
Guarding enterprise collaborations with trust decisions – the TuBE approach

Sini Ruohomaa – Lea Viljanen – Lea Kutvonen

Department of Computer Science

P.O. box 68

00014 University of Helsinki

Finland

{Sini.Ruohomaa, Lea.Viljanen, Lea.Kutvonen}@cs.helsinki.fi

ABSTRACT: Enterprise computing is currently moving towards more open, collaborative systems. It becomes essential for enterprise success that joining a business network is made efficient, despite the technical and semantic interoperability challenges involved in connecting different information and communication systems. Trust management is an important factor in the collaboration, as traditional trust-building over months of negotiations has become too slow a method in routine cases. As no business network is feasible without mutual trust between partners, the supporting technology should provide mechanisms for forming trust relationships, making automatic trust-based decisions on routine business transactions, and observing the business peers for malicious or incorrect behaviour on interactions. This paper describes a trust model to fulfil these needs, and gives a strategical overview of the system implementing this model.

KEY WORDS: trust management, B2B collaboration.

1. Introduction

Enterprise computing is moving from internal application integration towards more open, collaborative systems, enabling enterprise solutions to work together across organisational borders. This trend is supported by technical development around the SOA and Web-Services technology: They enable autonomously supported, application-level business services to be combined into collaborative business networks.

We view inter-enterprise collaborations as business networks that are governed by eContracts. For these contracts, the collaboration structure and semantics are modelled in terms of roles of the participating business services and interactions between them, and policy constraints to govern the roles and interactions. For an open collaboration, the partner to fill a particular role can be chosen freely from the open market, as long as it is able to fulfil the requirements for the role. These requirements describe the role's business responsibilities as well as more technical interoperability aspects, such as communication and data representation constraints. Functional interoperability requirements are accompanied by non-functional aspects, such as timeliness, availability, security and trustworthiness.

To support this kind of B2B middleware services, the CINCO group has developed partner selection and negotiation, interoperability tests for technical and business aspects of services, collaboration lifecycle management with partner changes, and breach management (Kutvonen, et al., 2005).

The fundamental metainformation element and active agent in the architecture is the eContract, which is created for each business network and which governs the collaboration with a combination of aspects rising from business strategies, legal and other regulatory systems, and technical interoperability needs such as sharing a business process model, and information representation and messaging techniques. The eContract defines not only the successful collaboration cases, but also defines what can be considered a breach, and what partners may or should run as a joint recovery process after a breach case.

To complement this work, trust management concepts, models and middleware facilities are needed. Breaches affecting trust or caused by lost trust become part of the overall behaviour governed by the eContract. The TuBE project aims to address these issues by providing the following:

- Definitions for trust-related concepts, such as trust decision and the context for it, and trustee reputation;
- A general architecture to create and distribute trust-related information;
- Middleware facilities for trust management; and
- Facilities for monitoring and reacting to misbehaviour and anomalies.

We focus on understanding and supporting trust between the partners, i.e. business services. A basic communication and security infrastructure is assumed to be in place, including e.g. identity management, as enterprises will need to identify each other in order to enter into legally binding contracts. Cryptography services deal with message integrity and eavesdropping, and the SOAP messages sent and received are assumed to follow format agreements and arrive in order.

The trust management system has two major tasks: first, it should act as a guard for the service application, applying trust decisions to protect the enterprise from taking too high risks. Second, it should upkeep reputation information on other peers, to allow the system to adjust to how the peers have been behaving in the past.

Trust decisions are needed in two very different situations. When a business network is already running, routine decisions determine whether a particular action should be allowed in the context of the network. Trust decisions are also made when deciding on joining a network or choosing the best partners for it, and whether to remove a peer or leave the network altogether, if it cannot function any longer due to insufficient trust. Both are addressed by the TuBE trust management system.

Trust decisions are built on a combination of situational risk analysis and a strategic viewpoint. Estimating the partner's future actions is key, but some limitations must be considered. For example, if the business network contract defines compensation clauses if some things are not done, the risk analysis should adjust to the strategic situation. Trust decisions should also be able to manage frequent temporary changes in the valuations of the enterprise. The guarded service application acts within the context of a business network, its host enterprise and a technical infrastructure, which should be considered in decision-making. The phase a business network is in can make some actions more important than usual, the enterprise may decide to weigh some risks more as a response to a particular market situation or the underlying system may be low on resources, and we need a means to communicate changes in the situation to the trust management system.

In order to meet these needs, a trust model represents the information to be made available for trust decisions. The representation of risk and the strategic viewpoint to compliment the tactical evaluation will guide the design of a decision-making algorithm. Facilities for describing, accumulating and evaluating experience information will provide a basis for dynamic risk analysis.

This paper provides an overview of the TuBE trust management system and the trust model behind it. Section 2 describes the trust model to fulfil the needs presented in this section, and Section 3 gives an overview of the TuBE trust management system to implement the model. Section 5 discusses implementation issues, and Section 6 concludes.

2. Trust model

The TuBE trust model defines trust 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 (2003) and situational trust by Josang et al. (2004). Our trust management system produces context-dependent and dynamic trust decisions, supported by estimations of the actual trustworthiness of a peer.

For the business network establishment phase, the web-Pilarcos platform provides a populator service (Vähäaho, et al., 2003). It uses a business network model for determining the collaboration structure and roles for participants, and a service offer repository for retrieving metainformation about potential partners. Based on these, the populator makes sufficient interoperability tests and suggests an eContract. The suggestion is then negotiated between the future partners. Both the populator and the partners are able to use trust information for decision-making.

During the operation of the network, the partners make local trust decisions on getting involved in interactions, based on a combination of local and shared information. The decision system realized as a guard in the communication channel lies on the organizational border at both sides between actors. A trust decision is triggered by in- or outbound messages that mark a risk-relevant commitment in an action. In cases where obliged interactions are missed because of insufficient trust, recovery processes encoded into the eContract are triggered.

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 the latter case, the decision must be passed to a higher level for further processing, ultimately to a human. The trustor denotes the party making the subjective trust decision. 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, and terminology is mixed (Viljanen, 2005b).

Reputation, as used in the TuBE model, is the measure of a peer's trustworthiness. Every trustor has its own view of what the reputation of a particular trustee is, so the measure is not bound to a global agreement and there is no need to build a representation of a global trust network. To build its subjective 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 (2000).

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. Some trust and risk models only consider two possible actions by the trustee: cooperation and defection. However, we find this view too simplified in business collaboration. There are different ways and degrees of defecting, such as slightly delayed delivery compared to no delivery at all, or varying quality of the product. The severity ranges of each risk and the weight ranges of each benefit are considered and stored per asset.

The risk parameter depends 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 probability 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 (1994). SECURE has applied continuous cost-benefit probability density functions for risk analysis, which squeezes all assets into one result function (Cahill, et al., 2003).

The importance parameter brings a strategic counterpart to the tactical evaluation contained in the risk parameter. While the risk analysis depends 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, with considerations such as the cost of denying an action, or the benefit of giving great service 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. Poblano (Chen, et al., 2001) uses importance as a strategic tool as well.

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 (Viljanen, 2005b). On the other hand, items such as the *reciprocity* of trust, as discussed by Marsh (1994), can be expressed as a contextual adjustment to the importance factor.

A system state context change may be needed when a denial of service attack has been detected: the perceived risk to service availability should be increased temporarily. Second, a business state context change is in order when storage space or funds are low: the importance of a "sell item" action is increased, which results in higher tolerated risk for that action. Third, should a business network go into a

renegotiation state due to problems with one peer, the risk of some actions strongly dependant on other peers may be increased.

3. The TuBE trust management system

In this section we describe two central subsystems of the TuBE trust management system, which implement the two tasks identified in the introduction: trust decision making and the management of reputation information. In the trust decision subsystem, the guard combines local and global trust information from the data processing component into a local trust decision whenever it is called for. The reputation management subsystem upkeepes reputation estimates used by the trust decision subsystem. It does this by combining experience information from local monitors and reports received from the global reputation network. An overview of the TuBE trust management system infrastructure is given in Figure 1.




Figure 1. Overview of the TuBE trust management framework. Connections to external systems are made through the guard and the reputation network.

3.1. Making trust decisions

The process to produce a trust decision begins when the guard intercepts a SOAP message on its way to invoke a service method. Each message type is part of a specific action, and the guard stores message types that should trigger the trust

decision for the action they represent. As noted in Section 2, the triggering message may also originate from the local application. The web-Pilarcos middleware makes sure that inbound messages form well-defined interaction patterns that conform to the eContract.

The guard examines its message type database to determine which parameters in the SOAP message are relevant to the decision. It then extracts these parameters and sends to the trust decision maker the set $\{trustee, action, parameters\}$. The trust decision maker passes this input to the data-processing component. It receives in return a risk analysis and a constraint set.

The risk analysis represents the calculated risk for the action, expressed as a set of cost-benefit probability distributions, one for each identified asset. For many assets, there are more than two categories of outcomes for different levels of costs and benefits. Others, such as money, can even have continuous ranges of possible cost or benefit, although a human user will probably find categories more intuitive in these cases as well. The constraint set represents the acceptable risk for this action in the current situation, with accepted probability ranges for different benefits and costs. Some aggregating constraints can also be provided in the set, such as the combined probability for either “minor” or “considerable” loss of security. The trust decision is made by comparing whether the risk estimate fits within the constraints. “Unsure” decisions are supported by their own constraint set.

The data-processing component includes four subcomponents: risk, importance and reputation evaluators, which are connected to a context evaluator. The component uses information from the action data storage, experience data storage and different context sources to respond to the guard’s queries.

Ontological information about the action and its parameters is held in the action data storage. This information may e.g. identify that a “sell a movie” action is strongly related to the more general “sell” action. The action data storage also contains a set of formulae for estimated default risk for *action*, and when action *parameters* are applied to it by the risk evaluator, a default risk estimate is created. The set of defaults provides parameter-dependent probability distributions for the possible cost or benefit to different assets possibly affected by the action. Experience information concerning the *trustee*, *action* and relevant other actions is held in the experience data storage, and the defaults given by the default probability distributions are adjusted based on experience information as well as context policies. To respond to a risk request, the data processing component uses the input triple, $\{trustee, action, parameters\}$, to gather relevant information for the three evaluators and to apply the relevant contextual adjustments to produce a completed risk estimate.

To generate the constraints to represent acceptable risk, a constraint policy set is retrieved from the action data storage for the *action* at hand. The importance evaluator then applies relevant action *parameters* and the *trustee*, following the applicable context policies for this operation, to produce the constraints for

acceptable risk. As there is a considerable number of different outcomes and possible risks to represent, they must all be combined to a single set of probability distributions beforehand for efficiency purposes. This can be done as risks are evaluated in the enterprise. SECURE combines the cost-benefit probability density functions for every possible outcome on the fly when performing risk analysis (Cahill, et al., 2003).

3.2. Reputation management

The aggregated experience information used in trust decisions is upkept by the reputation subsystem. A local reputation view is produced by combining experience from a local monitoring system with information gained from the reputation network. Experience captures the effects of past actions to identified assets.

The local monitoring system observes SOAP messages forwarded by the guard, but can also receive information from other data sources. An application-level monitor can detect anomalies in SOAP message exchanges, such as an unusually large quantity of ordered goods, and divergence from contractual specifications, such as requests for payment without a valid order. However, monitoring SOAP message exchanges alone does not capture many out-of-band events that can be considerably more important in the overall experience, e.g. delivering poor-quality goods behind schedule, leaving invoices unpaid or handling reclamations poorly. Many of these events do leave traces in information systems which are not directly connected to the service application. This information about violations can be captured by monitors connected to each of these separate systems. It can then be reported on to the local reputation system.

Some events do not leave traces in information systems at all. For example, poor reclamation handling by a partner can cause customers to take their money elsewhere. This kind of problem can result in a need to break the partnership, but the need is observed by people in the enterprise and dealt with by human interactions. It is essential that the monitoring system also accepts input from human “monitors”, which can capture certain kinds of violations much more efficiently than any analysis machinery could.

Most monitors used to gather local experience are far from infallible, and the information they search for may be of different value in determining an actor’s reputation. Especially anomaly detection is prone to false alarms (Viljanen, 2005a), as a change in behaviour is not always for the worse. The experience produced by local observation therefore also contains a measure of confidence in the report, on a percentage scale. This confidence is determined by the local reputation system in accordance with the monitor type and the kind of event it reports. The measure is used in determining the impact that the new item of experience has on the current local reputation; the higher the confidence measure, the greater the impact can be. The amount of information accumulated thus far plays a role here as well: minor

problems have more influence on a young business relationship than one that has gone on successfully for years.

The reputation system combines specialized event information from the monitors into experiences. It must consider the relationships between the events to build a larger view: an unusual message exchange for a purchase order is problematic only if payment does not follow, and the human user may provide overriding information. The monitors also provide active responses to threatening events e.g. through context management, but such adjustments are outside the scope of this paper.

The global reputation system combines the experiences of other peers. Sharing experiences with other actors helps businesses avoid making partner selection mistakes that someone else has already made. While the impact of word of mouth can be considerable, this “second-hand” experience brings problems of its own: it is difficult or impossible to check what a statement of experience is based on. Instead of an honest report, it could be a product of collusion or an attempt to make a competitor look bad.

The TuBE reputation subsystem accepts experience information from third parties, and before storing them attaches a percentage measure of credibility to each item of information. The process for deducing this credibility measure lies at the heart of successful use of the global reputation system.

A drop in reputation and its negative impact on trust can have serious effects on the partnership, such as partner removal. Partner removals, as well as locating a replacement partner and renegotiating the contract, are handled by a web-Pilarcos middleware service.

4. Implementation issues

Issues in the realization of the TuBE trust management system include: a) the representation and interpretation of trust information, b) the management of information sources in an effective way, c) performance penalties caused by the guards, and d) the cost of introducing the generic TuBE facilities into enterprise systems.

In the TuBE system, automated measurement of trusting belief alone is not sufficient, but it must be used for decision-making. The trust belief information, such as local reputation views, is seen in form of probabilities; trust is interpreted as a probability measure for success. Using a single probability as a trust belief measure would force an assumption that the outcomes of any action can be divided into two groups, cooperation and defection, in accordance with a game-theoretic world view. We consider this too broad a simplification for business applications. In the field of business interactions, there are clearly several levels of cooperation and defection, such as late delivery or no delivery at all. Our approach to consider each asset separately allows more specific policy in one sense, and in return allows free

choice of the specificity of cost and benefit categories according to the enterprise needs. Our aim is that a user enterprise could utilize risk analysis information gathered using separate tools in the configuration of its trust management system.

In each enterprise system, trust decisions are supported by evidence on the behaviour of known or potential trustees. The evidence includes both first-hand and second-hand experience. It is upkept through feedback loops from the local monitoring system and from external reputation systems. The challenges related to these sources are quite different: experience received from the global reputation network requires not only credibility analysis, but also analysis on how experience from very different activities relates to those relevant in the local system.

The monitoring system may cause major overhead in the system, even if the service monitor studies sent and received SOAP messages only. There are also intrusion detection and prevention approaches that hook into the operating system or platform the service application is running on, such as the Java VM, or build sensors inside the application (Viljanen, 2005a).

We must trigger trust decisions at relevant points of the exchange only, carefully adjust the width and focus of any anomaly and breach detection to actual enterprise needs, and limit the amount of information to transfer to the decision point. For example, experience information can be aggregated into compound items in different levels (English, et al., 2003; Liu, et al., 2004), which can alleviate both storage and transfer limitations.

The set of actions and the set of trustees are both large and dynamic, but at any time a single guard needs only a fraction of this information. Caching the data most relevant to the current interaction near the guard is a beneficial trade-off between transfer load and information freshness. The selection of the relevant information can be based on partnership information and business network activities information available from the web-Pilarcos facilities. The TuBE trust management system is strongly based on the concept of *action*, and information on action types and ontology is needed. Here we can utilize service typing information used in describing and matching Web Services in web-Pilarcos (Ruokolainen, et al., 2005b). The available information provides for both the action ontology needed for generalizing experiences, and the message typing used in the guard.

We require a reputation network to provide information classified by actions and trustees. In addition, in our model a trustee is a business service, as opposed to e.g. a human user, a computer or an entire enterprise. Current reputation systems are highly varied and incompatible, and there is nothing resembling a standard solution available. We have specified what kind of experience information we wish to receive from the network and the interface through which it is accessed. Studying the interoperability of different reputation systems is an item of future work.

5. Conclusion

This paper proposes a partner-to-partner trust model that is based on a global reputation flow and local trust decisions guarding inter-enterprise collaborations. The trust-aware guards influence actions taken at two levels: the establishment and negotiation of collaboration relationships, and significant inter-enterprise interactions. This kind of trust model is essential for federated architectures for inter-enterprise collaboration management, such as web-Pilarcos (Kutvonen, et al., 2005).

Significant features of the TuBE trust model include dynamic, multi-source accumulation of reputation information, and timely interception of business interactions. The solution compliments traditional security services by trust-based soft security, which is more applicable in open collaboration networks (Rasmusson, et al., 1996).

Although the field of trust management is still somewhat diverse (INTEROP-NoE, 2005; Viljanen, 2005b; Ruohomaa, et al., 2005), the TuBE trust model conforms to the commonly required main elements and furthermore elaborates the trust decision information to aspects relevant for inter-enterprise collaboration and business process management. The TuBE system combines the information collecting tasks traditionally most visible in reputation systems research (Liu, et al., 2004; Obreiter, 2004), and the automated decision making which the first trust management systems (Blaze, et al., 1998) have focused on.

Further work on the TuBE systems will bring the design into the existing web-Pilarcos prototype platform, and allow us to compare the effects of various trust decision algorithms to strategic business goals. Furthermore, trust and reputation information ontologies and reputation system interoperability are relevant areas of research.

This article is based on work in the web-Pilarcos and the TuBE projects (Trust based on evidence) at the Department of Computer Science at the University of Helsinki. The web-Pilarcos project is run in collaboration with VTT, Elisa, SysOpen, and in addition funded by the National Technology Agency TEKES in Finland, and Tellabs. The TuBE project is funded by TEKES, Nixu, and StoneSoft.

7. References

- Abdul-Rahman, A., Hailes S., "Supporting Trust in Virtual Communities", *Hawaii International Conference on System Sciences, HICSS*, Jan 2000.
- Blaze M., Feigenbaum J., Keromytis A. D., "KeyNote: Trust management for public-key infrastructures (Position Paper)", *Proceedings of Security Protocols: 6th International Workshop*, Springer-Verlag, LNCS 1550, Apr 1998, p. 59 – 63.

- Cahill V. et al., "Using trust for secure collaboration in uncertain environments", *Pervasive Computing*, vol. 2, num. 3, 2003, p. 52 – 61, IEEE.
- Chen R., Yeager W., "Poblano – a distributed trust model for peer-to-peer networks", report, 2001, Sun Microsystems.
- English C., Wagealla W., Nixon P., Terzis S., McGettrick A., Lowe H., "Trusting collaboration in global computing systems", *First International Conference on Trust Management*, LNCS 2692, Springer-Verlag, May 2003, p. 136 – 149.
- Gambetta D., "Can We Trust Trust?", *Trust: Making and Breaking Cooperative Relations*, University of Oxford, Department of Sociology, 2000, p. 213-237, Electronic edition.
- INTEROP-NoE Task Group 7, "Roadmap for TG7: Interoperability Challenges of Trust, Confidence, Security and Policies", 2005, In preparation.
- Jensen C., Poslad S., Dimitrakos T., Eds. *Proceedings of Trust Management: Second International Conference*, LNCS 2995, Springer-Verlag, Mar 2004.
- Josang A., Presti S. L., "Analysing the Relationship between Risk and Trust", Jensen et al. (2004), p.135 – 145.
- Kutvonen L., Metso J., Ruokolainen T., "Inter-enterprise collaboration management in dynamic business networks", *OTM Confederated International Conferences, CoopIS, DOA, and ODBASE*, LNCS 3760, Springer-Verlag, Nov 2005, p. 593 – 611.
- Liu J., Issarny V., "Enhanced Reputation Mechanism for Mobile Ad Hoc Networks", Jensen et al. (2004), p. 48 – 62.
- Marsh S., "Formalising Trust as a Computational Concept", PhD thesis, University of Stirling, Dept. of Computer Science and Mathematics, 1994.
- McKnight D. H., Chervany N. L., "Trust and Distrust Definitions: One Bite at a Time", *Trust in Cyber-societies: Integrating the human and artificial perspectives*, LNCS 2246, Springer-Verlag, 2001, p. 27 – 54.
- Obreiter P., "A Case for Evidence-Aware Distributed Reputation Systems Overcoming the Limitations of Plausibility Considerations", Jensen et al. (2004), p. 33 – 47.
- Rasmusson L., Jansson S., "Simulated Social Control for Secure Internet Commerce", *Proc. of the 1996 workshop on New Security Paradigms*, ACM Press, 1996, p. 18 – 25.
- Ruohomaa S., Kutvonen L., "Trust management survey", *Proceedings of Trust Management: Third International Conference*, LNCS 3477, Springer-Verlag, Apr 2005, p. 77 – 92.
- Ruokolainen T., Kutvonen L., "Service Typing in Collaborative Systems", accepted for publication in the proceedings of I-EISA 2006.
- Viljanen L., "A Survey of Application Level Intrusion Detection", report, 2005, University of Helsinki, Dept. of Computer Science.
- Viljanen L., "Towards an Ontology of Trust", *Proceedings of the 2nd International Conference on Trust, Privacy and Security in Digital Business (TrustBus'05)*, 2005.
- Vähäaho, M., Kutvonen L., "Enhanced trading service in middleware for inter-organisational applications" report C-2003-15, University of Helsinki, Dept. of Computer Science, 2003.