Security versus Usability in a Ubiquitous Environment

Walter Zimmer, Thomas Messerer, Rainer Steffen, Rudi Knorr

Fraunhofer Institute for Communication Systems
Hansastrasse 32
80686 Munich
{walter.zimmer,thomas.messerer,rainer.steffen,rudi.knorr}@esk.fraunhofer.de

Abstract: Security and usability are difficult to achieve simultaneously. As upcoming ubiquitous scenarios focus on usability, new security solutions are in high demand. We are following two approaches: "Inheritable Trust" which is an intuitive way of expressing and enforcing trust in a ubiquitous environment, and "Dynamic Trust Networks" which perhaps provide a mechanism for easy trust management in a longer timeframe.

1 Introduction

For the relationship between security and usability, there is a natural tradeoff: either a system is very secure and hard to use, or it is easy to use and not very secure. This results from the fact that investments have to be made to achieve security, and those investments have to be applied and administered, which impacts ease of use.

In some scenarios however, ease of use is valued more than the value of the objects which are to be protected. This occurs especially in future ubiquitous scenarios, where usability is the key factor to user acceptance and often the objects to be secured are of small value. For example, an electronic business card or the access to an electronic blind management system don't represent a high value, nevertheless the access certainly wants to be controlled. Therefore, we believe that there is demand for security architectures which accept a decrease in security in exchange for the advantage of being easy to use. In particular, we see this scenario happening in the office environment, where different employees have different rights, and occasional guests also require access to the provided resources.

A powerful approach for solving security issues is delivered by the "trust" paradigm. The different trust relationships between users are the primary cause why they want to secure or share resources at all. Therefore, it makes sense to evaluate the inclusion of individual trust into a security model.

Our approaches to apply the trust paradigm to a security architecture for ubiquitous computing are described in chapters 2 and 3, with chapter 4 containing the relation to other work and chapter 5 provides the final conclusion and outlook.

2 Inheritable Trust

The first concept we decided to elaborate further is what we named “Inheritable Trust”. The basic idea is a delegation system [Ka01] which enables users to express and enforce the trust they have in others by means of digital trust tokens. Each trust token incorporates specific rights. Tokens can be passed to other users, possibly with reduced rights. This key feature is the reason for the reference of the name to inheritage, as it resembles the main properties of the inheritage process: one inherits specific resources, and at a later time a subset of them can be passed to others again.

Our concept allows to distribute trust completely anonymous, without the (communication) overhead of a public key infrastructure (PKI). The advantages are that the user is not confronted with authentication, tokens can be passed on even if communication with the system itself is currently not possible, and revokation of issued tokens is possible by depositing revokation tokens at the service provider.

2.1 Inheritable Trust for Permissions

The obvious usage of the token architecture is to encode permissions into them. For example, for printing service tokens, this would include information about the allowed printing devices, expiration time, features like color, maximum number of pages per printjob, allowed time of day etc.

Every user can reduce the permissions before passing them to trustees. Cryptographic measures have been taken to protect the authenticity of the token, therefore no one can issue tokens by himself (forging).

This approach resembles a ticketing system, with the difference that our “tickets” are restrictable and copyable.

2.2 Inheritable Trust for Trust Values

The concept for Inheritable Trust shown above can be enhanced over its explicit nature, where every access to a resource must be encoded and manged. Lets consider the case that only one value is encoded into the token, which is not interpreted as a hard permission for a resource, but as a trust level. Therefore, the presentation of a trust token proves that the presenter is trustworthy to the extent encoded into the token. If someone is trusted, he can express his trust in another user by issuing him a token with a (possibly reduced) trust value.

This approach opens up new possibilities: Now, it is up to the service to decide what rights it grants to a user with a certain trust value. Also, it is possible to incorporate the actual context of the user into the calculation of the permission function. For example, a print service might be configured in such a way that only users with a trust level greater than 90% are allowed to print on all printers even out of office hours, in contrast to users with a trust level below 10%, which are not allowed to print at all.

Configuring these impairments manually would be time consuming for the user. By passing and configuring only the trust value, this burden is shifted to the adminsitrator, who can define the corresponding rights globally in accordance with the security policy. In exchange, the administrator is not concerned with user management.

2.3 Limitations of Inheritable Trust

The ease of use for Inheritable Trust certainly comes at a price: One limitation is, that physical copying of tokens is in fact not preventable. If all information which grant specific rights is encoded into a sequence of bits, these rights can be given to someone else by copying this sequence. The only way to deal with this situation is to make it as difficult as possible, in order to make the necessary investment needed to copy the bitstring much larger than the value of the encoded rights. If it takes twenty minutes to illegally get a permission to print which was initially valid for half an hour, the remaining ten minutes are not really worth the investment. Following the natural path of trust and asking someone for the permission might be more effective.

3 Dynamic Trust Networks

The second approach we are pursuing is the management of permissions through a network of distributed trust issuers, which we named „Dynamic Trust Network“ (DTN). The primary goal of DTNs is the simplification of the administration of a multiplicity of networked devices and the increased usability of those devices, through the application of a reputation system based „distributed trust“ mechanism as described in [GM02].

A DTN is a network of trusting nodes, with trust relationships as found in societies or (virtual) communities. One can trust an unknown person, to a certain extent, if this person is recommended by other persons who are already trusted. This can be used, for example, by a service provider to determine if a certain user is allowed to use a specific service. We consider trust as transitive like it is described in [Gr02], because the benefit of experienc from trusted nodes is substantial. The network behaves more dynamic and the risk for the single node is lower, compared to an approach where every node acts on his own responsibility.

The basic method for determining trust values is straightforward: After the arrival of a request, the service provider acquires the trust values respective to the requester from the DTN. With the results and his own trust values about the contributing nodes he can calculate a cummulative trust value and decide whether he will grant the access. If he decides to allow the access, he will monitor the behaviour of the accessing node and adjust his stored trust values concerning this node and the nodes which delivered false information about it.

The adaption of the node’s trust value leads to a dynamic trust relation, which can be enforced and weakened over time. If a new, unknown node behaves properly, its trust relations will grow, and it will become a contributing part of the trust network.

If the aquired trust value of a node is lower than the trust threshold for the particular service, then the service provider is forced to take a risk. To accept a limited risk is essential for the growth of a trust network, otherwise all the nodes will remain isolated.

Every Node in the DTN processes the same algorithm, and through the cooperation a self-organising trust network emerges. Long lasting and frequently used connections between devices become strong, and constitute the basis of the trust network. New nodes can be integrated in an already established DTN quickly, and cheaters will be eliminated in the long run because their trust relations weaken. Beyond this, in the case of cheating, the trust can switch to distrust, in order to warn other nodes.

For an administrator the DTN is an essential alleviation because he isn't forced to configure every single device or user in the network. The system shall ease the use of security mechanisms for the regular user by advising him in security affairs. After a training period it is possible that the software can make decisions for the user.

For a DTN it should be possible to act and grow autonomous without human interaction. The system shall be completely self initialising, but it will converge faster if there are some devices already trusting among themselves.

3.1 Limitations and fields of further research

One open issue is the occurence of missing determinism: a service user does not know whether his trust level is high enough to exceed the required trust threshold, which might both be unknown. Users demand determinism, therefore this is a field which has to be investigated further.

In many documents concerning distributed trust networks it is assumed that the trust relations are already established. But this is not the case in real life. A part of our future research will be to evaluate solutions for the initialisation of trust networks.

Feedback is essential for the dynamics of a trust network. In systems with user interaction, an approach might be to ask the user at the end of a transaction to evaluate if the invested trust was worthwhile. In the long run, however, it will be annoying for the user to rate every performed transaction. Therefore, to automate this process is a future field of interest.

To get an ample amount of trust values from the DTN it is important to maintain a lot of connections. In highly dynamic networks it's not always guaranteed that there are enough nodes within reach. With the trend towards rising connectivity however, this will eventually become a less limiting factor.

4 Related Work

Trust and distributed trust have become widely accepted research fields. Information about delegation systems can be found in [Ka01]. Secure ticketing systems are described in [Fu99]. All those depend on a PKI, however. The DTN-approach is based on trust relations in real life, as described in [Ko99], distributed trust networks are described in [GHP03] and [GH04].

5 Conclusion and Outlook

We are currently working on two approaches to ease the administration of user rights: *Inheritable Trust* which consists of an architecture to express and enforce trust between users, and *Dynamic Trust Networks*, which describes an approach to manage user permissions by means of distributed trust.

Inheritable Trust can be used for scenarios where the value of the objects is limited, as physical copying of tokens cannot be prevented. However, the ease of use makes it useful for applications which manage permissions for goods with low value or for a limited time. We have built a prototype based on bluetooth which implements the basic features of the system for demonstration purposes. After experience with this system has been gathered, the further path of this research track will be determined.

Dynamic Trust Networks is seen by us as an interesting concept to ease administration and application of user rights which still has some unsolved problems remaining.

References

[GHP03] Golbeck, J.; Hendler, J.; Parsia, B.: Trust Networks on the Semantic Web. Twelfth International World Wide Web Conference (WWW2003), Budapest, Hungary, 2003

[GH04] Golbeck, J.; Hendler, J.: Inferring reputation on the semantic web. Submitted to WWW'04

[GM02] Goecks, J.; Mynatt, E.: Enabling privacy management in ubiquitous computing environments through trust and reputation systems. In Proc. Conference on Computer Supported Cooperative Work, CSCW 2002, Workshop on Privacy in Digital Environments: Empowering Users, New Orleans 2002.

[Gr02] Gray E.; O'Connell P.; Jensen C. D.; Weber S.; Seigneur, J. M.; Chen, Y.: Towards a Framework for Assessing Trust-Based Admission Control in Collaborative Ad Hoc Applications. Technical Report, Trinity College Dublin, 2002

[Ka01] Kagal, L.; Cost, S.; Finin, T.; Peng, Y.: A Framework for Distributed Trust Management. In proceedings of IJCAI-01, Workshop on Autonomy, Delegation and Control, 2001

[Ko99] Kollock, P.: The Production of Trust in Online Markets. In Advances in Group Processes (Vol. 16), E. J. Lawler, M. Macy, S. Thyne, and H. A. Walker (eds), Greenwich, CT: JAI Press, 1999. http://www.sscnet.ucla.edu/soc/faculty/kollock/papers/online_trust.htm

[Fu99] Fujimura, K.; Kuno, H.; Tereda, M.; Matsuyama, K.; Mizuno, Y.; Sekine J.: Digital-Ticket-Controlled Digital Ticket Circulation. In Proc. of 8th USENIX Security Symposium, Washington D.C., 1999