Trust management for inter-enterprise collaborations

Sini Ruohomaa

University of Helsinki, Finland
Sini.Ruohomaa@cs.helsinki.fi,
WWW home page: http://cinco.cs.helsinki.fi/

Abstract. Enterprise computing is moving towards more open, collaborative systems, which involves issues in technical, semantic and pragmatic interoperability. Trust management focuses on pragmatic problems of whether two enterprises trust each other enough to want to collaborate in the first place, or whether they wish to end an existing endeavour due to perceived risks outweighing the trust between the partners. This paper presents my PhD research on the TuBE system for automated trust management and describes how trust and reputation are modeled in the system.

1 Introduction

Enterprise computing is currently moving towards more open, collaborative systems. Collaboration allows organizations to focus their resources on a few key fields of expertise, while continuing to provide broader services for customers. It also enables small and medium enterprises to compete in fields dominated by large corporations by joining together to gain more influence than they would have separately. The enterprises maintain their own independence during the collaboration, and make local decisions based on the enterprise policy.

There are technical, social and legal challenges in the way of this development, however. Information systems that should be connected are incompatible both technically and semantically, and systems integration is expensive and time consuming. Typically integration is also done according to the larger partner's system, and when partners change, the process must be repeated. Trust between new partners cannot be formed the same way as before when the entire process of setting up a collaboration is accelerated from several months, even years of negotiations to a few days or less. Finally, collaborators may be from different continents, with different cultural norms and legistlation, and intensive legal consultation for contract negotiations is costly. The environment is also fully distributed; there are no trusted third parties that would make sure all parties use a particular information system, or give consultation on whom to trust, or solve contractual disputes.

The value added by a collaborative business network must be balanced with its costs. It therefore becomes essential that joining a business network is made

efficient, despite the various difficulties. Connecting to new information systems should be straightforward, achieved through interoperability rather than manual integration, and locating new partners should be made automatic when the business need is defined. Trust management should be given simple and automated support which implements the enterprise's local policies and takes advantage of all the relevant information sources. Contract negotiations should preferrably be automated and based on machine-interpretable templates, which places new requirements on legal frameworks in questions of responsibility and contract validity.

Interoperability between information systems can be achieved by middleware instead of repeated integration: service-oriented computing suggests connecting the access interface of legacy systems into platform-independent service wrappers, which can in turn communicate in accordance to for example the document-based standard SOAP protocol used for Web Services [1]. Client code for accessing the service can be generated for various platforms once a standard Web Service description has been produced. and lower-level infrastructure will take care of technicalities such as reliable message passing or transaction management.

Defining business needs and locating partners to fulfil those needs requires support of more advanced middleware. The Pilarcos middleware services [2, 3] provide repositories for the public storing of service offers that fulfil a particular service type. These service types can be used to piece together business network models that correspond with a potential collaborator's needs. The business network models are matched with the available service offers and checked for interoperability by a populator service that can be provided by a third party. The middleware also supports automated contract negotiations between the parties, and monitoring fulfilment of the contract clauses and the business network's state.

No business network is feasible without an acceptable level of mutual trust between partners. Trust is required to balance risks, as an effective collaboration always contains an element of dependence and vulnerability. My research aims to provide support for trust decisions as a part of the existing Pilarocs business-to-business collaboration middleware. Agents represent their autonomous organizations in business networks, and need to routinely make decisions on whether to join new collaborations or withdraw from existing networks due to an increase in the percieved risk. If the risk is found untolerable, either the incentives must be increased or the risks limited.

The model for trust and reputation developed in the Trust Based on Evidence (TuBE) project explicitly connects risk analysis and business importance of a collaboration with the reputation of the other participants as the basic building blocks of a trust decision. Reputation information is built and constantly updated with new experience, gained both by participating in a collaboration and from other agents reporting their own experiences.

Section 2 presents the research question and gives an overview of my PhD research work. Section 3 briefly presents one result of the work so far, the trust

model. Section 4 elaborates on how a central factor of trust, reputation, is represented and built from experience. Section 5 concludes with some plans for future work.

2 Building a trust management system

My PhD research focuses on the upkeep of this local and external reputation information and its use to support choosing partners and guiding how a partnership evolves. Supporting research questions fall into two categories: how to organize the internal management of experience and reputation information, and how to use the gathered information to produce a trust decision.

Trust management is a broadly used term that defines a highly diverse set of research approaches, from inspiring trust in human minds [4] to managing certification webs of trust in the public key infrastructure [5]. My research builds on work that begun with trust expressed through certification and managed by policy [6]. The major development step for trust management has been the introduction of local learning through experience, and through that the independence from trusted third party certifiers. Reputation networks now share experience information that can be evaluated on a case-by-case basis, and information sources are also continuously evaluated for their credibility, that is, their trustworthiness as recommenders. Explicit evaluation of risks and benefits in trust decisions has a strong backing in research on the nature of trust [7], but it is a new enhancement in systems making computational trust decisions [8].

The currently existing trust management and reputation systems are mostly directed towards relatively low-risk environments focused on short-lived collaborations, such as electronical marketplaces or file sharing networks [9, 10]. My application environment involves services provided by organizations instead of private people, and the collaborations can last for months or even years. A single trust decision on whether to do business with another actor or not is therefore insufficient, as it should not be assumed that the trust relationship remains static through the collaboration. Business networks are governed by a negotiated eContract, which can define compensatory actions in case of breaches of contract or completeq withdrawal from the collaboration, which must also be taken into account when deciding whether to continue in a business network after e.g. a partner's reputation plummets. The Pilarcos middleware will provide valuable information that makes the TuBE trust management system capable of considerably more useful decisions than a standalone system that depends on the user feeding it all this information explicitly.

My PhD research work is divided into four phases:

1. Surveying the field and defining a model for trust,
2. designing the TuBE trust management system,
3. evaluating the design,
4. implementing the system and experimenting on it.

The goal of the first phase has been to learn to know the research field, to identify existing solutions and other useful results, and to build a model for trust

to use in the Trust Based on Evidence (TuBE) project. The work from this phase has produced two surveys [9, 10], and a description of the trust model used [9].

The goal of the second phase has been to design the TuBE trust management system and to determine which subsystem to implement as part of my PhD research. The system architecture was described during the TuBE project and published later in a joint paper [11]. After the one-year project was completed, I focused on how to integrate the trust management system into the existing Pilarcos architecture [12] and continued refining the architecture by information representation and algorithm aspects. I have decided to focus implementation efforts to the experience and reputation handling subsystem, and partially implement other parts of the system so as to enable experimentation.

The goal of the third phase is to evaluate the trust management system design. I have studied the various evaluation approaches available, and chosen to focus on examining the system's resilience against threat scenarios; I am currently involved in joint work identifying threats related to middleware supporting business-to-business collaboration, in which I am focusing on trust decision support and reputation management. This work will lay the basis on building relevant threat scenarios to use in the system design evaluation. I will also research possible performance bottlenecks in the system design before implementation. Work in this phase is ongoing.

The goal of the fourth phase will be to implement the relevant parts of the trust management system and to experiment on it. I consider two evaluation approaches most relevant for this phase: measuring the system's perforamnce to ensure that the introduced overhead is not inordinate, and analysing the system's actual response to scenarios implemented as input data sets. This scenario set will extend the threat scenarios used in the third phase.

3 A model for trust

Building trust through social activities requires continuous large investments in both time and people to first build trust relationships and monitor their development. In an open collaboration environment, where numerous trust relationships must be formed and upkept simultaneously, automation becomes the essential means to keep the costs of collaborating in check. This requires a computational representation of trust.

In the TuBE trust model, trust is seen as *the extent to which one party is willing to participate in a given action with a given partner in a given situation, considering the risks and incentives involved.* Similar viewpoints are referred to as trusting intentions by McKnight and Chervany [13] and situational trust by Jøsang et al. [14]. Our trust management system produces context-dependent and dynamic trust decisions, and estimations of the actual trustworthiness of a peer are simply a means to an end.

The focus of a business trying to decide on participation is generally not in managing trust, but managing risk. The connection between trust and risk has been widely noted [7, 8]. Indeed, the rather natural view of trustworthiness as

a subjective probability of successful collaboration (e.g. Gambetta [15]) clearly makes its measurement a tool for risk prevention.

A trust decision is a function of 7 parameters: *trustor*, *trustee*, *action*, *reputation*, *risk*, *importance* and *context*. It produces a decision with three possible values: *allow*, *deny* or *unsure*. In practice, "unsure" is very similar to "deny" in the short term, as the decision must be passed to a higher level, generally a human, and the final result may take considerably longer.

The trustor denotes the party making the trust decision; that is, as two peers rely on different local information in their decision, it is clear that the results also differ. The trustee is the source or target of the triggering message bound in or out, respectively. The action represents an ordered set of messages with content, and has a decision point determined in that set by when a risk-relevant commitment is being made.

The TuBE trust model elaborates the traditional factor basis of trustor, trustee and action by reputation, risk, importance and context factors. From these, a situational risk estimate and a representation of the risk tolerance for the particular situation are generated dynamically. A decision is produced from comparing the two. The choices for factors beyond the basic triple differ from one model to another [16]. In addition, terminology is mixed, so "context", for example, has several different meanings.

Reputation, as used in the TuBE model, is the measure of a peer's trustworthiness. It is not bound to a global agreement based solely on public information, which has been the traditional approach in standalone reputation systems. Instead, every trustor has its own view of what the reputation of a particular trustee is. To build this view, a trustor combines its own experiences with experiences reported by other peers, considering the credibility and information content of all statements. Such a combination is considered for example by Abdul-Rahman and Hailes [17].

The risk parameter contains a tactical risk estimate of the action. It consists of a set of identified risks and potential benefits to different assets, such as money, security, customer satisfaction and intellectual property. These risk and benefit estimates are speculations of the effect of a positive decision. There are trust and risk models which only consider two possible actions by the trustee: cooperation and defection. However, the world is not black and white: Not only can the trustee partially defect in multiple different ways, e.g. by not delivering goods or violating a contract clause, it can also defect at different degrees. For example, the costs of a product being three seconds late are often considerably different from it being three months late, and the quality of a provided service can be anything between excellent and abysmal. The severity ranges of each risk and the weight ranges of each benefit are considered and stored per asset.

The risk parameter depends solely on the action to be performed. However, the subjective probability that each risk manifests depends on the trustee's reputation. The risk analysis is completed by combining the structure of the risks and benefits with a set of probability distributions for them, derived from the trustee's reputation. The resulting estimate is a set of cost-benefit probabil-

ity distributions, one for each asset. Cost-benefit estimates have long been a part of trust models; e.g. Marsh considered several business value concepts in his work [18]. SECURE has applied continuous cost-benefit probability density functions for risk analysis, which squeezes all assets into one result function [8].

The importance parameter brings a strategic counterpart to the tactical evaluation contained in the risk parameter. While the risk analysis is directly dependent on what the trustee may do, the importance parameter directs what should be done independently of the trustee's possible behaviour in the future. This factor guides the tolerance of risk, which is represented by a set of constraints for the risk estimate. It represents considerations such as the cost of denying an action, or the benefit of giving great service to certain peers even when it is rather risky. For example, if denying service violates a contract, compensation is needed and the trustor's own reputation may suffer. In Poblano [19], importance is a way for the user to override a tactical risk-based decision, so it is similarly a strategic tool.

The context parameter represents a set of temporary adjustments to make to other factors. These adjustments either apply to risk or its tolerance, and their scope may be limited to a particular group of trustees, actions or their parameters. Context changes come from three sources: the internal state of the peer's system, the state of the peer's business and the state of the business network the peer is involved with. Context-aware systems in this sense seem rare [16]. On the other hand, items such as the *reciprocity* of trust, as discussed by Marsh [18], can be expressed as a contextual adjustment to the importance factor.

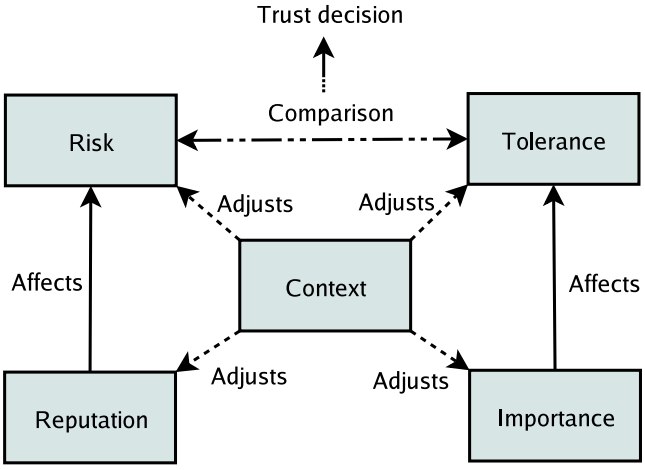

Fig. 1. Factors used to form a trust decision.

A trust decision is built on the estimated risk an action inflicts on the assets the trust management system must protect, compared with the risk tolerance for it. These depend on the reputation of the trustee, the importance of the action, and any adjustments needed in the context of the situation at hand. In the following, we will describe how these factors are represented, and outline the processes producing the information. The factors' interdependencies are depicted in Figure 1.

4 Representing experience and reputation

One major problem to all kinds of collaboration, but particularly eCollaboration, is the disruptive, opportunistic behaviour of partners and the difficulty of predicting it. Contract breach management is directed towards stopping misbehaviour within one collaboration, but more long-lived consequences are also needed to discourage unwanted behaviour. Some violations are forbidden by actively enforced legistlation, which is a relatively strong deterrent, although international collaborations cause challenges in knowing what laws apply in a given context. Social controls are needed to make unwanted but legally unpunishable behaviour sufficiently risky that acceptable behaviour becomes a more popular strategy.

Reputation is a strong social control mechanism in human collaboration, and forms a strong basis both for determining trust between actors who do not have much experince of working together as well as for updating perceptions with problems and successes others have experienced. A technical representation of reputation has improved business in specialized electronic marketplaces as well [20], and a generalized model of reputation can be expected to provide similar social controls in a business-to-business collaboration environment.

The group of potential, current and earlier partners is extremely large and dynamic, and these partners offer a wide and changing range of services. This means that whatever available first-hand experience one trustor has is going to cover only a small subset of actors in the network. Third-party experience is needed to support the local information, and to cover as many actors as possible, the information sources cannot be limited to a small ultimately trusted group. On the other hand, a wide array of sources brings along the problem of false or misleading reputation information, which needs to be addressed by a local critical evaluation of the credibility of the shared information [10].

Reputation is built from experiences. Experiences describe the outcomes of completed actions, and are tied to assets similarly to risk. Where risk estimates represent the probability of particular outcomes, experiences describe the actual result.

While the assets considered interesting in particular enterprises may vary, we believe that all organizations will have use of a handful of basic assets. Defining a set of standard assets also has the benefit that experiences based on these assets can be used across systems with less information lost due to unmatched

or unclearly defined assets, which may have a clear and valuable role in one organization, but are not understood in the other.

We have settled on a set of four standard assets: monetary, reputation, control and fulfilment. The monetary asset represents money and other things in the organization that have a well-defined monetary value. The reputation asset represents the trustor's good reputation in the eyes of the environment it operates in. The control asset is a joint representation for the trustor's security, privacy and general local self-protection assets. The fulfilment asset represents the trustor's expectations of the trustee's participation in the action, such as the quality of the service the trustee may be providing or its efficiency in fulfilling its own end of the agreement.

Our model for reputation represents outcomes as five categories of effects on each asset: large negative effect, slight negative effect, no change, slight positive effect, and large positive effect. The exact semantics of these effects depend on the enterprise: for example, the size of a typical business order affects whether a loss of a hundred euros is considered a large or a slight negative effect on the monetary asset for that particular action. As these outcomes are determined for all assets they can be determined for, they together form one item of experience. Some actions may also have an undefined or unknown effect on a particular asset, in which case a special "no outcome value" is recorded.

To aggregate experiences into a reputation view, the number of different outcomes for each asset is stored. The sum of these counters provides a measure on the number of experiences received, and therefore functions as a partial measure of the reliability of the information. On the other hand, the counter for no outcome values reveals the quality of the experiences when compared to the total count. An example reputation view is depicted in Figure 2.

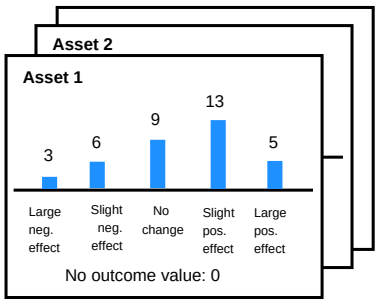

Fig. 2. An example reputation view.

5 Conclusion

Enterprise computing is moving towards more open, collaborative systems, which involves issues in technical, semantic and pragmatic interoperability. Trust man-

agement focuses on the pragmatic interoperability problem of whether two enterprises trust each other enough to want to collaborate. Trust decisions are made both upon setting up a collaboration and during it, to determine whether changes in trust cause the perceived risks to outweigh the incentives and trust between the partners.

The high-level trust model of the TuBE trust management system involves the interplay of four main components: reputation, risk, importance, tolerance, and context adjustments to them. The reputation model was presented in more detail, and the full information model will be further elaborated on in a separate paper.

The Pilarcos interoperability middleware provides a supporting environment for the TuBE trust management system to operate in, through useful facilities such as contract negotiation and monitoring, as well as central frameworks such as a service typing model and partner matching on the technical and semantic interoperability levels. The merging of the TuBE trust management system into the Pilarcos system has begun with an analysis on how the new system is placed within the processes of the existing middleware [12]. The process continues as the trust management system design is finalized and implementation begins.

References

1. Papazoglou, M.P.: Service-oriented computing: Concepts, characteristics, and directions. In: Proceedings of the 4th International Conference on Web Information Systems Engineering (WISE 2003), IEEE Computer Society (2003)
2. Kutvonen, L., Metso, J., Ruokolainen, T.: Inter-enterprise collaboration management in dynamic business networks. In: On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE: OTM Confederated International Conferences, CoopIS, DOA, and ODBASE. Volume 3760 of Lecture Notes in Computer Science., Agia Napa, Cyprus (2005)
3. Kutvonen, L., Ruokolainen, T., Metso, J.: Interoperability middleware for federated business services in web-Pilarcos. International Journal of Enterprise Information Systems, Special issue on Interoperability of Enterprise Systems and Applications **3** (2007) 1–21
4. Ishaya, T., Mundy, D.P.: Trust development and management in virtual communities. In: Trust Management: Second International Conference, iTrust 2004, Oxford, UK, March 29–April 1, 2004. Proceedings. Volume LNCS 2995/2004. (2004) 266–276
5. Karabulut, Y.: Implementation of an agent-oriented trust management infrastructure based on a hybrid PKI model. In: Proceedings of Trust Management: First International Conference, iTrust 2003, Heraklion, Crete, Greece, May 28–30, 2003, Springer-Verlag, LNCS 2692/2003 (2003) 318–331
6. Blaze, M., Feigenbaum, J., Lacy, J.: Decentralized trust management. In: Proceedings of the IEEE Symposium on Security and Privacy, IEEE (1996) 164–173
7. Jøsang, A., Presti, S.L.: Analysing the relationship between risk and trust. In: Proceedings of Trust Management: Second International Conference, iTrust 2004, Oxford, UK, March 29–April 1, 2004, Springer-Verlag, LNCS 2995/2004 (2004) 135–145

8. Cahill, V., et al.: Using trust for secure collaboration in uncertain environments. Pervasive Computing **2** (2003) 52–61
9. Ruohomaa, S., Kutvonen, L.: Trust management survey. In: Proceedings of the iTrust 3rd International Conference on Trust Management, 23–26, May, 2005, Rocquencourt, France, Springer-Verlag, LNCS 3477/2005 (2005) 77–92
10. Ruohomaa, S., Kutvonen, L., Koutrouli, E.: Reputation management survey. In: Proceedings of the 2nd International Conference on Availability, Reliability and Security (ARES 2007), Vienna, Austria (2007) Accepted for publication.
11. Ruohomaa, S., Viljanen, L., Kutvonen, L.: Guarding enterprise collaborations with trust decisions—the TuBE approach. In: Proceedings of the First International Workshop on Interoperability Solutions to Trust, Security, Policies and QoS for Enhanced Enterprise Systems (IS-TSPQ 2006). (2006)
12. Kutvonen, L., Metso, J., Ruohomaa, S.: From trading to eCommunity population: Responding to social and contractual challenges. In: Proceedings of the 10th IEEE International EDOC Conference (EDOC 2006), Hong Kong (2006) Best paper award.
13. McKnight, D.H., Chervany, N.L.: Trust and distrust definitions: One bite at a time. In: Trust in Cyber-societies: Integrating the human and artificial perspectives. Volume LNCS 2246/2001., Springer-Verlag (2001) 27–54
14. Jøsang, A., Keser, C., Dimitrakos, T.: Can we manage trust? In Herrmann, P., Issarny, V., Shiu, S., eds.: Proceedings of Trust Management: Third International Conference, iTrust 2005, Paris, France, May 23–26, 2005. Volume 3477 of LNCS., Springer-Verlag (2005) 93–107
15. Gambetta, D.: Can we trust trust? In: Trust: Making and Breaking Cooperative Relations. University of Oxford, Department of Sociology (2000) 213–237 Electronic edition.
16. Viljanen, L.: Towards an ontology of trust. In: Proceedings of the 2nd International Conference on Trust, Privacy and Security in Digital Business (TrustBus'05). (2005)
17. Abdul-Rahman, A., Hailes, S.: Supporting trust in virtual communities. In: Hawaii International Conference on System Sciences 33,HICSS. (2000)
18. Marsh, S.: Formalising Trust as a Computational Concept. PhD thesis, University of Stirling, Department of Computer Science and Mathematics (1994)
19. Chen, R., Yeager, W.: Poblano—a distributed trust model for peer-to-peer networks. Technical report, Sun Microsystems (2001)
20. Resnick, P., Zeckhauser, R., Friedman, E., Kuwabara, K.: Reputation systems. Communications of the ACM **43** (2000) 45–48