Rule Languages for Security and Privacy in Cooperative Systems

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

The open nature of modern network applications is potentially an excellent support to cooperative work of all sorts, and at the same time a majour source of security and privacy concerns.

It is now commonly recognized that traditional authentication techniques do not scale to the new open scenarios, and are not suitable for the dynamic nature of virtual organizations. Research is focussing on more flexible *trust management* models that let two or more peers interact securely even if they have never interacted with each other before and have no previous knowledge about each other's properties and protection requirements (see [14, 3, 7, 13, 5, 16, 8, 11, 15, 4, 9, 2, 17, 10], to name a few). This situation occurs whenever a new virtual organization is created or extended.

Trust management is based on electronic credentials, that constitute a flexible way of representing properties of individuals, groups and organizations. In principle, any entity can issue its own credentials, thereby signing some statements about some entities. Other peers may decide whether that entity can be trusted on those statements, in the framework of a specific task. Credentials are verifiable and unforgeable, so it can be robustly checked whether a given statement has been actually signed (hence issued) by a specific peer.

In this way, each partner of a virtual organization may state properties about itself and its members, and these properties can be used to make access control decisions. The policy of partner X may then specify permissions for partner Y's members by referring to what member Y says about them. Sophisticated delegation mechanisms are possible. In a similar scenario, security-related decisions can be based on rules formulated by and stored on different parties [2, 12, 6, 1].

The properties encoded in credentials do not necessarily suffice to identify users, therefore credentials are also an excellent way of specifying partial or multiple identities, and hence an excellent way of protecting privacy while disclosing the information required to access a resource.

However, balancing information disclosure and access permissions is not a trivial task. In order to minimize information disclosure, a peer should know under which conditions the other peers are going to share their resources; often, however, the other peers are not publishing all such conditions, because in general the policies themselves are sensitive resources. Access may be denied simply because the peers are not showing enough of their properties and their policies to each other [17]. Suitable credential negotiation strategies have to be adopted and implemented to minimize this problem.

This is not the only issue to be solved. Credential-based techniques should be brought to their full potential by spreading them in everyday applications, and by making them accessible to most users.[1] [2]

Common users, with no technical training in computer science or programming, should be able to formulate their own policies and understand their actual behavior. This goal requires high level specification and query languages, easy to learn and understand.

As credential exchange and access control decisions become more and more automatic, the need for *explanation* mechanisms increases. When a given request is denied, the user may still be able to obtain the desired service. Sometimes it is enough to interact with the system in a different way. Sometimes a particular procedure (such as a registration procedure) can be activated. Since in our reference scenarios peers often interact with each other only for a short time and we want to minimize the effort for integrating the procedures of the two peers, automated hints on how to get the required permissions are absolutely crucial to make the composite system really cooperative at a low cost.

Ideally, credential negotiation and access control should require no more than the policies formulated by the users. Then the system should automatically make a variety of decisions (on which credential to ask next, on granting the requested service, etc.) based on a single body of high-level requirements.

[1] http://rewerse.net/i2/

[2] http://cs.na.infn.it/rewerse/I2_FLYER_print.PDF

Proceedings of the 29th Annual International Computer Software and Applications Conference (COMPSAC'05)
0730-3157/05 $20.00 © 2005 IEEE

In order to reach these objectives, declarative rule-based languages are extremely promising as policy specification languages. Rules constitute a flexible and natural way of expressing a variety of policies, ranging from security policies to reputation management and business rules. Rules have clean and simple semantics, which helps in keeping complex policies under control and reduce the cost of maintaining policies. Rules constitute an explicit body of knowledge, well suited to automatically generating explanations of system decisions. Moreover, rules are relatively simple to formulate, and require no programming abilities. A controlled natural language front-end can be adopted to further enhance usability. Last but not least, a single rule base can be used in different ways depending on the current task (credential negotiation, access control, explanations, etc.) We argue that the language for tuning credential negotiation strategies to specific application scenarios in a declarative way should be of the same nature.

References

[1] Jim Basney, Wolfgang Nejdl, Daniel Olmedilla, Von Welch, and Marianne Winslett. Negotiating trust on the grid. In *2nd Workshop on Semantics in P2P and Grid Computing*, New York, May 2004.

[2] M. Y. Becker and P. Sewell. Cassandra: distributed access control policies with tunable expressiveness. In *5th IEEE International Workshop on Policies for Distributed Systems and Networks*, Yorktown Heights, June 2004.

[3] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. KeyNote: Trust management for public-key infrastructures (position paper). *Lecture Notes in Computer Science*, 1550:59–63, 1999.

[4] P. Bonatti and P. Samarati. Regulating Service Access and Information Release on the Web. In *Conference on Computer and Communications Security*, Athens, November 2000.

[5] Jan Camenisch and Els Van Herreweghen. Design and implementation of the *idemix* anonymous credential system. In *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*, pages 21–30, 2002.

[6] Rita Gavriloaie, Wolfgang Nejdl, Daniel Olmedilla, Kent Seamons, and Marianne Winslett. No registration needed: How to use declarative policies and negotiation to access sensitive resources on the semantic web. In *1st First European Semantic Web Symposium*, Heraklion, Greece, May 2004.

[7] A. Herzberg, Y. Mass, J. Michaeli, D. Naor, and Y. Ravid. Access control meets public key infrastructure, or: assigning roles to strangers. In *IEEE Symposium on Security and Privacy*, May 2000.

[8] A. Hess, J. Jacobson, H. Mills, R. Wamsley, K. E. Seamons, and B. Smith. Advanced client/server authentication in TLS. In *Network and Distributed Systems Security Symposium*, Feb. 2002.

[9] T. Jim. SD3: A Trust Management System With Certified Evaluation. In *IEEE Symposium on Security and Privacy*, Oakland, CA, May 2001.

[10] H. Koshutanski and F. Massacci. Interactive trust management and negotiation scheme. In *2nd International Workshop on Formal Aspects in Security and Trust (FAST)*, pages 139–152, Aug. 2004.

[11] N. Li and J. Mitchell. RT: A role-based trust-management framework. In *Third DARPA Information Survivability Conference and Exposition*, Apr. 2003.

[12] N. Li, W. Winsborough, and J.C. Mitchell. Distributed Credential Chain Discovery in Trust Management. *Journal of Computer Security*, 11(1), February 2003.

[13] B. Pfitzmann and M. Waidner. Federated identity-management protocols–where user authentication protocols may go. In *11th Cambridge International Workshop on Security Protocols*, Apr. 2003.

[14] S. Staab, B. 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 Journal*, 19(5):74–88, Sep/Oct 2004.

[15] Lingyu Wang, Duminda Wijesekera, and Sushil Jajodia. A logic-based framework for attribute based access control. In *2nd ACM Workshop on Formal Methods in Security Engineering (FMSE 2004)*, pages 45–55, Oct. 2004.

[16] M. Winslett, T. Yu, K. E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. The TrustBuilder architecture for trust negotiation. *IEEE Internet Computing*, 6(6):30–37, Nov./Dec. 2002.

[17] T. Yu, M. Winslett, and K. E. Seamons. Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation. *ACM Transactions on Information and System Security*, 6(1), Feb. 2003.