentiality.** Information is only made available to those who are authorised to have it.
- **Integrity.** Only authorised users may manipulate information.
- **Availability.** Information services must be accessible to those authorised.

Underpinning the notion of authorisation is that of *authentication*, which concerns proving the validity of an authorising claim. Traditional notions of authentication concentrate on the notion of proving the claim of an identity (if

identity can be proved, then this is a basis for authorisation). In Creese et al., 2003a we provided a critique of traditional identity authentication, arguing its unsuitability for pervasive networks because:

- Interaction would be between devices and it does not seem plausible that the identity of an arbitrary device, in an arbitrary environment, can be reliably determined. Furthermore in some applications mass-produced devices might not have unique identities.
- The value of authenticating an identity depends on the trustworthiness of the owner of the identity. If we do not know, either beforehand or by other means, that the owner of the identity is trustworthy, then little is gained by authenticating that identity. Thus, simply proving the identity of a device would be of limited value, since it provides little assurance that the device will behave in a trustworthy manner.

There were subsidiary reasons for doubting the value of identity authentication, such as the viability of certification infrastructures to support authenticating the identities of the huge numbers of devices that are likely to exist.

After presenting the above deconstruction we proposed that authentication for pervasive computing is revised to mean *attribute authentication*. Any device will have a range of attributes, such as its location, its name, its manufacturer, aspects of its state, its service history, and so forth. In a given situation some attributes will need authenticating and the attributes should be chosen to achieve assurance about *which* devices are the subject of interaction, and *what* those devices will *do*.

Protocols for authentication and authenticated key exchange have been the subject of intense study, a recent example of the growing literature is Boyd and Mathuria, 2003. Moreover, the subject of verifying such protocols has achieved significant advances, Ryan et al., 2001 contains some seminal work on the subject of verifying key agreement protocols.

For analysis and formal verification it is vital to be precise about the threat model which a given protocol must resist. The standard model of the attacker is due to Dolev and Yao, 1983, and it underpins a large portion of the research community's efforts. However, the Dolev-Yao threat model (as it is referred to) significantly predates the promulgation and widespread acceptance of the pervasive computing vision. In Creese et al., 2003b we proposed that such a threat model was too simplistic and unable to capture the authenticated key agreement protocols that might be required for pervasive networks. The principal amendment was to propose a "two-channel" threat model, as follows:

- 1 An *E*-channel which captures human or other "external" participation in bootstrapping an authenticated link. On the one hand, compared to the Dolev-Yao model, the attacker's capabilities on the *E*-channel were

significantly limited. But on the other hand the bandwidth for communication on the E -channel is assumed to be small.

- 2 An N -channel which captured the main medium for devices to create and perform secure electronic communications. The attacker would have similar capabilities to the Dolev-Yao attacker on the N -channel, but the bandwidth for communication is much greater.

A successful protocol for initialising a secure link in pervasive networks depends on sound use and interaction of the two channels. Our understanding of the literature to date has led us to believe that the two-channel threat model is a powerful abstraction capable of formalising a wide range of protocols.

Pervasive computing frequently makes the E -channel available thanks to the locality and context-dependent nature of authentication. And it may be *necessary* to use the E -channel due to the potential lack of ubiquitous PKI services and useful device identities.

Our two papers indicate how fundamental security parameters will change, as information services become pervasive. Both point to an increase in the breadth and heterogeneity of the problem space. Instead of authenticating identities, we may be obliged to authenticate any of a very wide range of attributes; and instead of the standard Dolev-Yao threat model, we have a matrix of threat models. This broadening of the problem space clearly indicates that ubiquitous, human-centric computing will make the problem of achieving trusted and trustworthy information services harder.

To structure our understanding of the broader problem space and to help organise discussion, we propose that the subject is factored into three sub-domains:

- **User Level.** This includes all the involvement of human users in achieving, violating or enabling the violation of security. It also includes the design of user interfaces. The user interfaces will themselves connect this level to the service level.
- **Service Level.** This level encompasses all applications, though our interest is primarily in security applications. The service level will make use of information resources and computing and processing capabilities offered by the infrastructure level.
- **Infrastructure Level.** This level contains the hardware present in the pervasive networks, the information resources, the communications architectures, the middle-ware and the software processing architectures.

Commonly, when such a layered factorisation is proposed, there is much debate and argument about the number of layers, the contents of the layers and so

forth. In this case such debate would miss the point: we do not prescribe this layered decomposition as canonical. It is merely a conceptual tool inspired by the fact that achieving trust in human-centric computing will not be possible without a careful consideration of the humans' role – thus the User layer to make this explicit. Furthermore, the delivery of ambient intelligence services will require a range of resources, communications and computing capabilities that will be globally standard and locally available – thus the need for an infrastructure layer. It should also be noted that the two-channel threat model, that we summarised above, implies two layers. Given that successful abstraction, we hope a layered decomposition of the problem will be a fruitful way to proceed.

3. The Issues That We Need to Understand

This final section contains an outline of some of the important issues that we feel need to be debated and understood, arranged according to the loose layering that we mentioned in the previous section.

3.1 User Level

1 The human's role in achieving trust needs to be clearly understood. The security requirements in the examples in Creese et al., 2003a, and the new modelling paradigm, in Creese et al., 2003b, derived from what assurances a human with a sound knowledge of information security would seek. In implementing them we made use of things a human would be willing and able to do to achieve such assurances. The important question here is this: to implement and achieve trustworthy interaction should the broad strategy be to minimise the human's role, or should it be assumed that humans can and should retain significant ownership of protecting their assets. Arguments for retaining the human's role include:

- the fact that people do care about their assets (and will continue to do so as they are digitised);
- people want to retain ownership of whatever they regard as precious; and
- the fact that people increasingly use electronic security mechanisms, especially PIN numbers.

Arguments for minimising the human's role are:

- the difficulty of designing trustworthy and effective human-computer interfaces;
- the general fact that most security violations involve irresponsible use or management by people;

- the fact that PIN numbers are frequently poorly managed and stolen; and
- the desirability of relieving the human user of tasks which might become very frequent and burdensome, or be necessary when the human is not in a position to do them.

This is clearly a fundamental question, but it may not be necessary to understand it as an exclusive choice.

- 2 With regard to the problem of enabling users to retain control of who and what they trust, this seems to define a whole service category of decision support tools. The tool will inevitably have some measure of control over the decisions that its owner makes. How well understood is the science of making such tools trustworthy? Clearly, for such tools to be effective their interface to the user must itself be effective. How well understood is the science of making the interface to such tools trustworthy?
- 3 Conversely, if the aim is to minimise the user's role in implementing security, then it should first be noted that this may make the problems of responsibility and liability harder.
- 4 The work in Creese et al., 2003a and Creese et al., 2003b lays the groundwork for understanding what a human security expert might require, and what is needed to establish authentication in pervasive environments. When using the concept of weakened Dolev-Yao channels as suggested in the second of these papers, it is important to investigate ways in which these channels can be realised both with, and more importantly without, human participation.

3.2 Service Level

- 1 An interesting issue arises whether tools and technologies for trust and security enable their users, or act on behalf of their human owners. This is whether to allow certain actions to go ahead in the presence of incomplete information. Of course the greater the importance of security, the less we will be inclined to do this.
- 2 Decision making about trust and security might be enabled, if it were possible to "quantify trust". This is hard, and any scheme will be prone to criticism, but the definition of trust as an "acceptable level of risk" might provide one basis for a way ahead. The notion of attribute authentication might provide an appropriate setting for trying to quantify trust and make decisions about acceptable risk. In any instance it is unlikely

that it will ever be possible to attain *complete* assurance about all the relevant attributes of the devices involved. Whether the user has control or not, a decision, based on incomplete information, will have to be taken about an acceptable level of risk.

- 3 However well we define our interface, we may still need to provide an underlying service which supports an appropriate authentication policy, depending on the context of use. Such a service or application would have to be able to tolerate heterogeneous user interaction, and still provide reliable security. However, would such a tool be considered trustworthy by users, given that it might be able to change the users command if it felt the user were mistaken? In addition, what data needs to be provided to such applications in order that they can provide the user with appropriate decision support regarding authentication policies?
- 4 Essential to pervasive computing will be the ease with which users can traverse distinct networks. This will require unparalleled levels of interoperability on the application level, heterogeneous devices and users will need to interact with a range of trust and security mechanisms. How can we enable such interoperability? Should we be subscribing to the top down approach of generating one standard, or ontology, to which all services subscribe, such as that being developed in the SWAD-Europe project¹? Is there a real alternative?
- 5 Where users and devices fail to authenticate should we provide services for broadcasting that fact, equivalent to revocation lists? If authentication means attribute authentication, then what would be the form and content of such “attribute revocation lists”?

3.3 Infrastructure Level

- 1 Section 2 provided a synopsis of the two-channel abstraction from Creese et al., 2003b, where the *E*-channel involves physical interaction and is critical to bootstrapping authentication. Typically the implementations of such channels will require things like physical contact, line-of-sight interaction, human intervention, and so forth. In any particular case the reliability of this channel will be crucial, and should be the subject of debate.
- 2 It is likely that any channel which assumes a weakened Dolev-Yao threat model will rely on (relative) contextual information about the processes using it. Therefore, we might regard such channels as a concise abstraction of the idea of context sensitivity.

- 3 Most authentication mechanisms currently rely on asymmetric encryption, which is computationally expensive, and requires larger keys – thus consuming more bandwidth. For pervasive computing, where many devices will be relatively weak in their computational and communication capabilities, it is highly desirable to find authentication mechanisms based on symmetric encryption, or one-way functions. Furthermore, the domination of asymmetric cryptography has, in part, been spurred by the need to implement identity authentication. Can attribute authentication provide the impetus for developing and deploying cheaper encryption techniques for authentication?
- 4 Will the trusted computing initiative² bring about solutions for supporting the authentication of device behaviours (these being some of the key attributes that will need authenticating)? If devices are reconfigurable in the field, then they are not necessarily the same as when they left the factory. What is the impact of this? Can we achieve “biometrics” for devices on which we could base our authentication of behaviours.
- 5 How can major global technology initiative, such as Grid computing³ and Semantic Web⁴ provide the information, computing and communication resources, to enable solutions to trust and trustworthiness in human-centric computing?
- 6 Finally, what can we do without infrastructure? Or, more precisely, what do we do fundamentally need, and what can we “create” spontaneously, on an on-demand basis?

4. Conclusions and Acknowledgements

In this position paper we have tried to organise and describe what we currently understand to be some of the major issues and problems that need to be understood, to achieve trust and security in human-centric computing. We look forward to healthy and active debate on the points noted above and we would like to thank Philip Robinson and Harald Vogt for organising the workshop and inviting us to make a submission.

Notes

1. www.w3.org/2001/sw/Europe/
2. www.trustedcomputing.org
3. www.gridforum.org
4. www.semanticweb.org

References

- Boyd, Colin and Mathuria, Anish (2003). *Protocols for Authentication and Key Establishment*. Springer-Verlag.
- Creese, S., Goldsmith, M. H., Roscoe, Bill, and Zakiuddin, Irfan (2003a). Authentication in pervasive computing. In Hutter, D. and Ullman, M., editors, *First International Conference on Security in Pervasive Computing*, Boppard. Springer LNCS.
- Creese, S. J., Goldsmith, M. H., Roscoe, A. W., and Zakiuddin, I. (2003b). The attacker in ubiquitous computing environments: Formalising the threat model. In *Formal Aspects of Security and Trust*, Pisa. Springer.
- Dolev, D. and Yao, A.C. (1983). On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2).
- Ryan, P.Y.A., with M.H. Goldsmith, S.A.Schneider, Lowe, G., and Roscoe, A.W. (2001). *The Modelling and Analysis of Security Protocols: the CSP Approach*. Addison-Wesley.