The Technology of Trust

Ravi Sandhu
SingleSignOn.net
and George Mason University

Assuring adequate security in cyberspace is one of the more challenging problems we face as we try to leverage the Internet's power to increase productivity and provide competitive advantage. Security has never been an easy goal. It requires tangible expense to implement while its payback is less certain. Return on investment is notoriously difficult to calculate and justify. Security practitioners understand there is no such thing as absolute security. Every system has real vulnerabilities and is likely over time to suffer some security breaches. A rapidly changing threat environment and steady emergence of new applications and connectivity make it hard to assess the state of security at any given moment. Moreover, security has a paradoxical quality: The more we have of it, the more of it we desire. Thus, the key word here is "adequate."

Consider, as a rough analogy, health care.

Even if our health is good, it can still be improved, right? If our personal lifestyles and habits are healthy, we can still get sick. There's no shortage of diseases to treat. We can even proactively prevent disease rather than reactively try to cure it.

No matter how secure our systems, they can still be more secure. No matter how much we protect our systems, they are still vulnerable to attack. Many security problems, both known and unknown, exist. We can spot-treat them (patches), prevent them using firewalls and similar technologies, or cure them (virus checkers or fixers) and the more complex our systems, the more opportunities for holes or attack.

Security and Trust

Think of security as comprising two disciplines: *security hygiene* and *security as business enabler*. Security hygiene tends to dominate discussion today. We deploy firewalls, virus checkers, and security patches and obtain cooperation from users and vendors to support security goals. Good security hygiene is worth pursuing, but it cannot be information security's primary goal. Security professionals must also view, research, develop, and sell security as a business enabler. This requires a fundamental reorientation for most of us.

 1089-7801/02/$17.00©2002 IEEE

Security has classically developed in the context of a single organization seeking to control its employees' access to its information resources. The need for security that cuts across multiple organizations has emerged in tandem with the Internet's growth. The upshot is that what we once regarded as an esoteric, high-end piece of the security puzzle has become the dominant driving force.

Taking primarily inward-focused security technologies and extending them outward is not easy. Establishing trust across organizational boundaries is a qualitatively different problem from establishing it within an organization because different organizations have different interests and objectives. Even when engaged in a cooperative enterprise, their interests can be contradictory and competing. So, how do we articulate these concerns and reconcile conflicting objectives to arrive at adequate security?

Trust technology involves many security issues. To build trust in cyberspace, we must consider software, systems, people, computers, processes, and so on. We must also consider transference of trust from one entity to another. Likewise, traceability of trust across different entities is important. Ultimately, trust is a human and social issue, so we must also factor in laws, regulations, and policy.

We are far from understanding how to effectively engineer systems that deliver adequate trust to users. Nevertheless, we have made considerable progress and can expect this field to be more sharply defined over the coming years.

The Articles

This timely issue presents three articles on the general theme of the technology of trust. Two of the articles look at the technology of information security. The third focuses on human factors issues, which are extremely important for security but are all too often overlooked or poorly engineered. These three pieces provide an excellent cross-section of current concerns and research in the trust arena.

"Negotiating Trust on the Web" by Marianne Winslett and her colleagues identifies trust negotiation as an iterative exchange of digital credentials and requests for credentials between two parties seeking to establish mutual trust to conduct a transaction. The authors propose an approach for automated establishment of such trust between entities in different security domains by means of cryptographically signed credentials used like letters of introduction. Because these credentials can contain sensitive information, they are protected like other sensitive resources. Trust is established incrementally, through a bilateral iterative exchange of credentials.

David Scott and Richard Sharp's "Framework for Secure Web Application Development" describes a set of tools developed at Cambridge University to secure Web applications, especially those including third-party components. The authors identify a set of the most common attacks on Web-based applications such as form modification, SQL attacks, cross-site scripting, and control-flow tampering. They describe how their Spectre family of tools can help secure applications from such attacks. Web developers specify policies in the Security Policy Description Language, from which the Spectre tools generate wrappers deployed in an application-layer firewall to enforce the policies. Using wrapper software to protect legacy code is a proven technology in the security community. The authors take it one step further by also providing a mechanism to generate appropriate wrapper software for Web applications.

Andrew Patrick's "Building Trustworthy Software Agents" deals with the softer (but no less important) human side of security technology. There is considerable interest these days in deploying agents to make the Internet a more friendly and productive resource for individual users. Many people have written about the technical issues of securing such agents and their behavior in this context. However, this is one of the first articles to focus attention on the perception of trust issues. The author builds on prior work about user attitudes to e-commerce transactions to develop a model of agent acceptance based on feelings of trust and perceptions of risk.

Ravi Sandhu is a professor of information and software engineering and director of the Laboratory for Information Security Technology at George Mason University. He is also the chief scientist and cofounder of SingleSignOn.net. He has published more than 150 technical papers on computer security in refereed journals, conference proceedings, and books. His technical interests include access control and authorization models and protocols. At SingleSignOn.net, he is the principal designer of the security protocols and role-based authorization in the Secure Identity Appliance product. Sandhu has BTech and MTech degrees in electrical engineering from the Indian Institute of Technology, Bombay and Delhi, and MS and PhD degrees in computer science from Rutgers University, New Jersey. He is a fellow of the ACM and the IEEE.

Readers can contact Sandhu at sandhu@gmu.edu.