



Identity Management

Duncan A. Buell
University of South Carolina

Ravi Sandhu
George Mason University
and NSD Security

Identity management has recently emerged as a critical foundation for realizing the Internet's business benefits in terms of cost savings, management control, operational efficiency, and most importantly, business growth. Enterprises need to manage access to information and applications scattered across a wide range of internal and external computing systems. Moreover, they must provide this access for a growing number of users, both inside and outside the organization, without compromising security or exposing sensitive information. Managing multiple versions of users' identities across multiple legacy applications makes the task even more daunting.

For this issue of *IEEE Internet Computing*, we invited researchers and practitioners to submit articles describing all aspects of identity management technology and practice. Together, the articles present both background and some in-depth coverage of some fundamental research questions that must be addressed.

Technical Hurdles

Among the major challenges to effective identity management is controlling information when the entities that need to access it are dispersed and highly diverse. The field also faces some standard technical questions, including how to control access

to information in databases. The best answers to such issues often depend on whether the requestor is someone inside or outside the organization. Other questions relate to cryptography and associated protocols. We must be able to verify an electronic sign-on's authenticity, for example, and digital signatures must be applicable when nonrepudiation is required.

As the articles in this issue show, many of the challenges to identity management come from a desire to grant single-sign-on access to a collection of resources that might well have different, even contradictory, access-protection rules. Thus, the single-sign-on mechanism must also permit that access-control to function as it would in a multiple-sign-on environment.

In addition to purely technical issues, we must consider the human factors in any e-commerce situation. Some compare to existing behaviors in traditional situations, but others are new to the domain of electronic transactions. These factors differ according to the transacting entities' natures.

The standards and practices under which business-to-business (B2B) transactions operate — usually between a limited number of entities at different enterprises — are quite different from those needed for electronic retailing, in which customers appear at random.

These situations also differ from those involving organizations' internal information needs; although we can plan to grant access to information only as needed, codifying such access rules can be difficult because individuals often have multiple responsibilities.

Furthermore, as legislation and public policy evolve, new technical issues arise and expectations increase; if electronic transactions are regulated by law, then users expect those transactions to be conducted in accordance with the regulations. B2B transactions are usually subject to audits, for example, and many database transactions that involve personal data, such as financial or health records, are increasingly subject to governmental regulation. Not even individual retail transactions are entirely unregulated. Finally, certain situations dictate that an electronic identity remain anonymous. That is, we must be able to verify a specific identity, but the only attribute information that should be transferred is its validity.

Because the legal status of electronic data and transactions is in a rapid state of flux, we elected not to delve too deeply into the public policy questions that surround identity management. In the US, for example, the Graham-Leach-Bliley Act and the Health Insurance Portability and Accountability Act of 1996 (HIPAA) have mandated certain levels of accountability with regard to the privacy of electronic financial and health data, but laws and legal precedents are different around the world. In many cases, we still lack operational definitions for standards of behavior because the problems are new, jurisdictions are different, and an adequate set of test cases has yet to appear on which to base accepted practice.

About this Issue

In response to the call for this issue, we received a significant number of submissions addressing a range of issues as broad as the topic itself. We selected the following three articles for publication because they represent the theme well, presenting both background and relevant research results.

In "Managing Multiple and Dependable Identities," Damiani, De Capitani di Vimercati, and Samarati discuss an approach to controlling multiple robust identities in an electronic world, a crucial issue in developing the next generation of distributed applications. As described here, a *digital identity* is the electronic representation of an individual's or organization's sensitive information. A system that manages identity must, of course, be reliable. It must also permit the user to control the

disclosure of information associated with the digital identity, and it must move with the user rather than be fixed in location. To a great extent, identity-management solutions have required the creation and use of trusted third parties as the authority or intermediary. Despite some differences, this is not unlike the authority often assumed as part of a public key infrastructure. Three protocols now describe how to establish such third-party intermediaries in single-sign-on authentication: Microsoft's .NET Passport, Oasis's Security Assertions Markup Language (SAML), and the Liberty protocol.

Many challenges come from the desire to grant single-sign-on access to collections of resources that might have contradictory access-protection rules.

In the second article, "Analysis of Liberty Single-Sign-on with Enabled Clients," Pfitzmann and Waidner examine a (subsequently fixed) flaw they discovered in the Liberty protocol, and discuss the general nature of third-party authentication methods. In a scheme like the one the Liberty Alliance has proposed, users each sign on with *identity providers*, which then authenticate them to subsequent services. The Liberty protocol lets clients remain unaware of the cryptographic specifics, other than that secure channels (such as SSL) will be used in the usual way. A disadvantage of secure-channel protocols such as Liberty is that the third party's intervention with asymmetric encryption is slower than direct communication using symmetric-key encryption. In some situations, notably intra-enterprise transactions, we could use the concept of "enabled clients" (as described in the article) to take advantage of existing, dependable information to improve efficiency without compromising security or authenticity.

Finally, Skogsrud, Benatallah, and Casati describe their proposed trust negotiation framework in "Model-Driven Trust Negotiation for Web Services." In their approach, trust negotiations involve signed assertions regarding the owners' attributes rather than the outright transfer of requesters' identities. Historically, trust management systems have been difficult to get right; they have also been difficult to adapt in changing environments that involve enterprise policies and reg-

ulations. This article describes a state-based language for trust management that offers a potential solution to these problems.

Future Work

Security professionals often view security objectives as “keeping the bad guys out” on one hand and “letting the good guys in” on the other. Keeping the “bad guys” out is important and has received considerable attention in recent years. It makes for good press and cops-and-robbers stories, and it is perhaps the first objective that should be addressed. All the same, organizations can obtain true productivity gains only by letting the “good guys” in. Identity management is the cornerstone technology for achieving this goal. As such, it is likely to remain a pressing issue for many years to come.

Traditionally, identity management has been concerned with managing an organization’s employees to ensure that their authentication and authorization information is consistent and up to date within the organization’s information systems. This traditional arena continues to pose many challenges for security architects and designers, especially given the large base of legacy systems. However, the true value of identity management comes into play with business partners and consumers. The ability to federate identity across organizations while maintaining clear trust, liability, and cost responsibilities is a major challenge for enterprises as we continue to pursue efficiency and cost savings in cross-organizational business and customer-relationship processes.

There are many research and development challenges to address before seamless identity management becomes a reality. We need stan-

dards to build trust, authority, and policy relationships across organizational boundaries. First, end consumers need assurances regarding privacy of sensitive information, particularly given the prevalence of “identity theft.”

We must further elaborate the interplay between authentication and authorization rather than following the classical approach and treating them as orthogonal issues. We must also refine existing access-control models to reflect the obligations on the provider and consumer of identities in multiparty transactions. We need lightweight and user-transparent protocols with zero, or very small, footprints to run on end users’ machines. At the same time, we need to look ahead to emerging client-side platforms that permit some degree of trust in the end systems themselves. We are just beginning to come to grips with these issues. They should ensure interesting research problems and scenarios for many years to come. □

Duncan A. Buell is professor in and chair of the Department of Computer Science and Engineering at the University of South Carolina. His research interests include high-performance computing, reconfigurable computing. He received a PhD in mathematics from the University of Illinois at Chicago. He is a member of the ACM and a senior member of the IEEE. Contact him at buell@cse.sc.edu.

Ravi Sandhu is professor of information security and assurance at George Mason University, and Chief Scientist at NSD Security. His research interests are primarily in authorization, authentication and access control. He received a PhD in computer science from Rutgers University. He is a fellow of IEEE and a fellow of ACM, and serves on *IEEE Internet Computing’s* editorial board. Contact him at sandhu@gmu.edu.

SET
INDUSTRY
STANDARDS

IEEE Computer Society members work together to define standards like
IEEE 802, 1003, 1394, 1284, and many more.

HELP SHAPE FUTURE TECHNOLOGIES • JOIN AN IEEE COMPUTER SOCIETY STANDARDS WORKING GROUP AT

computer.org/standards/

wireless networks
gigabit Ethernet
enhanced parallel ports
802.11 FireWire
token rings