How much is “enough”? Risk in Trust-Based Access Control

Nathan Dimmock*
University of Cambridge Computer Laboratory
JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom
Nathan.Dimmock@cl.cam.ac.uk

1 Introduction

Weiser’s vision of ubiquitous computing[8] is of a massive, and largely unseen, population of cooperating networked devices. This new computing environment will lead to new security challenges far removed from today’s fortress mentality of firewalls and strictly controlled fixed network infrastructure. Cooperation requires trust between participants. In a connected world, trust may be obtained by deferring the security decision to a trusted-third party, such as in OASIS[2]. In a disconnected environment, an entity may carry tokens to assert their trustworthiness[3], but both of these systems require an existing trust infrastructure.

Such a trust infrastructure might be SPKI[5], but as a hierarchical infrastructure, the requirement to have an ultimately trusted (root) certificate makes it too inflexible for our ad-hoc collaboration environment. A PGP “web of trust”[5] infrastructure allows non-hierarchical trust relations similar to those formed in a human community, but PGP supports only a <name,key> mapping; it says nothing about the access control (authorisation) rights held by the principal. A further problem with SPKI certificates in this respect is that authorisation is part of the certificate and so all privileges must be determined when the certificate is issued. This restricts collaboration to those interactions which have been foreseen. In a poorly or vulnerably connected environment, it may be that the cost of verifying the authorisation is greater than the risk of allowing the operation to proceed and if a web-of-trust for authorisations is allowed, a mechanism is needed for resolving conflicting recommendations.

2 Trust, Risk and Privilege

Trust-based access control is the idea of using a model of human notions of trust and community as the basis for assigning privileges. However, although we may make general statements such as “I trust Brian”, or even quantitative ones such as “I have 90% trust in Brian”, these are not particularly meaningful. How much trust do I need to have in my fellow commuter before, for instance, I let him borrow my telephone?

In human communities, the amount of trust required seems to depend on the nature of the action, or more specifically, the *risk* involved. The likelihood that a person will lend a particular amount of money to a friend depends on the amount involved, how good a friend they are, and perhaps other factors such as past experiences of lending money to this person. Effectively, the person is asking: *How likely is this person to pay me back and, if they do not, does it matter?*

What is Risk?

The safety-critical programming field defines risk in terms of the likelihood and severity of an accident [6]. The insurance industry uses the definition that (economic) risk derives from variation from the expected outcome and one measure of this is the standard deviation of all the possible outcomes [1].

3 A Risk Model for TBAC

When a trust-based access control system grants privileges to a principal it is with the expectation that they will use them in a particular manner (for example, to update patient records according to what treatment was administered to them) but there is also the possibility that the principal will deviate from this expected behaviour and the combined likelihood and severity of that variation is the risk of granting them a privilege.

In the business world, risk is seen not only as the potential economic loss but also the source of one’s profit and this leads us to use a probabilistic cost-benefit analysis to determine the level of privilege that should be assigned to a principal. Access control decisions are determined on a case-by-case basis: each principal, before making a decision on whether to interact with another principal, analyses all the possible outcomes from that interaction and assesses the possible variation from the expected one. This is done by calculating

*This research is part of the EU-funded SECURE project, IST-2001-32486. Special thanks to Brian Shand, Jean Bacon and the anonymous reviewers for their helpful comments.

the likelihood and maximum potential cost or benefit (in financial terms) of each outcome. It is very likely that some outcomes will actually be distributed over a space such that there is a range of potential costs with corresponding probabilities. Therefore each outcome will be represented by a "cost-PDF", that is, a probability density function with cost on the x-axis.

Deriving Costs and Probabilities for an Outcome

The cost of potential outcomes may also be extrapolated from any historical data, otherwise it may be estimated or arbitrarily assigned, depending on the application[7]. A precise cost is not always needed — simply ordering the outcomes may be enough to show the correct course of action. Theoretically the probability of an outcome can also be naïvely derived by examining the results of previous interactions with this principal. Unfortunately this extrapolation may be inaccurate due to too few previous interactions, or a recent change in circumstances. We therefore invoke our trust-model[4] to evaluate all the known information about a principal (including the relevant previous interactions) to determine how *trustworthy* they are considered. A trustworthy principal is one we believe is likely to act in a positive manner towards us while an untrustworthy principal is likely to choose a course of action that is detrimental to us.

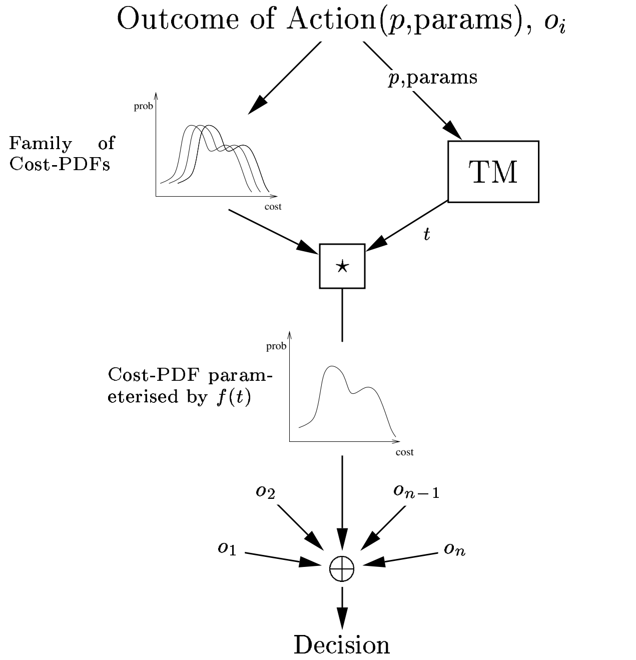

Figure 1. The interface between trust and risk.

Suppose a principal, p, wishes to enter into a parametrised interaction with a user (fig. 1). For each possible outcome, the user has a parameterised cost-PDF (that is, a family of cost-PDFs) which represent the range of possible costs and benefits that may be incurred by the user should this be the actual outcome. The user passes the name of the principal and any parameters associated with the action to the trust model which returns some trust-information, t. The $\star$ function in fig. 1 then uses t to determine the values of the parameters of the cost-PDF and therefore selects which of the family of cost-PDFs is most appropriate in this situation. For example, suppose t is a vector consisting of t_1 and t_2, indicating a judgement of how reliable p is in this context. For this particular outcome, the cost-PDF might be a Gaussian distribution with mean determined by t_1 and variance determined by $t_1 \oplus t_2$.

Once a cost-PDF for each outcome has been determined, the cost-PDFs for all the possible outcomes are combined and analysed according to the user's security policy; a decision is made on whether to allow this action to be taken. This "answer set" may include more than two values, for example allowing the indication of too little information to make a decision ("not sure"). A key benefit of this approach is that uncertainty (in costs and probabilities) is measured explicitly and carried through to the end of the calculation. The security policy may be based on expected cost, variance and/or maximum possible loss, depending on the type of risk analysis called for by the application. For example, if one of the principals is a company there may be a loss value that they cannot risk incurring at all (the value of the company). Non-linearity in mapping financial loss to perceived cost is quite possible and represents the principal's *risk sensitivity.*

4 Conclusions and Future Work

A model for using explicit risk analysis to determine how much trust is required to assign a particular privilege in trust-based access control has been outlined. Further work includes validation of the model, investigating how continuous probability density functions may be reasoned about in resource constrained environments and the creation of a general purpose policy language to allow the specification of policies.

References

[1] J. F. Anderson and R. L. Brown. *Risk and Insurance.* Number 1-21-00 in Study Notes. Society of Actuaries, 2000.

[2] J. Bacon, K. Moody, and W. Yao. Access control and trust in the use of widely distributed services. In *Proc. Middleware 2001, LNCS 2218*, pages 295–310, 2001.

[3] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proc. IEEE Conference on Security and Privacy.* AT&T, May 1996.

[4] V. Cahill and J.-M. Signeur. Secure Environments for Collaboration among Ubiquitous Roaming Entities. Website: `http://secure.dsg.cs.tcd.ie`.

[5] IETF RFC-2963. *SPKI Certificate Theory*, 1999.

[6] N. G. Leveson. *Safeware: System Safety and Computers*, chapter 9. Addison-Wesley, 1995.

[7] B. Shand, N. Dimmock, and J. Bacon. Trust for transparent, ubiquitous collaboration. In *IEEE Conf. on Pervasive Computing and Communications*, 2003.

[8] M. Weiser. The computer for the 21st century. *Scientific American*, pages 94–104, Sept 1991.