

Expressing Privacy Preferences in terms of Invasiveness

Patrik Osbakk and Nick Ryan

Computing Laboratory, University of Kent, Canterbury, CT2 7NF, UK
{pjo2,n.s.ryan}@kent.ac.uk

ABSTRACT

Dynamic context aware systems need highly flexible privacy protection mechanisms. We describe an extension to an existing RBAC-based mechanism that utilises a dynamic measure of invasiveness to determine whether contextual information should be released.

Keywords: Context, Privacy, RBAC, PIV, P3P, CCS

INTRODUCTION:

There are serious privacy issues that must be addressed before publicly acceptable context-aware systems can be deployed on a large scale [1, 2, 3]. Experimental systems have been developed for use in areas such as homes [4] or offices [5] where, arguably, restricted numbers of participants and closed environments limit privacy concerns. In more public areas, including shops, transport, gyms [6] or tourist attractions [7], systems will require privacy mechanisms with explicit policies to manage dynamic personal information as participants interact with, and move between, environments.

Previous work, including our privacy enhancing infrastructure [8], and others, e.g. [2], are based on classical privacy protection mechanisms such as Role Based Access Control (RBAC) [9]. Such mechanisms can protect a participant's privacy but, used alone, are insufficiently flexible to address the needs of dynamic, context-aware, systems. Here, we outline an extension to our privacy protection mechanism that permits a closer reflection of our everyday privacy decisions.

BACKGROUND

Privacy refers to a person's right to control the flow of information about them. Ownership of information is with the subject and its release should be at their discretion. Once released in an intelligible form, the subject ceases to have any control over their personal information. The trust placed in the recipient is therefore an important factor in determining whether or not to disclose information. Legislation that requires recipients to honour any agreement upon which information was released is perhaps the only possible protection mechanism beyond the point of release. In an imperfect world, it must be assumed that this can at best be only a discouragement to improper use.

The ideal level of privacy offered by a system equates to that which the participant enjoys offline. Compromising privacy should not be a default requirement in order to enjoy the benefits of context-aware environments. Where personal information is revealed, a privacy protection mechanism must permit sufficient customisation and flexibility to handle the multitude of situations occurring in the real world.

In earlier work [10] we examined a simple classification and clearance scheme (CCS) for privacy protection. Each context element was assigned a classification level indicating its sensitivity. Sites, services and other participants were assigned clearance values representing

a level of trust and determining which elements they could access. Whilst easy to understand, the approach does not scale to deal with large numbers of context elements or consumers. More recently, as part of a general privacy-enhancing infrastructure [8], we have sought to address this problem using Role Based Access Control (RBAC). A list of access controls, each referring to a specific element, allows a combination of read, write and history access to be granted.

Both approaches support requests by previously unknown consumers, provided that they can express their intended use with extended Platform for Privacy Preferences Project (P3P) policies [11]. Rulesets specifying how information must and must not be used, enable automatic decisions to be made on what clearance or roles are assigned to a consumer.

LIMITATIONS

Although the current privacy protection mechanism performs well it does have limitations. In static and clearly defined environments it is possible to setup RBAC to accurately represent privacy preferences because the context in which the preferences are set is known. In a dynamic environment privacy preferences may change with context. For example, activity information may be public when at work but not when at home.

Privacy preferences may also vary according to the context of the potential recipient. Conference attendees may be prepared to expose more information than usual but only to other attendees. Secondly, whilst occasional requests for particular context elements, e.g. location or velocity, may be permitted, repeated requests may represent an unacceptable level of surveillance.

Similarly, the risk in exposing an element may depend on which other elements have been exposed previously, either publicly, or to the same recipient. Finally, the potential impact of exposure may vary with the precision and reliability of the information. It may well be acceptable to reveal a rough location, say at town or region level, but not exact coordinates or street address.

PRIVACY INVASIVE VALUE

To address these limitations we introduce the concept of a *Privacy Invasive Value* (PIV). The idea behind PIV is that whilst any release of context information invades privacy, the extent of the invasion depends on many factors including what information is being released, to whom, when, under what circumstances, how often, etc. We retain "about what" and "to whom" as the primary factors for determining access, so the PIV concept can be used to extend the RBAC mechanism. Instead of simply assigning read and history access, the mechanism will also take account of the PIV. Similarly, in addition to roles and personal permissions, participants are assigned a maximum level of invasiveness.

By considering how invasive context consumers are allowed to be, a number of things can be achieved. Firstly the privacy invasion (PI) of a request for any single element may be capped. Secondly the aggregate PI of any individual consumer, and its rate of increase, can be limited. Thirdly, the aggregate PI may be made to decay over time in line with the temporal validity of component elements.

A key benefit is that PIVs and PIs need not be constants. The PIV of any context element and the PI of any participant can be modified at runtime according to pre-defined rules. In its simplest form the aggregate PIV is the sum of those of the released elements. In a more complex situation, the effective PIV of a requested element may be modified according to the current context of the owner and, if possible, that of the requestor. Similarly the PIV may be modified depending on previous actions, e.g. the release of a combination of location and velocity may be far more invasive than either alone. Finally reducing the precision of an element may incur a lower PIV, and so enable its release.

IMPLICATIONS

The introduction of PIV offers a gain in flexibility. Not only can many more privacy preferences be expressed, the extension also models real life decision more accurately. We can now make decisions based on information sensitivity in addition to “about what” and “to whom”. Also by introducing the PIV into our existing RBAC mechanism the system can be setup to mimic the behaviour of the previously evaluated CCS mechanism or the current RBAC, if desired.

The added functionality has a performance cost. The present mechanism can effectively cache access decisions as they need only be evaluated when preferences change. By basing access on dynamic information, the effectiveness of caching is significantly reduced. However, we anticipate that the additional load will be well within the increasing computational power of personal devices.

DISCUSSION

By using levels of invasiveness, rather than simply allowing or denying access, many new doors are opened. One of the attractive features is that the complexity of the mechanism is primarily determined by the preferences being described. This ensures that the mechanism is sufficiently simple for a novice user, yet flexible enough for an expert.

In our initial evaluation of the approach, PI is treated as a single value. We are also considering whether it might be beneficial to take a multi-dimensional view of PI. The balance between any benefits and increased complexity will need to be carefully evaluated.

The PIV extension may also benefit from explicit role activation, though this needs to be evaluated against the extra knowledge needed by the context consumers.

At this stage it is still unclear what the full potential of the PIV approach is, but we believe that it is sufficiently interesting to be worthy of further investigation.

RELATED WORK

Whereas we have extended the protection mechanism by increasing the responsibility of the permissions others have extended the RBAC model [12, 13] itself. We have found, though, that even the simple RBAC in our infrastructure is considerable more difficult to setup than the CCS previously evaluated. By retaining the simplest RBAC model possible and introducing the concept of PIV we can address the current limitations without further complicating the privacy protection mechanism. The added flexibility may even make administration easier by reducing the number of roles required.

CONCLUSION

We have here presented an extension that has been developed to address the limitations of our current RBAC-based privacy protection mechanism. We suggest that by describing privacy preferences in terms of privacy invasiveness, sufficient flexibility is gained to alleviate these limitations and to allow further control. Since the PIV concept represents our day to day privacy decisions more closely than the current RBAC's allow/deny states, the effectiveness of the protection mechanism is improved. Its simplicity, and the benefits found so far, makes the approach sufficiently interesting to warrant further research into its potential.

REFERENCES

- [1] Langheinrich, M. (2002). “Privacy Invasions in Ubiquitous Computing”. *Workshop on Socially-informed Design of Privacy-enhancing Solutions*, UbiComp 2002, Göteborg, Sweden.
- [2] Ebling, M., G. Hunt, et al. (2001). “Issues for Context Services for Pervasive Computing”. *Advanced Topic Workshop Middleware for Mobile Computing*, Heidelberg, Germany.
- [3] Brown, P. J. & G. Jones. “Context-awareness and privacy: an inevitable clash?”. [Version #1]. Last accessed: 26/02/2004. http://www.dcs.ex.ac.uk/~pjbrown/papers/ieee_privacy.pdf.
- [4] Kidd, C. D., R. Orr, et al. (1999). “The Aware Home: A Living Laboratory for Ubiquitous Computing Research”. *2nd International Workshop on Cooperative Buildings - CoBuild'99*, Pittsburgh, USA.
- [5] Harter, A., A. Hopper, et al. (2002). “The Anatomy of Context-Aware Applications”. *Wireless Networks* 8(2-3): 187-197.
- [6] McCarthy, J. F. (1998). “MusicFX: An Arbiter of Group Preferences”. *AAAI Spring Symposium on Intelligent Environments*, Palo Alto, USA.
- [7] Cheverst, K., N. Davies, et al. (2000). “Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences”. *Conf. on Human Factors and Computing Systems*, The Hague, Netherlands.
- [8] Osbakk, P. and N. Ryan (2003). “A Privacy Enhancing Infrastructure for Context-Awareness”. *1st UK-UbiNet Workshop*, London, UK.
- [9] Sandhu, R. S., E. J. Coyne, et al. (1996). “Role-Based Access Control Models”. *IEEE Computer* 29(2): 38-47.
- [10] Osbakk, P. and N. Ryan (2002). “Context, CC/PP, and P3P”. *UbiComp 2002 Adjunct Proceedings*, Göteborg, Sweden.
- [11] “The Platform for Privacy Preferences 1.0 (P3P1.0) Specification”. [W3C Recommendation 16 April 2002]. *World Wide Web Consortium (W3C)*.
- [12] Covington, M. J., W. Long, et al. (2001). “Securing Context-Aware Applications Using Environment Roles”. *Proc. 6th ACM Symposium on Access control models and technologies*, Chantilly, Virginia, USA.
- [13] Zhang, G. and M. Parashar (2004). “Context-aware Dynamic Access Control for Pervasive Applications”. *Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS 2004)*, San Diego, CA, USA.