ISSUE PAPER

“We Have Met the Enemy, and He is Us.”

John Burke, Principal Research Analyst

Executive Summary

Organizations are being forced to move towards identity-based access lifecycle management tools in order to support IT agility and flexibility, safely empower staff, meet compliance requirements, and cope with the rapid evolution of data centers and core systems. Robust identity management demands breadth in the entity being identified, granularity in managing the attributes associated with an identity, and easy integration into access control and auditing tools. An important corollary to identity is context: modifying user privilege based on where, when, and how an entity attempts to use it. Such capabilities are especially important and useful in ferreting out or preventing the abuse of access privileges, one of the main modes of insider system or data compromise. Identity-driven access management, by making it possible to provide to each entity all and only those services it should receive, enables rather than impedes change, growth, and innovation. Service delivery that can be trusted even in the midst of great change makes it possible to take full advantage of service-oriented architectures, virtualization, and increasingly fluid organi-zational structures, without taking on new and unacceptable operational risk.

The Issue

Organizations of all sizes and across all industries are caught in the turbulent intersection of seemingly contradictory trends: the movement to allow freer access to data and the movement towards ever tighter control of access to data. At the same time, trends in the data center are reinforcing the former and hampering the latter, whether or not it is what IT or management wants.

The drivers pushing organizations to increase access are many: agility, flexibility, and empowerment. Agility, the ability of the organization to respond to changes in its environment, whether for defense or for gain, depends increasingly on broader access to data. In order to quickly change the way it operates in the market, perhaps by implementing new customer service processes centering on creating a single point of contact for a customer with a problem, the organization very often has to make previously inaccessible information available to the people in that process. In order to address a new opportunity, the organization must often pull together streams of data that had not previously intersected, and again make them newly available to staff or customer or prospect.

Similarly, IT flexibility, the ability to do more and new things with only minimal additional resources, requires that IT be easily able to repurpose tools and processes in the face of new challenges. To avoid having to add entire new application stacks and associated data stores, IT needs to be able to provide new or modified access to existing data stores and systems, quickly, easily, and broadly. To use the customer relationship management system to track activities among internal departments, for example, it may be necessary to link in data about the internal structure of the organization, and about staff with no direct relationship with customers.

IT-based empowerment aims to make it easier for staff in the organization to be flexible and agile as well. Empowerment technologies allow staff to create ad-hoc and ephemeral structures – temporary teams, working groups, shared work spaces. Such ad-hoc support tools let them explore new ways to accomplish their regular duties (and to find out if those new ways are better and deserving of preservation). It may also help them address a new challenge or explore an opportunity with little or none of the formal machinery of project and process management.

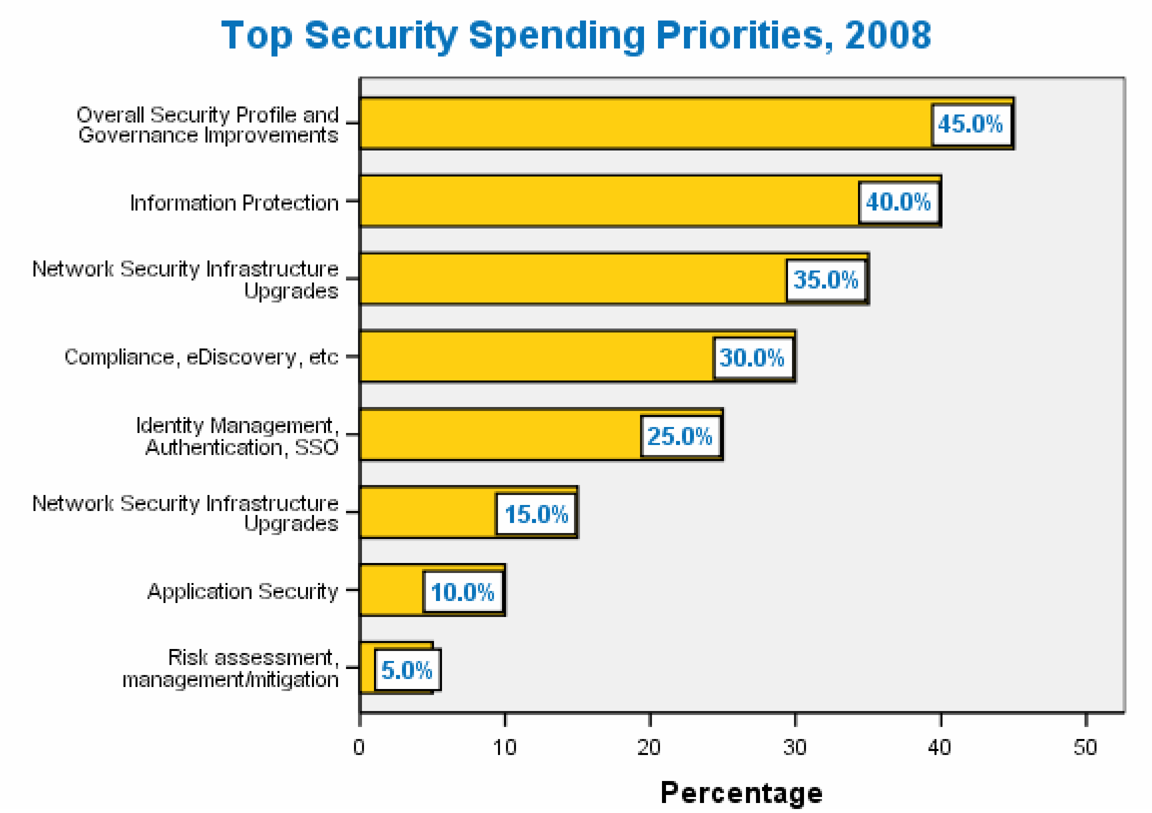

Figure 1: Security Spending Priorities for 2008

Whether it is to be agile, or flexible, or to unleash the untapped creative problem solving skills of staff, organizations feel an urgent need to allow broader and freer access to data. It can be a matter of survival, if the stakes are meeting a market threat or finding and exploiting a new opportunity in a stagnant market; or a matter of renewal, reinvention, and growth.

Drivers to limit and control access are many, and equally urgent: risk management, governance, and compliance. Risk management is perhaps the greatest motivation for controlling access to systems and data. Whether to protect the integrity of systems and the data they manage, or to preserve the availability of those systems and data, or to guarantee the confidentiality of that data, managing operational risk to the organization, in part, entails controlling access to— and the flow of—information.

Good governance is built on data flowing according to the "need to know," with that need defined as granularly as possible, allowing for "least possible privilege" access rights to be granted. Specifying those access privileges based on—and in support of—separation of duties is one of the strongest protections available against insider threats to confidentiality, integrity, or availability. Only 48% of participants in Nemertes' *Security and Information Protection* benchmark had policies for separation of duties, however, and even fewer were able to consistently enforce them, due to staff size and turn over. Happily, though, 45% of their security spending priorities for 2008 centered on such improvements to practice. (Please see Figure 1: Security Spending Priorities for 2008, page 2.)

Compliance with the law, with government regulations, and with professional-association and trade-group standards echoes and reinforces that drive to control access by creating external motivations to caution, in the form of fines, criminal charges, or public exposure. In extreme circumstances, the consequences of allowing improper access to data could threaten the viability of the organization itself: loss of sanction by the government (in a regulated business), loss of reputation so significant that continuing operation was impossible, or a precipitous drop in stock prices and mass customer exodus.

So, many a move to free up access to data is, potentially, fraught with peril. Unfortunately, technological trends are not just supporting but driving many such changes, as pursuit of a new architecture often drags along with it a changed relationship to organizational data. New technologies, new architectures, new business strategies, and business events all conspire to complicate the problem of managing access to information. Service-oriented architectures (SOA), for example, break down information as well as programmatic silos; virtualization decouples function from location in new ways, breaking older models for controlling access to data based on physical connectivity; software as a service (SaaS) in many cases requires that confidential or sensitive information flow across the Internet and reside on the systems of other corporate entities. More than 63% of participants in Nemertes' *Service-Oriented Architectures and Applications* benchmark already use SaaS to deliver some complete enterprise application to their users.

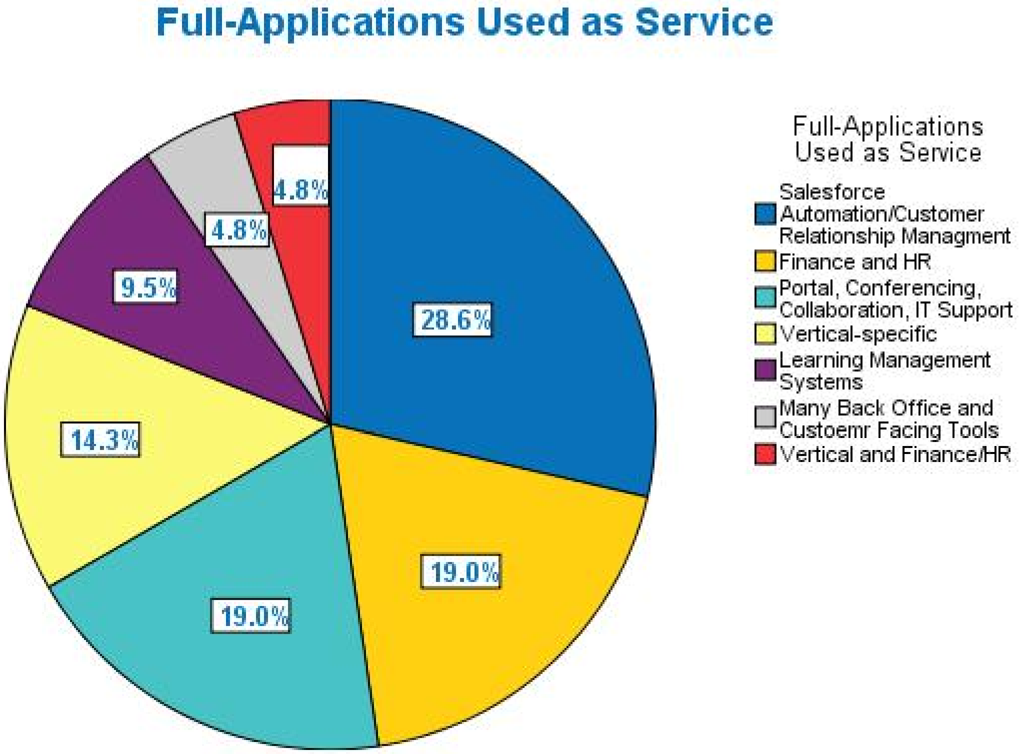

Figure 2: Breakdown of Applications Delivered via SaaS

(Please see Figure 2: Breakdown of Applications Delivered via SaaS, page 4, for a breakdown of which complete enterprise applications participants used SaaS to deliver.)

On the organizational and business side, hiring freezes, staff reductions, and staff churn all drive IT and business lines to relax strict controls on access based on role, as the roles flow among staff temporarily during transitions, or to new staff repeatedly in times of high turnover, or are permanently combined as staff size is reduced. Mergers and acquisitions highlight the intersection of technological and organizational stressors, requiring the combination of systems and staffs, often resulting in data bases merging and staff being reduced or turned over.

In order to navigate among these challenges successfully, organizations need to shift to a consistent basis for granting, enforcing, tracking, and revoking access privileges. To ensure that users have all – but only – the access privileges they require in order to perform their current duties, and that all systems have the ability to manipulate the data their end-users need, security and control based on physical and functional silos have to give way to systems based on a robust and flexible notion of identity.

Identity is the new foundation

When systems, roles, and the outline of an organization itself can change rapidly and frequently—when agility requires that previously separate systems interact, when organizational change requires that staff or even clients and customers take on new roles, when function and data both can reside inside or outside the organization in a way that varies over time—all that remains in the end as a stable foundation for access management is identity. Whether of people, groups of people, roles, corporate entities, devices, or processes, identity is the key to robust and future-ready security, and the ability to reliably deliver all and only the correct services to the entities entitled to use them.

Robust identity management demands breadth in the entity being identified, granularity in managing the attributes associated with an identity, and easy integration into access control and auditing tools.

A flexible notion of identity is critical in forward-looking identity management systems. The dissociation of software from specific hardware and specific network and storage paths that is the hallmark of fully virtualized infrastructure, and the mobility of desktop and other systems, undermine any place-based security. Since IT is rapidly losing the ability to identify systems and applications by their location, or simply control their access to data by physical means, it must extend to systems and applications the same notion of identity and identity-based access that it applies to users. This becomes especially important in SOA environments, where a service may be instantiated in one, many, or varying places within the infrastructure, but at all times requires the same access to the same data in order to perform its functions.

In order to serve in an ecosystem defined by SOA and incorporating SaaS, in a business environment requiring close cooperation and even deep interoperation with the systems of partners, customers, and suppliers, identity systems must be able to federate. Federation allows different organizations trust boundaries and management of access to overlap. SOA especially also requires that identity and access control based on identity be able to propagate throughout an infrastructure, so that access controls can be applied anywhere, as distributed as the services themselves.

Granularity is important to allow the fine tuning of access privileges to data and resources. A server, real or virtual, must have certain access privileges simply in order to be a part of the infrastructure (chiefly, the right to connectivity for network and storage); it can have many more privileges based on what applications or application components it is running. An application or component may have a general right to access information related to, for instance, client health records, but that right may be modified dynamically based on who is making a request for information. This illustrates both the need for fine-grain rights assignment and the requirement that access privileges be reckoned cumulatively, with the most restrictive set dominating.

Identity must also be managed according to the lifecycles of the entities in question: roles with associated privileges assigned to people, applications, or systems, as soon as they are needed and for as long as needed—but no longer.

You can't get *there* from *here*....

With a robust basis of identity, organizations can build a common basis for access policy around enterprise assets. An important corollary to identity is context: allowing user privilege to be altered based where and when and how an entity attempts to act on it. That is, the granularity of the identity and access control systems need to be able not just to say, “You can get there, he cannot,” but rather “You can/can't get there from here, at this time.” For example, an admitting nurse in a clinic may need to see patient health records when working in an intake examination room and from the intake application, but organizational policy, driven by HIPAA and the concern for patient confidentiality, may need to deny that same access when it comes from a computer in a public space like a cafeteria; or from a different application; or during hours when the nurse is not supposed to be working, no matter what the location; and should always deny access to financial data in the record.

An important addition to the idea of context sensitivity is the notion of behavioral context. Ideally, systems will also support an ability to say, “You can’t get there now given what else you have been doing in the last few minutes.” Consider as an example drawn from real life the case of a biochemist in a pharmaceutical research lab: that person, as a matter of course, will be reading and even downloading documents related to recent research and containing important intellectual property. Such activity, taking place in an office or lab, a few times a week, presents no problem: it is the expected, normal behavior associated with the role. If that same person undertakes 400 downloads in a row, from a lobby or lounge, the behavior is anomalous enough to be suspect, and should lead to immediate and broad revocation of privileges.

...and Where Have You Been All This Time?

Along with basic access privileges, context, and especially behavioral analysis, comes compliance and auditing. Basic audit trails will record who makes use of which access privileges. This sensitivity to who can do what is only a part of the picture, though. Since policy concerning access is so often driven by context—the need to know derives from who, where, and when an entity is—that robust awareness of all these factors is essential to full auditing. Certainly it is possible for human eyes to find the patterns of usage implicit in basic access records, but access control systems that track those patterns themselves, especially if they are also capable of acting on them by raising alarms or locking down accounts, can enable more complete, faster, and more scalable identity-based access management.

Such tools and capabilities are especially important and useful in identifying, tracking, or preventing the internal abuse of access privileges that constitutes one of the main modes of insider system compromise. The prevalence of insider breaches – more than 26% of participants in Nemertes’ benchmark, S*ecurity and Information Protection*, reported internal security breaches in the last 3 years, and the FBI continues to identify them as the number one form of security breach across the board – makes this an especially attractive set of features to deploy. Although more than half of the participants in this benchmark say they are deploying tools specifically aimed at reducing insider threats, they are rarely built on such a full featured integration of identity and access management and behavioral analysis.

Conclusion

For reasons of agility, flexibility, and organizational empowerment, swept along by the rapid evolution of data centers and core systems, and in response to the requirements of compliance and risk management, organizations are being forced to move towards identity-based access lifecycle management tools. These systems, far from being impediments to the easy realignment of systems, groups, and functions, become the necessary and empowering means by which such changes can be not just survived or sustained, but instead embraced and exploited without a crippling introduction of risk and uncertainty regarding who has access to what. By making it possible to take advantage fully of service-oriented architectures, virtualized infrastructures, and increasingly fluid organizational boundaries (internal and external), identity-driven access management enables rather than impedes change, growth, and innovation.

About Nemertes Research: Founded in 2002, Nemertes Research specializes in analyzing the business value of emerging technologies for IT executives, vendors, and venture capitalists. Recent and upcoming research includes Web services, security, IP telephony, collaboration technologies, and bandwidth optimization. For more information about the analyst, please contact Christine Zimmerman at research@nemertes.com.