

An Integration of Reputation-based and Policy-based Trust Management

* Piero Bonatti¹, Claudiu Duma², Daniel Olmedilla³, and Nahid Shahmehri²

¹ Università di Napoli Federico II, Napoli, Italy
bonatti@na.infn.it

² Department of Computer and Information Science, Linköpings universitet
{cladu,nahsh}@ida.liu.se

³ L3S Research Center and University of Hannover, Hanover, Germany
olmedilla@l3s.de

Abstract. Trust management is currently being tackled from two different perspectives: a “strong and crisp” approach, where decisions are founded on logical rules and verifiable properties encoded in digital credentials, and a “soft and social” approach, based on reputation measures gathered and shared by a distributed community. We analyze the differences between the two models of trust and argue that an integrated approach would improve significantly trust management systems. We support our claim with real world scenarios and illustrate how the two models are integrated in PROTUNE, the core policy specification language of the network of excellence REVERSE.

1 Introduction

Trust management has been an important research line in the development of modern open distributed and decentralized systems. Trust has been studied in the context of decentralized access control [5, 16], public key certification [4, 9], and reputation systems for P2P networks [2, 14, 10].

There exist currently two different major approaches for managing trust: policy-based and reputation-based trust management. The two approaches have been developed within the context of different environments and targeting different requirements. On the one hand, policy-based trust relies on objective “strong security” mechanisms such as signed certificates and trusted certification authorities (CA hereafter) in order to regulate the access of users to services. Moreover, the access decision is usually based on mechanisms with well defined semantics (e.g., logic programming) providing strong verification and analysis support. The result of such a policy-based trust management approach usually consists of a binary decision according to which the requester is trusted or not, and thus the service (or resource) is allowed or denied. On the other hand, reputation-based trust relies on a “soft computational” approach to the problem of trust. In this case, trust is typically computed from local experiences together with the feedback given by other entities in the network (e.g., users who have used services of that provider). For instance, in eBay buyers and sellers rate each other after each transaction. The ratings pertaining to a certain seller (or buyer) are aggregated by the eBay’s reputation system into a number reflecting

* In alphabetical order. This work is partially supported by the Network of Excellence REVERSE, IST-506779, <http://reverse.net>.

seller (or buyer) trustworthiness as seen by the eBay community. The reputation-based approach has been favored for environments, such as Peer-to-Peer or Semantic Web, where the existence of certifying authorities could not be always assumed but where a large pool of individual user ratings was usually available.

The two trust management approaches address the same problem - establishing trust among interacting parties in distributed and decentralized systems. However, they assume different settings. While the policy based approach has been developed within the context of structured organizational environments, the reputation systems have been proposed to address the unstructured user community. Consequently, they assume different sources for trust (CAs and community opinion, respectively) and accordingly employ different mechanisms. Due to this, in the past years, researchers have targeted scenarios focusing on requirements which they could address with only one of these approaches. However, real life scenarios are not split in a way that they can just fit one of these approaches and in many cases, a mixed approach is required. For example, users might be interested in knowing whether a provider has a certificate from a CA but also in experiences other users had in the past while performing transactions with it. In addition, a seller might be interested in protecting an item on sale in different ways depending on the value of the item: based on reputation if the price is of a few euros (e.g., a T-shirt) or based on policies if it is of thousands (e.g., requiring a credit card for a flight).

Therefore, in this paper we propose the integration of policy based and reputation based approaches into a versatile trust management language capable of addressing both the structured organizational environments as well as the unstructured user communities. By combining the two different approaches, our integrated trust mechanism enhances the properties of the existing trust management tools.

2 Policy based vs. Reputation based Trust Management

The term *trust management*, introduced in [5] as “a unified approach to specifying and interpreting security policies, credentials, and relationships which allow direct authorization of security-critical actions”, has been given later a broader definition, not limited to authorizations [12]: “Trust management is the activity of collecting, encoding, analyzing and presenting evidence relating to competence, honesty, security or dependability with the purpose of making assessments and decisions regarding trust relationships”. Two main approaches are currently available for managing trust: policy-based and reputation-based trust management.

2.1 Policy-based Trust Management

This approach has been proposed in the context of open and distributed services architectures [6, 15, 11, 7] as a solution to the problem of authorization and access control in open systems. The focus here is on trust management mechanisms employing different policy languages and engines for specifying and reasoning on rules for trust establishment. The goal is to determine whether or not an unknown user can be trusted, based on a set of credentials and a set of policies.

In addition, it is possible to formalize trust and risk within rule-based policy languages [18, 13] in terms of logical formulae that may occur in rule bodies.

Currently, policy-based trust is typically involved in access control decisions. Declarative policies are very well suited to specifying access control conditions that are eventually meant to yield a boolean decision (the requested resource is either granted or denied). Systems enforcing policy based trust typically use languages with well-defined semantics and make decisions based on “non-subjective” attributes (e.g., requester’s age or address) which might be certified by certification authorities (e.g., via digital credentials). In general, policy-based trust is intended for systems with strong protection requirements, for systems whose behavior is guided by complex rules and/or must be easily changeable, as well as for systems where the nature of the information used in the authorization process is exact.

2.2 Reputation-based Trust Management

This approach has emerged in the context of electronic commerce systems, e.g. eBay. In distributed settings, reputation-based approaches have been proposed for managing trust in public key certificates, in P2P systems, mobile ad-hoc networks, and, very recently, in the Semantic Web. The focus here is on trust computation models capable to estimate the degree of trust that can be invested in a certain party based on the history of its past behavior.

The main issues characterizing the reputation systems are the trust metric (how to model and compute the trust) and the management of reputation data (how to securely and efficiently retrieve the data required by the trust computation) [3].

Marsh [17] made one of the early attempts at formalizing trust using simple trust metrics based on linear equations. This model has been further extended by Abdul-Rahman and Hailes to address reputation-based trust in virtual communities [1]. A number of reputation mechanisms for P2P systems, such as [3, 14, 10], followed similar trust and reputation models.

Typically, reputation-based trust is used in distributed networks where a system only has a limited view of the information in the whole network. New trust relationships are inferred based on the available information (following the idea of exploiting world’s information). In these scenarios, the available information is based on the recommendations and the experiences of other users, and it is typically not signed by certification authorities but (possibly) self-signed by the source of the statement. This approach supports trust estimates with a wide, continuum range and allows the propagation of trust (e.g., transitive propagation) along the network as well as weighting of values (e.g., fresher information vs. older information).

3 Integrated View of Trust Management

As described in previous sections, policy-based and reputation-based trust management address the same problem but from different perspectives. However, these points of view are not always just black or white and in many cases it would be desirable to combine them. In this section we propose an approach in which both of them can be integrated, based on the policy language PROTUNE [7].

First, reputation-based trust can be formalized by relations between *trustors*, *trustees*, *actions*, and *trust levels* [18]. For instance, a fact like

$$\text{trust}(P, S, \text{diagnosis}(\text{viral}), 80-100)$$

would model the fact that patient P trusts specialist S on diagnosis of viral diseases with an estimated confidence level belonging to the interval $80 - 100$.

Such trust statements can be the basis for trust propagation (e.g. via rules such as “trust X as a bike mechanic if X is trusted as a car mechanic”), and for access control decisions such as

```
allow(download(contents/pre_release)) ←  
    user( $X$ ),  
    trust(self,  $X$ , download(contents/pre_release), 90–100).
```

Such decisions may consider a notion of *risk*, as in

```
trust(ProgramX, Server, storeData(Server), 80–100) ←  
    Server.owner:CoXYZ,  
    risk(fail(Server), 0–0.1).
```

These examples (taken from [18]) show how trust and recommendations can be modelled and applied through a small set of predicates. The problem is: How should the basic *facts* about trust and risk be gathered and maintained?

In some cases, such facts can be defined by standard policy rules, for example:

```
trust( $A$ ,  $B$ , download(file), 80–100) ←  
    credential( $X$ , VISA),  
     $X.type : credit\_card$ ,  $X.owner : B$ .
```

However, the main current approaches are based on numerical models (see [8] for an extensive illustration of the main approaches) and ad-hoc algorithms for gathering, processing, and propagating historical data about past interactions and the resulting trust measures. In perspective, it may be possible to apply probabilistic, possibilistic or annotated logics to handle such numbers, but so far there is no clear indication that this is the right direction, nor any hint on how to do it.

In many approaches, the trust relationships we used as facts (not the inferred ones) are computed automatically based on experience and on the declarations of other users, using a numerical model. On the contrary, in policy based trust, all trust relationships are declared manually (e.g. an entity trusts another entity explicitly creating a statement in FOAF).

We argue that policy based decisions can be enhanced by numerical-based ones and viceversa. For example, in a policy where we protect our credit card, we could think of a policy like the following:

```
allow(visaCard) ←  
    credential(member(Requester), bbb),  
    trust(self, Requester, buying,  $X$ ),  $X > 0.8$ .
```

specifying that we will give our credit card only to entities that are certified by the Better Business Bureau (company that certifies that a company behaves according to the policies it published) and only if the server has a good reputation (this value is extracted from our personal experiences or inferred by using a reputation based algorithm on the community).

Further difficulties are: (i) data are application dependent, as well as the procedures for obtaining them; (ii) trust is a dynamic concept, i.e., it changes over time.

The above difficulties suggest a modular approach, namely, the computation and distribution of the basic facts on reputation and risk are delegated to suitable external packages. The results of their processing can be imported via HERMES-like [19] predicates such as

```
in(trust( $X$ ,  $Y$ ,  $A$ ,  $L$ ), reputation_pkg : eval_trust())
```

(more details available in [7]). In the above examples the functions `eval_trust()` wrap queries to the underlying reputation management algorithms, whatever they are. The wrapper collects and return the results of those subsystems as a set of terms matching the first argument of the `in` predicate. Then non-rule-based reputation and risk models can be integrated in policies without any ad-hoc language primitives. Moreover, the semantics of the `in` predicate depends on a time dependent state [7, 19], and this makes it possible to address the dynamic aspects of reputation.

Another advantage of this approach is that a single policy may simultaneously apply different approaches to reputation simply by invoking different packages and combining their results with suitable rules. This kind of flexibility is particularly important in a stage where it is not yet clear which of the competing models of reputation-based trust will become widely accepted, and which application domains they will prove to be good for. It is also possible to change the number and type of parameters of the `trust` and `risk` predicates, if needed by a particular reputation model.

This flexible architecture is compatible both with on-demand trust computation and with proactive propagation of trust evaluation, as reputation packages may receive asynchronous messages from other peers, concerning warnings and reputation evaluations.

3.1 An Application Scenario: Electronic Business

Transaction policies must handle expenses of all magnitudes, from micropayments (e.g. a few cents for a song downloaded to your iPod) to credit card payments of a thousand euros (e.g. for a plane ticket) or even more. The cost of the traded goods or services typically contributes to determining the risk associated to the transaction and hence the trust needed for performing it. For instance, for micro-payments of a few euros or cents, a seller could just check the reputation of the buyer within the community. If the buyer's reputation is high, the risk that he or she would not pay is very low, and thus the transaction can be conducted with a simple check. On the contrary, if a buyer's reputation is low or the amount of money involved in the transaction is high, risk is higher and thus the seller may require stronger guarantees, such as a verified credit card number to ensure that the buyer can and will pay.

The buyer's point of view is dual. If the amount of the transaction is high, the buyer may require strong and objective guarantees that the seller will deliver the goods and that the credit card will not be misused. For example, the buyer may require a secure connection, BBB (Better Business Bureau) certificates, blacklist checks, etc. In addition, the buyer may consider the seller's reputation in the community to increase the chances of successful transaction completion and privacy protection.

4 Conclusions

In this paper we have identified the advantages and limitations of policy-based and reputation-based trust management and described how the two approaches can improve each other. The need for an integrated approach has been motivated with real world scenarios. We proposed an integrated trust management approach that combines rule-based and credential-based trust with numerical trust estimates based on a large number of sources (e.g., user community). Our formalization privileges flexibility and extendibility as a design goal. The extension of the traditional crisp, boolean policies with a continuum range

of trust levels, and the extension of numerical trust models with well defined trust combination and propagation rules, yield a versatile trust management framework capable of addressing the complexity and the variety of semantic web scenarios, involving both structured organizational environments and unstructured user communities.

References

1. A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Proceedings of 33rd Hawaii International Conference on System Sciences*, 2000.
2. K. Aberer. P-grid: A self-organizing access structure for p2p information systems. In *Proceedings of Ninth International Conference on Cooperative Information Systems*, 2001.
3. K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *Proc. of 10th International Conference on Information and Knowledge Management*, 2001.
4. T. Beth, M. Borcherding, and B. Klein. Valuation of trust in open networks. In *Proc. of the 3rd European Symposium on Research in Computer Security*. Springer-Verlag, 1994.
5. M. Blaze, J. Feigenbaum, and J.Lacy. Decentralized trust management. In *Proceedings of IEEE Conference on Security and Privacy*, 1996.
6. P. Bonatti and P. Samarati. Regulating service access and information release on the web. In *Proc. of the 7th ACM conference on computer and communications security*, 2000.
7. P. A. Bonatti and D. Olmedilla. Driving and monitoring provisional trust negotiation with metapolicies. In *IEEE 6th International Workshop on Policies for Distributed Systems and Networks (POLICY)*, pages 14–23, Stockholm, Sweden, jun 2005. IEEE Computer Society.
8. P. A. Bonatti, N. Shahmehri, C. Duma, D. Olmedilla, W. Nejdl, M. Baldoni, C. Baroglio, A. Martelli, V. Patti, P. Coraggio, G. Antoniou, J. Peer, and N. E. Fuchs. Rule-based policy specification: State of the art and future work. Report I2:D1, EU NoE REVERSE, sep 2004.
9. G. Caronni. Walking the web of trust. In *Proceedings of 9th IEEE International Workshops on Enabling Technologies (WETICE)*, pages 153–158, June 2000.
10. C. Duma, N. Shahmehri, and G. Caronni. Dynamic trust metrics for peer-to-peer systems. In *Proc. of 2nd IEEE Workshop on P2P Data Management, Security and Trust*, August 2005.
11. R. Gavriloaie, W. Nejdl, D. Olmedilla, K. E. Seamons, and M. Winslett. No registration needed: How to use declarative policies and negotiation to access sensitive resources on the semantic web. In *1st European Semantic Web Symposium (ESWS 2004)*, pages 342–356, Heraklion, Crete, Greece, may 2004. Springer.
12. T. Grandison. *Trust Management for Internet Applications*. PhD thesis, Imperial College London, 2003.
13. T. Grandison and M. Sloman. Specifying and analysing trust for internet applications. In *Towards The Knowledge Society: eCommerce, eBusiness, and eGovernment, The Second IFIP Conference on E-Commerce, E-Business, E-Government (I3E 2002)*, IFIP Conference Proceedings, pages 145–157, Lisbon, Portugal, oct 2002. Kluwer.
14. S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. Eigenrep: Reputation management in p2p networks. In *Proc. of 12th International WWW Conference*, pages 640–651, 2003.
15. N. Li and J. Mitchell. RT: A Role-based Trust-management Framework. In *DARPA Information Survivability Conference and Exposition (DISCEX)*, Washington, D.C., Apr. 2003.
16. N. Li and J. C. Mitchell. Datalog with Constraints: A Foundation for Trust-management Languages. In *Proceedings of the Fifth International Symposium on Practical Aspects of Declarative Languages (PADL 2003)*, pages 58–73, January 2003.
17. S. Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, Uni. of Stirling, 1994.
18. S. Staab, B. K. Bhargava, L. Lilien, A. Rosenthal, M. Winslett, M. Sloman, T. S. Dillon, E. Chang, F. K. Hussain, W. Nejdl, D. Olmedilla, and V. Kashyap. The pudding of trust. *IEEE Intelligent Systems*, 19(5):74–88, 2004.
19. V. Subrahmanian, S. Adali, A. Brink, J. Lu, A. Rajput, T. Rogers, R. Ross, and C. Ward. *HERMES: Heterogeneous reasoning and mediator system*. 1995.