

# AN IT ARCHITECTURE FOR NIMBLE ORGANIZATIONS: MANAGING ACCESS FROM CYBERSPACE

Raju Kocharekar

In today's world of "co-opetition" among organizations, where one organization competes and cooperates simultaneously with business partners, all organizations must have IT architectures that are nimble and flexible. This article takes a look at one important aspect in achieving a flexible IT architecture: access management of information resources over the Web. It describes an access management architecture that can change as business does and recommends how best to implement this nimble architecture.

**T**HE CONCEPT OF A VIRTUAL ORGANIZATION is not new.<sup>1</sup> As part of the business process reengineering effort, organizations had to think about their core competencies, and many outsourced those that did not add high value to the business. New business relationships emerged among suppliers, partners, and consumers. Reduced trade barriers across nations helped these organizational relationships become global, especially in manufacturing. Ever since, all forms of business relationships have been existent in all areas of corporate activities, from product design to operations to sales and client support. The virtual organization that exists today blurs the organization boundary, as it tries to be flexible and nimble to adapt to changing business and technology environments. This trend now continues with the second wave of global outsourcing in services.

As more business is conducted in cyberspace, how organizations provide access to information technology assets will be a key competency. However good the organization's IT assets are, they are of low value if the access mechanisms are broken. Organizations can

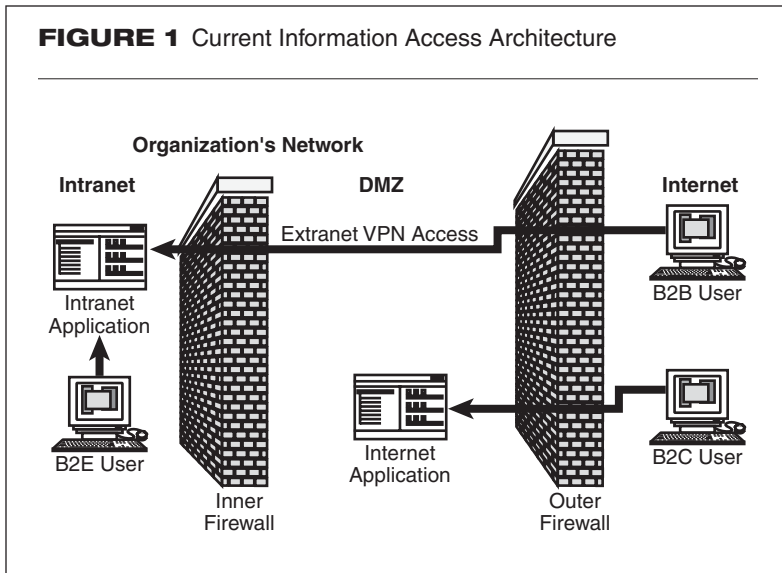
lose business opportunities if the access mechanism provides too little access to the assets. On the reverse side, overly broad access poses all kinds of risks, ranging from legal challenges to loss of competitive edge.

Access management to a virtual organization's information assets therefore requires a comprehensive architectural approach. This article articulates the current issues in access management for virtual organizations and the emerging architecture for addressing those issues. This architecture is pragmatic in recognizing current limitations in deployable technology but remains consistent with the future anticipated evolution, both in technology supply and business demands. The viewpoint articulated in this article is therefore primarily applicable for IS management and practitioners.

## ISSUES WITH THE ACCESS MANAGEMENT

Undoubtedly, information technology has been one of the major drivers behind the virtual organization.<sup>2</sup> It is ironic, however, that current IT architecture has not yet caught up with

*RAJU KOCHAREKAR is production manager for Internet Services Program in The World Bank's Information Solutions Group. He has more than 20 years of professional experience in IT management and has published several articles in academic and professional journals covering this subject. He can be contacted at rkocharekar@worldbank.org.*

**FIGURE 1** Current Information Access Architecture

organizational changes. It still heavily relies on security protection at the network layer in managing access to the organization's IT resources. It assumes that what is inside the network-level firewall is tightly protected and managed. It also assumes that internal employees always access the organization's information assets from inside the firewall. In reality, both are rapidly becoming false assumptions.

The scope of intranets has expanded globally through remote access, third-party Internet service providers, and other means, and now reaches the homes and other places where employees travel. Wireless access to an organization's network poses security issues of its own.<sup>3</sup> On the flip side, employees, business partners, and consumers are demanding access to organizational information from anywhere in the cyber and physical space.

As the need to provide access to organization information to a wider audience has grown, organizations hosted servers in a "demilitarized zone" (DMZ), outside the organization's intranet, where they copied sections of their applications, including back-end data sources. This created large administrative and resource overhead to maintain duplicate sets of information. As demand grew to provide direct transactional access to the business partners as well as employees in cyberspace, organizations set up virtual private networks (VPNs) to provide network-level authentication and encryption and bring network traffic directly inside the firewall.<sup>4</sup> These scenarios are depicted in Figure 1. This mechanism again relied on network-level access control. Users have to share the encryption keys and have access only from

the VPN endpoints. This solution is not scalable to support thousands of users, due to high administrative overheads. The solution provides no identification of the user accessing the VPN, and this is assumed to be done outside the network layer access control.

Organizations are now building portals<sup>5</sup> that access back-end systems and databases to aggregate and syndicate information and tailor it to users' roles and preferences. Front-end portal code sometimes has relied on special accounts when accessing back-end applications. This was because single sign-on (SSO) mechanisms were not available. Fortunately, products that offer SSO with session management are now plentiful. This, however, also requires identity and access provisioning management across multiple applications. Access management for a rapidly growing population of users to an increasing number of organizational IS assets needs central, externalized authorization that was traditionally implemented directly in applications.

Finally, portal application programs themselves must be authenticated and have access to communicate with back-end applications and data sources.

## A NEW ARCHITECTURE

Increasing business pressures to provide seamless information access and the inadequacies in the existing model have forced the use of a different approach for architecting access. This new model turns the organization's information access architecture upside down. It no longer first focuses on lower network-level access control. Instead, it focuses on users accessing the information first. This architecture's design considers users' need to access information and the level of security required. The architecture covers user identity and access provisioning management, authentication, and authorization management, starting from user definitions down to applications, servers, and networks.<sup>6</sup>

As one moves further down the network stack, the number of entities (i.e., the users and resources and their access relationships) explodes beyond a manageable level. This architecture therefore tries to reduce the number of identities and access relationships to manage at an acceptable level.

## User Segregation

The new model segregates users requiring access to the organizational information into four

**A**ccess provisioning products with delegated administration will soon obviate the need for the proxy approach.

different categories. These user categories map closer to the different types of users in a virtual organization model: Business-to-employee (B2E), business-to-business (B2B), business-to-consumer (B2C), and anonymous users.

**B2E Users.** These are the traditional staff and contractors working for an organization. As the bulk of an organization's information systems have been centered on access to these users for years, there is an ample number of existing policies and procedures for all aspects of this type of access management (i.e., the identity management, access provisioning, authentication, and authorizations of these users). These policies and procedures provide a baseline for extending the same to other types of users.

**B2B Users.** Organizations have various forms of relationships with other organizations, as suppliers, joint venture partners, or downstream value-add customers. As an organization takes its business relationships into cyberspace, individuals affiliated with partner organizations require online access to an organization's information. Today, prenegotiated offline agreements on policies and procedures are necessary between two organizations. Partner organizations and agreement attributes must be registered in an organization's enterprise directory. As part of these agreements, liaison or contact persons from the partner organizations are identified in managing other user accounts within those organizations. Their responsibilities include taking necessary actions or notifying the partner about changes in employee status that affect access controls. B2B user access management requires the use of products and procedures to manage workflow and delegated administration capabilities.

Also keep in mind that one organization's users require access to partners' information systems. An organization is obligated to cooperate in the management of user access to partner organizations' information assets. The organization must therefore record and keep track of the information access provided by partner organizations. As responsibilities or employment status changes, identity management systems must transmit information about these changes to partner organizations. Failing to do this could jeopardize business relations with the partner and expose an organization to legal challenges if any security breaches occur.

**B2C Users.** These are the users with whom an organization does business directly as individuals. They do not need any organization affiliation for this purpose. An obvious example is a consumer buying products over the Web, but the same applies for other cases such as job seekers or new suppliers submitting bids for E-procurement tendering. Procedures and policies to manage these users' access are certainly less burdensome than the previous two categories. Also, self-service registration processes are well established and off-the-shelf products exist that have embedded best practices.

**Anonymous Users.** This is the simplest form when dealing with access management. But even access by these users requires a minimum level of policies and procedures. With the deluge of available information, users accessing a site can get frustrated if they do not find the information for which they are looking. A prevalent technique is to track user access during a visit or multiple visits, and then enhance their experience by personalizing the information presented to them. At a minimum, an organization needs privacy policies for this level of access.<sup>7</sup> Also, an organization should want anonymous users to be aware of copyright and other regulations of the information a site provides.

Segregating users into these categories helps in building broad definitions and implementation of policies and procedures in identity and access management, within each user type. But it also requires procedures when users migrate from one category to the other. In such cases, the task is to preserve user identity and relevant profile information (personalization, etc.) while ensuring that the access controls match the new status.

#### **Users with Access Grant Privileges**

Almost orthogonal to the user segregation previously discussed is another special class of users worth mentioning. General users just have access to resources. Special users have privileges to grant access to resources to other users. In today's environment, these are typically B2E users, but that will not be the case for long. Also, systems and resource administrators sometimes act as proxy to business or application managers, when they grant access to other users. Access provisioning products with delegated administration will soon obviate the need for the proxy approach.

**T**oday, the right balance between central and application-level authorization models is specific to each organization's IT asset portfolio.

Because of the nature of access rights as well as their high affinity for IT resources and the users they target, the management of these access rights is complex and requires substantial investments in access provisioning policies and procedures. For example, what are the procedures when a user who has granted access rights to other users leaves an organization? Should all the rights given by this employee be reviewed and revalidated?

#### **A NEW IT ARCHITECTURE**

Once user segregation and the corresponding identity management processes are analyzed, an architecture for user access to an organization's IT assets can be designed. As mentioned, many access management products now offer SSO capability across multiple front-end and back-end applications and products. This facilitates straight-through processing with transaction integrity and fully auditable user interaction across multiple systems. It also substantially improves the user experience. SSO pushes toward a centralized infrastructure for user identity and access management across different applications. User account and access provisioning systems are also available; they can update user access rights across multiple systems with transactional integrity. They provide workflow and delegated administration capabilities that are critical for access provisioning. Finally, they themselves require and have to rely on base-level authentication and authorization services provided in access management products.

Justification for central authentication and account administration infrastructure is relatively simple. The difficult part is to comprehend and design the authorization model. Currently, the authorization process is buried deep in applications, because all the information needed to establish authorization is only available just prior to resource access. This approach is costly to maintain when authorization rules are changed. It raises concerns for auditors, especially when once validated, authorization modules are subsequently changed. Finally, it also makes access provisioning difficult to implement because it has to interface with many different applications.

Today, the right balance between central and application-level authorization models is specific to each organization's IT asset portfolio. Three different factors that dictate this balance are:

1. User attributes management
2. Information resource attributes management
3. Authorization policy management

**User Attributes Management.** Because authorization rules make use of user attributes, it is critical to determine how they are maintained and made available for rules processing. For example, user attributes, such as job grade, are maintained centrally by the Human Resources (HR) department. They would be maintained in HR databases and made available through a central enterprise user directory. Other attributes such as the reporting relationship may be left to line management to administer within local units. These attributes may be made available through virtual user directory products that process rules. The virtual directory product dynamically accesses departmental directories, where attribute information is stored. Both types of attributes can be used for authorization in different applications that span multiple business units. On the other hand, some user attributes are managed and used in a local departmental application domain, such as department collaborative applications. There is little cost justification for managing these attributes centrally.

**Information Resource Attributes Management.** Just like user attributes, authorization rules make use of information resource attributes. For example, resource attributes such as the product category are maintained centrally and then made available through the central enterprise resource directory. Other attributes can be managed locally, but made available through a virtual resource repository.

**Authorization Policy Management.** Either user or resource attributes can be made available through central physical or virtual repositories for policy evaluation. However, rules or policies that utilize attributes may be specific and within applications themselves, or they could be managed external to the applications. One big advantage of using a central infrastructure for authorization is the use of a common authorization rules engine. It is easier to manage authorization rules by just reconfiguring the rules rather than modifying application code. The rules engines also allow for quick validation and audit functions.

If authorization policies are devised and administered at the organizational level, central

**FIGURE 2** Access Management Matrix

	User Attributes Management	Information Resource Attributes Management	Authorization Policy Management
Update	Central or Local		→
Use	Central or Local		→

policy evaluation makes sense. For example, procurement or staff travel policies and functions can be centrally managed. Many times, the entire process and function are implemented in one monolithic ERP system, including its own authorization policy evaluation engine. In such cases, an organization must rely on the vendor to provide appropriate policy management and subsequent technology enhancements. In other cases, authorization policies can be decentralized to business units. In such cases, it could either be provided through central policy management infrastructure with decentralized policy administration features, or it could be provided through virtual policy processing. In the latter case, a virtual policy evaluation engine dynamically aggregates rules evaluation results from decentralized policy evaluation engines.

Figure 2 presents a matrix form that must be completed for each combination of user and information resource attributes and authorization management policy. The individual cells indicate if an activity (e.g., update of an attribute or its use in determining access) for that attribute is central across an organization or local to individual business units or applications. The decision on where and how authorization should be done depends on the completed matrix. At one extreme, if user and information resource attributes and authorization policies are all managed centrally, the obvious choice is to have a central infrastructure. On the other extreme, if user and information resource attributes and authorization policies are confined to local units, the cost justification for a central infrastructure does not exist. All other scenarios fall in between, with central physical or virtual attribute repositories and a central or virtual policy engine.

A central authorization model works best when job roles are well-defined. The Roles-Based Access Control (RBAC) Model assigns

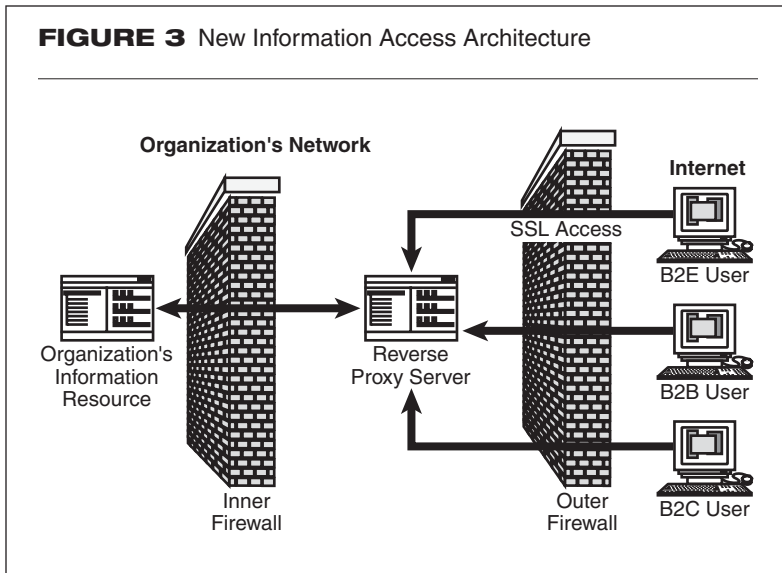
roles to users, thus reducing maintenance costs for managing authorization.<sup>8</sup> However, at the organizational level, role definitions are neither easy to configure to begin with, nor are they easy to maintain. In many cases, roles are defined but only for specific business applications. ERP systems now stipulate roles that can be assigned to users. These roles, however, may not necessarily map directly to broader organizational-level jobs and roles.

One successful approach to manage the balance between central and local authorization is to perform broad-level authorization centrally, while granular authorization is done within an application. This approach, however, is only useful if the broad authorization check first done centrally is not repeated at the application level. Applications provide incremental authorization evaluation beyond the broad authorization granted centrally.

### Security across the Application and Network Stack

Only after user-level identity and information access management are fully designed at the top level can attention be paid to security across the application and network stack. Normally, user requests for access to an information resource are always interpreted and processed through one or more applications. For example, as previously discussed, front-end portals aggregate information from various back-end applications and data sources, and process it before presenting final results to the user. Although the user identity information is passed through one application to another in an SSO environment, there must be a level of trust among program objects that communicate with each other when processing a request. Program objects must have identity and access management just like end users.

Identity management and authentication of program objects is currently rarely implemented when two program objects are communicating with one another within an intranet. That is, network-level protection is entrusted with these tasks, as is the case with determining user types. Although the same mechanisms implemented for user identity and access management could also be utilized for program objects, there are differences. Typically, the identities of program objects are currently supported by code signing and other certificate-based approaches. Many organizations have had limited success in using certificate-based end-user identities. However, they are suitable

**FIGURE 3** New Information Access Architecture

for programs and other resource objects where certificate administration is less costly.<sup>9</sup> Even so, identity management and authentication continues to be a challenge for programs and other non-user entities.

The number of identities and access rights to be managed could potentially explode if any user and any program is directly communicating with another. In this new architecture, all user access to the organizational information resources is provided only through the reverse proxy level, irrespective of the type of user or where he or she is located on the network. The new model now relies on access based on user identity rather than network-level attributes. This eliminates the user dependence on physical or network location. There is no need for VPN or intranet access for B2E or B2B users. Access can be directly associated with such real-world user attributes as user identity, group membership, and organization affiliation.

Restricting resource access only through reverse proxy substantially reduces the number of identities and access relationships that must be managed. Reverse proxy acts as proxy for all of an organization's information resource entities, including program objects. Users can then be assured that a resource is authentic by validating the identity of the reverse proxy. On the other side, only the proxy server communicates with the downstream organization's applications.

This scenario is depicted in Figure 3. Downstream applications and servers could be hardened for security by closing all other network ports and access to these applications and servers. In addition to managing communications

with the proxy servers, access must be managed when application servers such as portal servers communicate with other application servers as well as back-end database servers. Today, we must manually analyze and map these access relationships. The only workable approach is to reduce the number of access relationships to a manageable level, as tighter security is built around more critical resources such as corporate databases.

What happens when a program object in another organization, instead of a user, is seeking access to one organization's resources? Technically, even user access is through a browser. An industry-standard Web services protocol<sup>10</sup> is now supported in many off-the-shelf enterprise resource planning (ERP) and supply-chain management (SCM) products. IBM and Microsoft have also proposed a Web services security model.<sup>11</sup> In such cases, if a partner has a front-end proxy interface that directly communicates with an organization's reverse proxy, it would be easy just to authenticate the partner proxy for all access requests that are channeled through the partner's program objects. That is, organization proxy servers act as authentic sources for all organization information flows to and from the organization IT assets. This scenario is depicted in Figure 4.


Finally, this new model is not immune to the virus and other intrusion attacks. Sophisticated intrusion detection and management systems<sup>12</sup> with advanced firewalls are still needed, as well as routers that can analyze the network packets deeper in the application levels to detect and block intrusion.

## IMPLEMENTING THE ARCHITECTURE

Any architecture, however promising it may appear on paper, will remain just on paper if there is no good project implementation strategy that also addresses organizational constraints.

Before putting together a project plan, business and IT organization structures must be reviewed. A decentralized IT organization makes it more difficult to implement any organizational security architecture. One way to address this issue is to build and offer security infrastructure service to other IT units within an organization.

It is very important to formulate identity management policies early during a project. It is difficult to get agreement on policies before gaining some experience with implementation.

**FIGURE 4** Program-to-Program Access Architecture

However, getting this consensus early during the project anchors the implementation plan, achieves wider communication, and quickly gains legitimacy for the project. The trick is to design policies in a way so that they remain flexible during implementation.

Because of regulatory responsibilities, the growing sophistication and frequency of cyberattacks, and an increasing reliance on IT for conducting business at all levels, security is a high-priority project issue.<sup>13</sup> However, justifying funding for security remains a challenge because security cannot be justified purely on the basis of return on investment (ROI). Security straddles infrastructure and application aspects of a project. If funding approval authority lies with individual business sponsors, it is more difficult to articulate an IT security architecture from an infrastructure investment point of view, where emphasis is more on such business-related issues as streamlining access provisioning for internal and external users. On the other hand, if security is part of the infrastructure budget, security funding may be sacrificed to cover such infrastructure costs as multiple user directory administration. The source of funding influences top-down implementation more than a bottom-up approach.

The IT architecture discussed in this article is very broad and has significant consequences on overall IT resources and configuration. It is unrealistic to accomplish this architecture in one step or even in one project. Implementation of the architecture should be segmented into multiple manageable steps or projects, and management should review each step after its completion and make adjustments as needed to mitigate risks. As mentioned, SSO authenti-

cation and account administration are relatively easy to handle and should be implemented first to gain knowledge about this architecture. Such difficult issues as authorization management and decentralized access administration should be tackled in later phases. Communication and change management during the implementation project are also vital to project success.

#### THE FUTURE OF ACCESS MANAGEMENT

Technology and products continue to evolve in the relatively new cyber-security and access management space. Future products will probably evolve to assist in the administration of organizational access agreements as well as B2B user affiliations among partner organizations. There is no priority to register B2B users in an organization directory; instead, federated identity and other new security information exchange open protocols would be used.

Because of this growing interdependence between identity and access management, there will probably be products that encompass both areas of functionality. Also, these products will likely manage identity and access of not just end users but also of such other non-user entities as programs or other computing equipment, including servers. Future products will take a holistic management approach.

If upcoming products use open standards available to exchange and manage security, organizations will not be tied to a vendor. Security Assertion Markup Language (SAML)<sup>14</sup> is the standard for sharing authentication and SSO across organizations, and eXtensible Access Control Markup Language (XACML)<sup>15</sup> is the standard for sharing authorization credentials.

***The more widely used the user and information resource attributes and authorization policies are, the better it is to provide centrally managed access management.***

Service Provisioning Markup Language (SPML) is the standard for communicating access provisioning information.<sup>16</sup> With standard interfaces, identities and their attributes as well as security policies could be shared among products from different vendors.<sup>17</sup> As identity and access management technologies rapidly mature and become available as products, costs for access management shift downward.

Different modes of sourcing services are now possible across the globe, as has been the case for manufacturing for years. In today's world of "co-opetition" among organizations,<sup>18</sup> where an organization competes and cooperates simultaneously with business partners, organizations must become more nimble and flexible.<sup>19</sup> Virtual organizational structures would require integration to happen in real-time across different organizations to provide just-in-time access. Therefore, demands on information access management will shift upward. IT architectures for access management will continue to evolve in response to these trends.

## CONCLUSION

This article described a new approach in designing an organization's IT architecture for information resource access management. This approach starts with a focus on the business users, the organization's critical information resources and then on the technology.

This step-by-step approach begins with the segregation of users accessing an organization's information resources. Next, user identities need to be managed and authenticated for information access; then all attributes of the organization's information resource base are analyzed from management and usage perspectives. Once user and information resource management is understood, authorization policies are designed on the basis of user and information resource attributes.

The level of investment needed in managing access is an economic decision. The goal of an organization is to provide the right information to the right person at the right time and in the right place. Any implementation is an approximation toward this goal. For example, the right user is determined by the user's membership to a group or role that could be narrowed to one person or could cover an entire organization. As an organization tries to get closer to this goal, costs start to increase and at some point begin to outweigh the benefits.

The combination of how user, information resource, and authorization policies are managed dictates an access architecture. In particular, it tells how much of an access architecture should be provided centrally, outside the other business logic in various software programs. The more widely used the user and information resource attributes and authorization policies are, the better it is to provide centrally managed access management.

After access management is designed at the business entities level, the computing and communication technologies lower in the network stack, which support access to information, are examined. Just like user and information resources, access to and by these technological entities must be controlled. An approach that segments user groups from the most critical information resources reduces the possible number of access combinations to a manageable level. This approach therefore provides security in the most optimal way.

Issues in implementing the architecture range from ensuring a proper match with organizational structure to implementing specific access management components first to reduce risk. This architecture design approach and implementation is a solid instructional template that IT managers can readily use in all aspects of the organization's access management initiatives, from a strategic to an operational perspective. ▲

## Notes

1. Charles Handy, "Thinking About...: Trust and the Virtual Organization," *Harvard Business Review*, 73(3), 1995.
2. The following article describes how knowledge acquisition and business processes span across organization's boundaries: R., Kocharekar, K-Commerce: Knowledge-Based Commerce Architecture with Convergence of E-Commerce and Knowledge Management, *Information Systems Management*, 10(2), 29-34, 2001.
3. Xianjun Geng, Yun Hang, and Andrew B. Whinston, "Defending Wireless Infrastructure against the Challenge of DDoS Attacks," *Mobile Networks and Applications*, 7(2), 213, June 2002.
4. The following article describes the use of the VPN and firewalls in secure information access over the Internet: D.J. Gooch, S.D. Hubbard, M. W. Moore, and J. Hill, "Firewalls — Evolve or Die," *BT Technology Journal*, 19(3), 89, July 2001.
5. M.A. Roth, D.C. Wolfson, J.C. Kleewein, and C.J. Nelin, "Information Integration: A New

- Generation of Information Technology," *IBM Systems Journal*, 41(4), 563, 2002.
6. A comprehensive, but not so holistic (architected) approach is also discussed in the following article: Huong Ngo Higgins, "Corporate System Security: Towards an Integrated Management Approach," *Information Management & Computer Security*, 7(5), 217, 1999.
  7. Privacy and security are the number-one reasons why customers are not purchasing on the Internet, according to a recent survey: Godwin J. Udo, "Privacy and Security Concerns as Major Barrier for E-Commerce: A Survey Study," *Information Management & Computer Security*, 9(4), 165, 2001.
  8. R. Sandhu, D.F. Ferraiolo, and D.R. Kuhn, "The NIST Model for Role Based Access Control: Towards a Unified Standard," *Proceedings, 5th ACM Workshop on Role Based Access Control*, July 26-27, 2000, first public draft of proposal for an RBAC standard.
  9. The following article provides a good overview of the public key infrastructure and its overall acceptance: K.P. Bosworth and N. Tedeschi, "Public Key Infrastructure — The Next Generation," *BT Technology Journal*, 19(3), 44, July 2001.
  10. Web Services Activity Home Page from W3C <http://www.w3.org/2002/ws/>.
  11. M. Hondo, N. Nagaratnam, and A. Nadalin, "Securing Web Services," *IBM Systems Journal*, 41(2), 228, 2002.
  12. Steven M. Furnell and Paul S. Dowland, "A Conceptual Architecture for Real-Time Intrusion monitoring," *Information Management & Computer Security*, 8(2), 65, 2000.
  13. Taz Daughtrey, "Emerging Issues in Information Security Management," *Quality Congress, ASQ's Annual Quality Congress Proceedings*, Milwaukee, 2001, pp. 406.
  14. Home page of OASIS Security Services Technical Committee: [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=security](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security).
  15. Home page for OASIS XACML Technical Committee: [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=xacml](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml).
  16. Home page for OASIS Provisioning Services Technical Committee: [http://www.oasis-open.org/committees/tc\\_home.php?wg\\_abbrev=provision](http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=provision).
  17. Products based on Liberty Alliance's specifications for an interoperable federated identity management framework are already available on the market. Home page for Liberty Alliance: <http://www.projectliberty.org/>.
  18. Adam Bradenburger and Barry Nalebuff, "The Right Game: Use Game Theory to Shape Strategy," *Harvard Business Review*, 73(4), 1995.
  19. The following article discusses the nature of information goods from economic utility perspectives and its implications for organizations: Raju Kocharekar, "Without the Speed Limit, but Within the Limit: Managing Knowledge in Organizations," *Information Strategy*, 17(3), 10-15, 2001.