Trust Aspects in the Architecture of Interoperable Systems

Lea Kutvonen

University of Helsinki, Department of Computer Science
Lea.Kutvonen@cs.helsinki.fi

Abstract. By introducing trust concepts to the enteprise computing arena, a more user-oriented view to trustworthy services becomes available. We can consider a business service to be trustworthy, if it is likely to provide us the right functionality and to deliver it in a manner that is, for example, timely, secure, and privacy-preserving.
In order to achieve this goal, a lot of progress has to be made on research, development and standardisation. The existing trust-management and reputation-management systems provide a good selection of different solutions for making informed decisions on entering collaborations or participating business interactions. However, no common standards exist.
This paper introduces the work performed in INTEROP NoE TG7 on trust related aspects, and elaborates it towards a roadmap relevant for federated, inter-enteprise computing environments.

1 Introduction

The present trend in business management is towards networked organisations and business networks across enterprise boundaries. The development of suitable enterprise systems for these needs require development of a coherent computing architecture that matches the needs of the networked enterprise architecture. This paper discusses trust related aspects of such enterprise computing architectures. In addition to the technical challenges on this area, there are still unsolved problems, for example, in the business management area for defining the level and points of automation that can be tolerated and exploited, and in the legal domain for defining the effect of electronic contracts and breaches of them.

In both scenarios, the computing architecture comprises of a set of business services working together according to a set of business processes in order to jointly provide an added value service. The business services are realised by different parts of an organisation, or by different enterprises. The degree of joint coordination and mutual trust between the business services that can be assumed, depends on the degree of independence of the service providers.

When there is mutual coordinator or a project providing shared coordination rules, an integrated system can be formed. When there is a shared model that independent service providers can agree to use for the basis of interoperation, an model-based solution or unified system view can be achieved. However, the shared model is there to define an initial state of affairs, and is not necessarily

agile, for example, for changes in the partnership, technology, context or phase of collaboration, and breaches. Eventually, a fully dynamic, federated computing support can be reached if the computing platform is enhanced with common negotiation and contract management protocols, and common models and metrics for contractual contents.

While the level of automation in contract negotiation can be increased, there will be additional needs for defining clearly in which extent the automated agents of the enterprise are allowed to make commitments.

The present development of enterprise systems is much focused on the functional aspects of the collaborations. However, management of the collaborations should address pragmatic aspects as well, i.e., willingness of participation in a community with given partners and given business rules, and detection of breaches in terms of contents of exchanged documents, timeliness of the service and other qualitative aspects. One of the fundamental aspects in the pragmatic level is that of trust-based decisions on collaboration participation and behaviour.

The INTEROP NoE TG7 has produced a roadmap that covers the state of the art on non-functional aspects [1]. In this paper, the trust-related issues are briefly introduced, concentrating on the inter-enterprise collaboration theme. An additional focus of interest for the subgroup was on trustworthiness and privacy-preserving of data integration from distributed, independent data sources.

In the following, first the various dimensions of trust are discussed, separating out needs for a trustworthy infrastructure with societal services and the peer-to-peer reputation management needs for business level considerations. Section 3 mainly discusses the still occurring variety of trust models on the field, addressing the need of standardising the concepts used and metrics for them. Section 4 addresses the trust management functionalities needed. Section 5 reflects trust as one of the non-functional system aspects and returns to the role of layers of system trust. The conlusion addresses the benefits of the research and the future work issues.

2 Dimensions of trust

In an inter-enterprise collaboration case, the set of potential trustors and trustees includes persons, organisations, infrastructure agents, or application services or information. Trustor is the role for an entity deciding to participate a collaboration or an individual activity; trustee is the identified entity considered as a peer in the collaboration or activity.

Some of the important trust-related terms can be listed, for example according to McKnight and Chervany [2] as follows.

- *Trusting belief* is the extent to which we believe another party able and willing to act in out best interest.
- *Trusting intention* is the extent to which we are willing to depend on another party in a given situation, taking into account the potential risks involved.

- *System trust* is the extent to which we believe the infrastructure and societal services to be in place and support our interests.
- *Dispositional trust* denotes our trusting or distrusting attitude towards other parties.
- *Situational decision to trust* denotes our readiness to to trust other parties in general, in a given situation.

Models and systems developed for supporting trust decisions and activities requiring trust management mechanisms are concerned with a variety of dimensions, as shown in Figure 1. In the figure, layers of interest can be seen as follows.

- Human users trust each other in some extent to have shared goal for collaborations, and to choose collaboration pattern that lead to win-win situations. The collaboration is enabled by trusting belief and trusting intention, and restricted by limitations of dispositional trust and situational trust.
- Users trust the system or business services they use; these services can be created by a community of networked enterprises, and represented as an agent for the user. This is a phenomenon of system trust.
- As part of the ICT system, the applications trust the computing facilities and the communication solutions to provide an accurate, unchanged, private service in terms of information exchange and processing. This is a phenomenon of system trust, but is built on an implicit trusting belief that the lower layers include sufficient security aspects and have situational trust to each other. Awareness of the level of trusting intention is minimal in the present systems.
- The networked enterprises and the distributed computing infrastructures involved consist of agents working on behalf of the business applications, in a collaborative way, and those agents need to trust each other for accurate information, services, unviolated integrity and accurate metainformation about the management of the collaboration. This involves requirements for situational trust and trusting belief. In the present systems, the trust relationship to peer systems is implicit, but in future systems, elaboration on security mechanisms and control on the "overlay network" partnership must be developed.
- The infrastructures for collaboration management must trust that the issuers of credentials, security and privacy policies, identification providers, etc. are trustworthy and follow joint juridical, contractual, and business-oriented regulations and do that in a technically sound manner. This is a phenomenon of system trust. Especially, users trust the infrastructure to provide accurate information for making trust decisions based on the information available on the other parties against the private knowledge about the situational and dispositional aspects.

For the future inter-enteprise collaboration support systems, we should build infrastructure level services that provide the users and the infrastructure level agents with suitable, trustworthy information for trusting belief decisions that take into consideration the reputation of the identified services and the situation in which the client for that service finds itself [3].

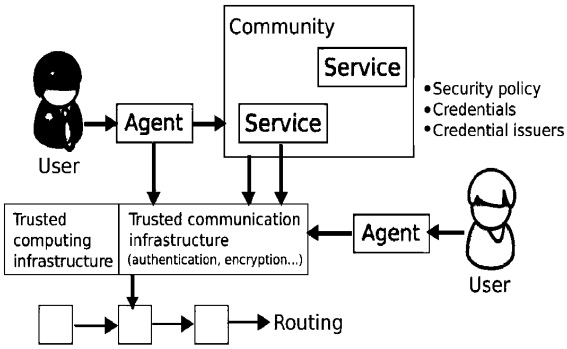

Fig. 1. Dimensions of trust.

3 Trust decisions

For the roadmap work, the TG7 group performed state-of-the-art surveys on systems claiming to have a trust management system or trust model developed [1, 4–10]. Systems like SECURE [11], TrustCom [12], iTrust [13], TuBE [14], T-SAS [15], EigenTrust [16] and Trustbuilder [17] were covered. Later, we surveyed reputation systems to complete the picture, covering eBay [18], Unitec [19], FuzzyTrust [20], REGRET [21], NICE [22], Managing the Dynamic Nature of Trust (MDNT) [23], PeerTrust [24], Managing Trust [25], Maximum Likelihood Estimation of Peers Performance (MLE) [26], and Travos [27]. It appeared, that there is significant variance on the conceptual models still, in addition to the alternating usage of terms like trust and reputation information.

For comparison we concentrated on a few features of the trust-decisions; the trustor being one of the following: a person, organisation, service process, information source infrastructure agent, or a credentials user. For the decision-making different types of properties were assumed to be know of the trustee and the situation in which the decision was made. For example, the qualities of the trustee could include any combination of the following: identity of the trustee, conformance to a named security or privacy policy, competence, or dependability.

Furthermore, there were differences in the indications the trust decisions had. The trust decisions could be part of service or information provider selection, establishment of contract, or restriction of information visibility based on the trustee identity or trustee properties. The indications were either considered to lead to long-term trust-relationships or affecting a single transaction only.

The most variety was, as expected, found in the situational considerations: Some systems refer to moral states involving intentions, potentially considering legal and societal restrictions. Meanwhile others concentrate on computational

systems and discuss system components providing audit trails, authorisation, identification, integrity and availability.

Architecturally, the approaches can be divided into graph-based systems and unstructured systems. In networked business in general, the means of predicting the behaviour of a new partner in a collaboration is by reputation or recommendation. For open service markets, also the reputation information should be freely flooded to interested parties. In contrast to this, the graph-based systems realise a recommendation graph between the essential partners.

In a graph-based system, the trust relationships are captured into a trust-graph that is visible throughout the networked system. Important questions deal with making the trust graph public or private and the semantics for example of the transitivity of trusting beliefs. The graph can be open or closed; in an open graph it is possible to introduce new service providers to the graph by recommendations of already linked participants. In a closed system, there is no dynamic method for changing the graph; this would suit for example for a graph indicating infrastructure level services and their relationships.

The unstructured systems rely on a flood of reputation information. There is different levels of credibility information available for the reputation information items, depending on the used system. Also for reputation management systems, the variety of solutions is still wide [28].

4 Trust and reputation management

Trust and reputation information and other elements affecting trust decisions change over time as the context of the trustor changes and the trustee properties either change or become observed. Therefore, the trust management facilities include categories of:

- initialisation of trust information for a trustee;
- observing or measuring properties of the trustee and accumulating that information as trust or reputation values;
- use of trust information for trust decisions;
- managing trust relationships (contracts, graphs) and delegations;
- managing and interpreting the integrity of trust and reputation information.

The creation and aggregation of trust or reputation values can follow a number of methods, for example statistical models and tools such as regression analysis of feedback received from users; probabilistic models for more accurate representation of uncertainty between alternative behaviours; social network-based models; and game-theoretical models.

Automated aggregation of reputation information requires that there were a coherent ontology on services, information types, and providers that could be used as vocabulary for reputation information exchange. Further, for this conceptual world, metrics of expressing the good and bad experience and exchanging that as reputation values is needed. So far, there is no commonly accepted ontology or set of metrics [28].

For the graph-based solutions, trust can be either formed by receiving recommendations from trusted partners, or by negotiating by exchanging credentials.

Using reputation as part of the trust decision making process requires understanding of the credibility and trustworthiness of the reputation information. However, many of the surveyed reputation systems did not take this aspect into consideration [23, 28]

The trust management facilities depend on infrastructure services that provide security and trustworthy identification and traceability of trustees, and secure and private communication.

5 Trust as a non-functional aspect in enterprise computing

In terms of the natural location of management in the overall federated architecture, the non-functional aspects can be categorised into openly negotiated aspects and private decision aspects.

First, aspects such as QoS, secure communication and processing, transactionality (sic!), and various forms of distribution transparency can be embedded to the supporting abstract communication platform and can be selected by mutual collaboration contract. To make this work, we still need to go further in standardisation of the communication platforms and modularisation of them to form suitable service-level modules that can be selected based on a contracted model [29]. Further, there are non-functional aspects such as business policies, and choices over alternative behaviours, related to the business view of the collaboration. Like the communication agreements, these aspects can be selected via mutual negotiations, and the agreement captured in the contract in terms of selected functional models and parameters refining them [29].

Second, there are private decisions involved at each collaboration establishment or business interaction started. Trust decisions and business policies guarding, for example, preservation of information privacy belong to this group.

To be able to provide solutions for these needs, we need to address the two dimensions of trustworthiness or dependability in the system. First, we must provide facilities for capturing the trust a business service places towards a peer service in a collaboration. Second, we must provide for the layers of system trust involved.

Above, only the trust between business services have been discussed. However, equally pressing problem is to create a trustworthy platform for making routine trust-decision on behalf of the service users, and providing credible trust-related information for decision-making. Two key issues here are the reputation information creation, and association of it to a stable enough identity management system. The identity management system should become a trusted-third-party societal service, for example in a form of an Internet overlay network. The network should provide information that is non-repudiable, protected, traceable, accurate, privacy-preserving, and so on.

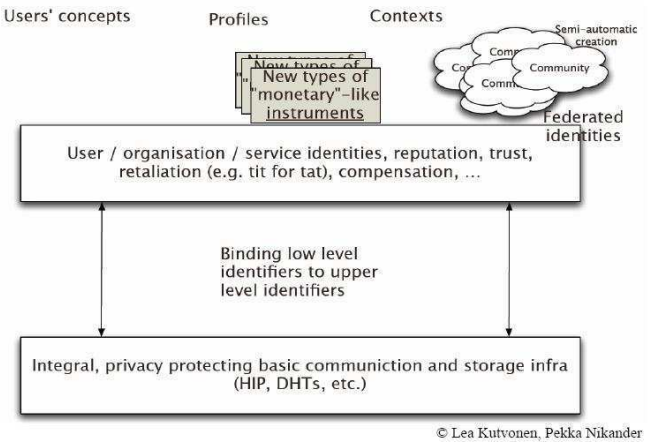

Fig. 2. Need for new concepts for trust, reputation, and identities.

To mature the viability of the reputation system area for routine trust decisions in business-to-business collaboration management, a number of challenges has to be addressed. We have envisioned a global system where trust decisions are made locally, but reputation information is shared in a global reputation management network.

First, for this vision, the reputation information should be standardised to achieve interoperability between systems and services that use them. The granularity of targets to which reputation information is associated should be first determined and then, suitable identification mechanism for these targets provided. The granules of interest depend on the application area, but can involve for example humans, machines, and business services.

Second, experience-based reputation information should be based on a commonly acceptable framework of concepts, ranging for example from successful and correct performance in business transactions to illegal transactions or breaches of technical criteria. For all these axes, ontologies should be developed to capture the metrics to be used.

Figure 2 illustrates how there is need for creating more user-oriented or manager-oriented concepts and access interfaces for trust-related information and guidance of the semiautomated trust-decisions. These concepts should be cleanly mapped to the technologies in use.

Third, the role we envision for reputation systems in the open collaborations creates new vulnerabilities. We have started a comprehensive threat analysis of systems supporting trust, reputation and privacy management, but additional work is still needed for creating a system that would resist these new threats. One of the essential aspects of this development is the extensive use of credibility metainformation on exchanged reputation information and development of trust

decision algorithms sensitive to both the credibility measures and changes in them.

Finally, we have noted the absence of benchmarks suitable for comparing the effectiveness (performance) of making trust decisions, or causing changes in trust decisions depending on the reputation information.

Trust and dependability issues are often left as an add-on-property to be dealt with late in the system development time. However, dependability cannot be added afterwards, but must be considered as first-class feature that must be addressed at the architecture design time, taking into account trust management quality aspects [6].

As an industry driven approach, the Web Services technology family provides a topical framework where an architecture with identification authorities and credential token issuers is presented, and federation between authorities defined [30, 31]. Other recommendations in the group provide for dependable service provision and secure messaging between service providers; still, the scheme is less rigorous than is visible in the research arena.

6 Conclusion

By introducing trust concepts to the enteprise computing arena, a more user-oriented view to trustworthy services becomes available. We can consider a business service to be trustworthy, if it is likely to provide us the right functionality and to deliver it in a manner that is, for example, timely (QoS), secure (non-repudiable, untampered, traceable), and privacy-preserving (privacy-policy guarded access, encrypted communication).

In order to achieve this goal, a lot of additional research and development work is required. We need standard concepts for ownership, accessibility, trust, reputation and commitment. We need to create suitable interfacing styles for the user and manager views to the various assets in the supporting system and application domain that affect these concepts. Especially, we need to develop an automated, role-based authorisation method for structured, multi-owner information items. This is a relevant facility for example for applications around health-care where pieces of patient record are owned by different members of the medical personnel, and for the trust-management infrastructure itself. Furthermore, identification of services needs its overlay networks.

Finally, the management of trust information by the supporting computing and communication infrastructure creates a new level of privacy problem; not only is the information about the activities of a person, organisation, or agent potentially to be considered private, but also the accumulated information about the trustworthiness of the entity must be.

References

1. INTEROP-NoE Task Group 7: INTEROP-NoE Task Group 7: Roadmap for TG7: Interoperability challenges of trust, confidence, security and policies (2005)
2. McKnight, D., Chervany, N.: The meanings of trust. Technical report, University of Minnesota, MIS Research Center (1996)
3. 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, IEEE (2006) 199–210
4. Ruohomaa, S., Kutvonen, L.: Trust management survey. In: Proceedings of the iTrust 3rd International Conference on Trust Management, Rocquencourt, France, Springer-Verlag (2005)
5. Viljanen, L.: Towards an ontology of trust. In: Proceedings of the 2nd International Conference on Trust, Privacy and Security in Digital Business (TrustBus'05), Springer-Verlag (2005)
6. Sindre, G.: Trust-related requirements: a taxonomy. In: Proceedings of the 15th International Conference on Information Systems Development (ISD'06). (2006)
7. Zhang, Q., Yu, T., Irwin, K.: A classification scheme for trust functions in reputation-based trust management. In: International Workshop on Trust, Security, and Reputation on the Semantic Web, Hiroshima (2004)
8. Suryanarayana, G., Taylor, R.N.: A survey of trust management and resource discovery technologies in peer-to-peer applications. Technical report, ISR (2004)
9. Despotovic, Z., Aberer, K.: Maximum likelihood estimation of peers' performance in P2P networks. In: Second Workshop on the Economics of Peer-to-Peer Systems. (2004)
10. Scannapieco, M., Missier, P., Batini, C.: Data quality at a glance. Datenbank-Spektrum **14** (2005) 6–14
11. Cahill, V., et al.: Using trust for secure collaboration in uncertain environments. Pervasive Computing **2**(3) (2003) 52–61
12. Dimitrakos, T., Wilson, M., Ristol, S.: TrustCoM – a trust and contract management framework enabling secure collaborations in dynamic virtual organisations. ERCIM News **59** (2004)
13. iTRust: iTrust Working Group on Trust Management in Dynamic Open Systems. (2005)
14. 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)
15. Lo Presti, S., Butler, M., Leuschel, M., Booth, C.: A trust analysis methodology for pervasive computing systems. In: Trusting Agents for Trusting Electronic Societies. Number 3577 in Lecture Notes in Artificial Intelligence, Springer-Verlag (2005) 129–143
16. Kamvar, S., Schlosser, M., Garcia-Molina, H.: The EigenTrust algorithm for reputation management in P2P networks. In: Proceedings of the 12th International World-Wide Web Conference (WWW03). (2003) 446–458
17. Winsborough, W., Seamons, K., Jones, V.: Automated trust negotiation. In: Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX '00), IEEE (2000) 88–102
18. eBay: The online marketplace (2006) http: //www.ebay.com/.

19. Kinateder, M., Rothermel, K.: Architecture and algorithms for a distributed reputation system. In: Proceedings of Trust Management: First International Conference (iTrust 2003), Springer-Verlag (2003) Volume 2692 of LNCS.
20. Song, S., Hwang, K., Zhou, R., Kwok, Y.K.: Trusted P2P transactions with fuzzy reputation aggregation. IEEE Internet Computing (6) (2005)
21. Sabater, J., Sierra, C.: Reputation and social network analysis in multi-agent systems. In: AAMAS 02: Proceedings of the First International Joint Conference on Autonomous Agents and MultiAgent Systems. (2002)
22. Lee, S., Sherwood, R., Bhattacharjee, B.: Cooperative peer groups in NICE (2003)
23. Staab, S., Bhargava, B., Lilien, L., Rosenthal, A., Winslett, M., Sloman, M., Dillon, T.S., Chang, E., Hussain, F.K., Nejdl, W., Olmedilla, D., Kashyap, V.: The pudding of trust. IEEE Intelligent Systems **19**(5) (2004)
24. Xiong, L., Liu, L.: PeerTrust: Supporting reputationbased trust for peer-to-peer electronic communities. IEEE Transactions on Knowledge and Data Engineering **16**(7) (2004)
25. Aberer, K., Despotovic, Z.: Managing trust in a peer-2- peer information system. In: Proceedings of the 10th International Conference on Information and Knowledge Management (2001 ACM CIKM). (2001)
26. Despotovic, Z., Aberer, K.: Maximum likelihood estimation of peers performance in P2P networks. In: The Second Workshop on the Economics of Peer-to-Peer Systems. (2004)
27. Patel, J., Teacy, W.L., Jennings, N.R., Luck, M.: A probabilistic trust model for handling inaccurate reputation sources. In: Proceedings of Trust Management: Third International Conference (iTrust 2005), Springer-Verlag (2005) Volume 3477 of LNCS.
28. 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.
29. Kutvonen, L.: Automated management of inter-organizational applications. In: 6th International Enterprise Distributed Object Computing Conference (EDOC). (2002)
30. Gudgin, M., Nadalin, A.: Web Services Trust Language (WS-Trust). (2005)
31. Kaler, C., Nadalin, A.: Web Services Federation Language (WS-Federation), Version 1.0. (2003)