The TuBE approach to trust management

Lea Viljanen, Sini Ruohomaa, and Lea Kutvonen

University of Helsinki, Finland
{lea.viljanen, sini.ruohomaa, lea.kutvonen}@cs.helsinki.fi,
WWW home page: http://www.cs.helsinki.fi/group/tube/

Due to fierce competition, organizations must focus on their key strengths to survive. In order to avoid a complete fragmentation of services, this tendency is mitigated by cooperation with other organizations to provide a composed service. The resulting communities of autonomous service-provider peers suffer from uncertainty: the other partners are not controllable, and the service provided may be of varying quality or even nonexistent. The partners sometimes are competitors in related fields, and a peer may be uncooperative, compromised or malicious.

In this kind of environment, making access-control and partner-selection decisions based on a static policy is unsustainable. A continuous application-level re-evaluation of peer trustworthiness is needed. For situation-aware decisions, the risks and potential benefits related to a positive decision should also be considered.

The TuBE project (Trust Based on Evidence) aims to provide middleware services for trust-based decisions on authorization and partner selection. Trust in a partner is updated based on experience gathered from earlier collaboration. These services enhance the B2B middleware from the web-Pilarcos project [1], which provides lifecycle management and interoperability support for open business networks.

Trust decisions are made locally, and local observation is used to update the dynamic trust views. The TuBE draft architecture includes local experience gathering in the form of anomaly detection and checking compliance to expected information and behaviour patterns. To take advantage of other partners' experiences as well, peers exchange their views of other peers' trustworthiness in the network to build a better-informed view of their collaborators. This external reputation information, once analyzed for credibility, can be used to enrich the local view. The architecture includes an engine to update local reputation based both on experience from local observation and external reputation information from other peers.

The TuBE trust information model consists of a 7-tuple of trustor, trustee, action, trustee reputation, action risk, action importance, and trustor context. A trust decision depends on the trustor, the trustee and the action—typically a service request. Different actors and actions require their own trust analysis.

The next three factors, describing the trustee and action in the trustor's view, are the main inputs to a decision policy. The trustee reputation represents estimated trustworthiness, and is also action-specific. The factors describing the action are the amount of risk related to authorizing the action, and the action importance. The latter measures potential benefits as well as the degree of necessity. For example, a slightly misbehaved partner may be allowed to participate in a low-risk delivery of gravel to a customer, but not to deliver diamonds. On the other hand, if a delivery of dairy products is needed to enable the quite profitable production and sale of cheese at another point of the virtual organization, even a somewhat unreliable partner may be allowed a second chance, as the potential benefits are high.

The final factor, trustor context, indicates what temporary modifications are needed for the values of the previous triple of reputation, risk and importance, in order to adjust the basic setting to the current situation. For example, if cash reserves are low, a context setting can be applied to increase the importance of selling, and possibly increase the risk of buying. The TuBE services include subsystems to upkeep context information as well as risk and importance values for different actions.

Compared to other systems, characteristic for the TuBE system is a dynamic trust view, which is updated from experience and adjusted by a dynamic context. To simplify reputation gathering and exchange, an identity management infrastructure is assumed to be in place. Decisions about a willingness to trust aim to balance several factors, including reputation and risk. After a trusting decision, the outcome is observed to determine the quality of cooperation or lack thereof.

Early approaches to trust management, such as Keynote [2], focus on trust based on authentication and credentials. Unless credentials change outside the system, the trust view remains static.

The SECURE project [3] considers trust and risk in a ubiquitous computing environment. Identification builds on pseudonyms instead of identity, and dynamic context is not addressed. Authorization decisions are based on the results of risk analysis, which is directed by local and external reputation.

Marsh provides a computational model of trust for agent systems [4], which shares many factors with the TuBE trust model. Marsh divides experience into two outcome types: cooperation and defection. The TuBE experience view is closer to a continuum, considering different levels of cooperation as observed by the means of anomaly detection.

References

1. Kutvonen, L., Ruokolainen, T., Metso, J., Haataja, J.P.: Interoperability middleware for federated enterprise applications in web-Pilarcos. In: Proceedings of INTEROP-ESA'2005. (2005) 196–208
2. Blaze, M., Feigenbaum, J., Ioannidis, J., Keromytis, A.D.: The KeyNote trust-management system, version 2 (1999) Request For Comments (RFC) 2704.
3. Cahill, V., et al.: Using trust for secure collaboration in uncertain environments. Pervasive Computing **2** (2003) 52–61
4. Marsh, S.: Formalising Trust as a Computational Concept. PhD thesis, University of Stirling, Department of Computer Science and Mathematics (1994)