

# Identity Management & Trust Services

Foundations for Cloud Computing

By Jack Suess and Kevin Morooney

In the recent EDUCAUSE book *The Tower and the Cloud*, Richard Katz organized a set of essays by higher education community leaders who explored the emergence and impact of “cloud computing.” The book, which uses a broad definition of cloud computing, compellingly highlights how the Internet-based technologies of cloud computing are changing higher education. Katz, in his introductory chapter, stresses that institutions need to develop a cloud strategy, and he argues that institutional leaders, especially CIOs, must build into their strategic planning the role of virtualization, software-as-a-service (SaaS), open resources, shared community infrastructure, and commercial cloud computing offerings.<sup>1</sup> Increasingly, IT organizations will move from providing IT services locally to becoming an integrator of IT services—some provided locally and others provided outside the institution. As a result, institutions must immediately begin to plan for shared services and must understand the essential role that identity management and trust services play in making integration possible.

*Jack Suess is Vice President of Information Technology and CIO at the University of Maryland, Baltimore County. Kevin Morooney is Vice Provost for Information Technology and CIO at The Pennsylvania State University. Morooney is Vice-Chair and Suess is Treasurer of the InCommon Steering Advisory Group, for which they co-chaired a subcommittee that recently looked at options for advancing federated trust services in the higher education community.*



One of the most tangible areas where this shift from local provider to consumer is playing out is in the provision of e-mail services. At most institutions, e-mail is still considered a mission-critical application and is fundamental to administration, teaching, and research. Failures or outages in e-mail systems severely disrupt the day-to-day operations of the campus. Not surprisingly, since the announcement of Google Apps for Education in 2007, the subject of e-mail outsourcing has been a constant discussion topic on the EDUCAUSE CIO constituent group list. Likewise, in a July 2008 EDUCAUSE Center for Applied Research (ECAR) survey, nearly 20 percent of the 351 responding colleges and universities had outsourced their student e-mail to a commercial provider.<sup>2</sup> What may be surprising about this development is how seemingly well the transitions have gone. Obviously, there have been issues, but whenever questions about the process have been posed to the CIO constituent group list, the overwhelming response has been that the central IT organization is delivering service as good as or better than when e-mail was hosted locally.

E-mail has been the service most discussed by the CIO community in moving from locally operated to cloud-based applications, but it represents nothing more than the visible part of the iceberg. Over the last few years, hosted applications, sometimes referred to as software-as-a-service (SaaS), have taken off. Institutions can now procure externally hosted services that support residential housing management, event management, faculty productivity reporting, fa-

cility management, institutional assessment, student billing services, parking services, learning management, emergency notification, and alumni services—to name just a few. As the web becomes the common user interface, consortiums and state university systems are deploying administrative solutions that support a shared infrastructure for multiple institutions to operate their human resource management, payroll, financial, and student information services.

Although the list of potential externally managed services is long, and growing each day, most institutions are utilizing only a few. In many instances, the decision to procure these services was driven not through any coherent IT strategy but by campus departments making local decisions and procuring these services with varying degrees of central IT involvement. In some cases, the decision is made by individuals, such as faculty, who may procure a hosted service, such as WebAssign, as part of the requirement for their course. The benefits to institutions in utilizing the hosted applications are often focused on deployment and cost:

- *Faster deployment cycles:* days or weeks—versus months or years to develop or implement an application locally
- *Lower initial cost:* paid for on an annual basis, with no large upfront cost for hardware and implementation

In the last five years, institutions have begun to effectively utilize external services for mission-critical needs through an evolving set of standards and technologies.

- *Virtualization.* The advent of high-quality virtual operating systems has allowed external service providers to achieve economies of scale. They can use virtualization software to run multiple organizations on the same physical infrastructure as if each were independent. This allows the external service provider to leverage economies of scale and reduce the cost and management overhead.
- *eXtensible Markup Language (XML).* XML provides a standard way to share information and data. By utilizing XML, an external service provider can implement one standard method for the interchange of data between the institution and the service provider.
- *Web 2.0 technologies.* Web technologies now allow web-based applications to have sophisticated user interfaces. External service providers adopting these web standards can provide access to any person connecting to the service with a standards-compliant browser.
- *Network bandwidth cost reduction.* Over the last decade, the higher education community, led by regional and state networks, has leveraged its collective buying power to greatly reduce the cost of commodity Internet bandwidth.<sup>3</sup> As a result, most institutions can procure adequate bandwidth and do not see commodity bandwidth as an issue when utilizing external services.

These four components make it possible for external service providers to offer cost-effective and compelling hosted services. But it is important to recognize



**In the last five years, institutions have begun to effectively utilize external services for mission-critical needs through an evolving set of standards and technologies.**

that these technologies have enabled service *providers*—not service *consumers*. The challenge for IT organizations is to integrate these disparate services in a coherent and effective manner. Issues such as authentication, access control, and the user experience in moving from one hosted service to another are all important factors for long-term success.

What has been missing is a way for institutions to quickly and effectively

integrate these external service offerings. At present, many external services require a separate username and password combination that is stored by the external service provider. Managing the creation and deletion of users, ensuring application access rights, and integrating the external services into the broader campus computing environment are all left to the institution. Worse, external service providers tend to differ in their approaches

to these actions, requiring institutions to develop and maintain multiple methods as they add new external services to their offering.

To support both local and external service-delivery models, institutions need a comprehensive approach to identity management and trust services—an approach that allows external service providers to leverage campus identity management and trust services. This comprehensive approach should focus on three activities:

1. *Developing an identity management system.*

The numerous articles written about the development of identity management systems can serve as useful starting points.<sup>4</sup> In addition, a number of commercial products—such as Microsoft's Identity Lifecycle Manager (<http://www.microsoft.com/ilm>) and Sun Identity Management (<http://www.sun.com/software/identity/>)—now make this task less daunting than in the past.

2. *Creating a standard set of attributes for each person.* Personal attributes have been defined by the eduPerson schema (<http://middleware.internet2.edu/eduperson/>), developed by the Middleware Architecture Committee for Education (MACE) and the higher education academic community in consultation with outside groups such as the American Association for Collegiate Registrars and Admissions Officers (AACRAO). The eduPerson schema lists some common elements, such as campus role or username, that can be requested by outside applications.

3. *Enabling external access through a federation such as InCommon.* Working to link identity providers (such as higher education institutions) with service providers (such as other higher education institutions, commercial entities, and government/nongovernment agencies), InCommon (<http://www.incommonfederation.org/>) presently uses two community-developed products: the XML-based Security Access Markup Language (SAML)<sup>5</sup> and Shibboleth (<http://shibboleth.internet2.edu/>), a web-based service

that supports authentication for remote service requests for trust services.

Through these three activities, institutions can work with local and external service providers to create a common standards-based approach to authentication. In addition, because the final authentication is performed through a local campus service interconnected to Shibboleth, institutions have much better security control. The following section offers a more detailed explanation of how these three activities fit together.

### A Primer on Identity Management Systems

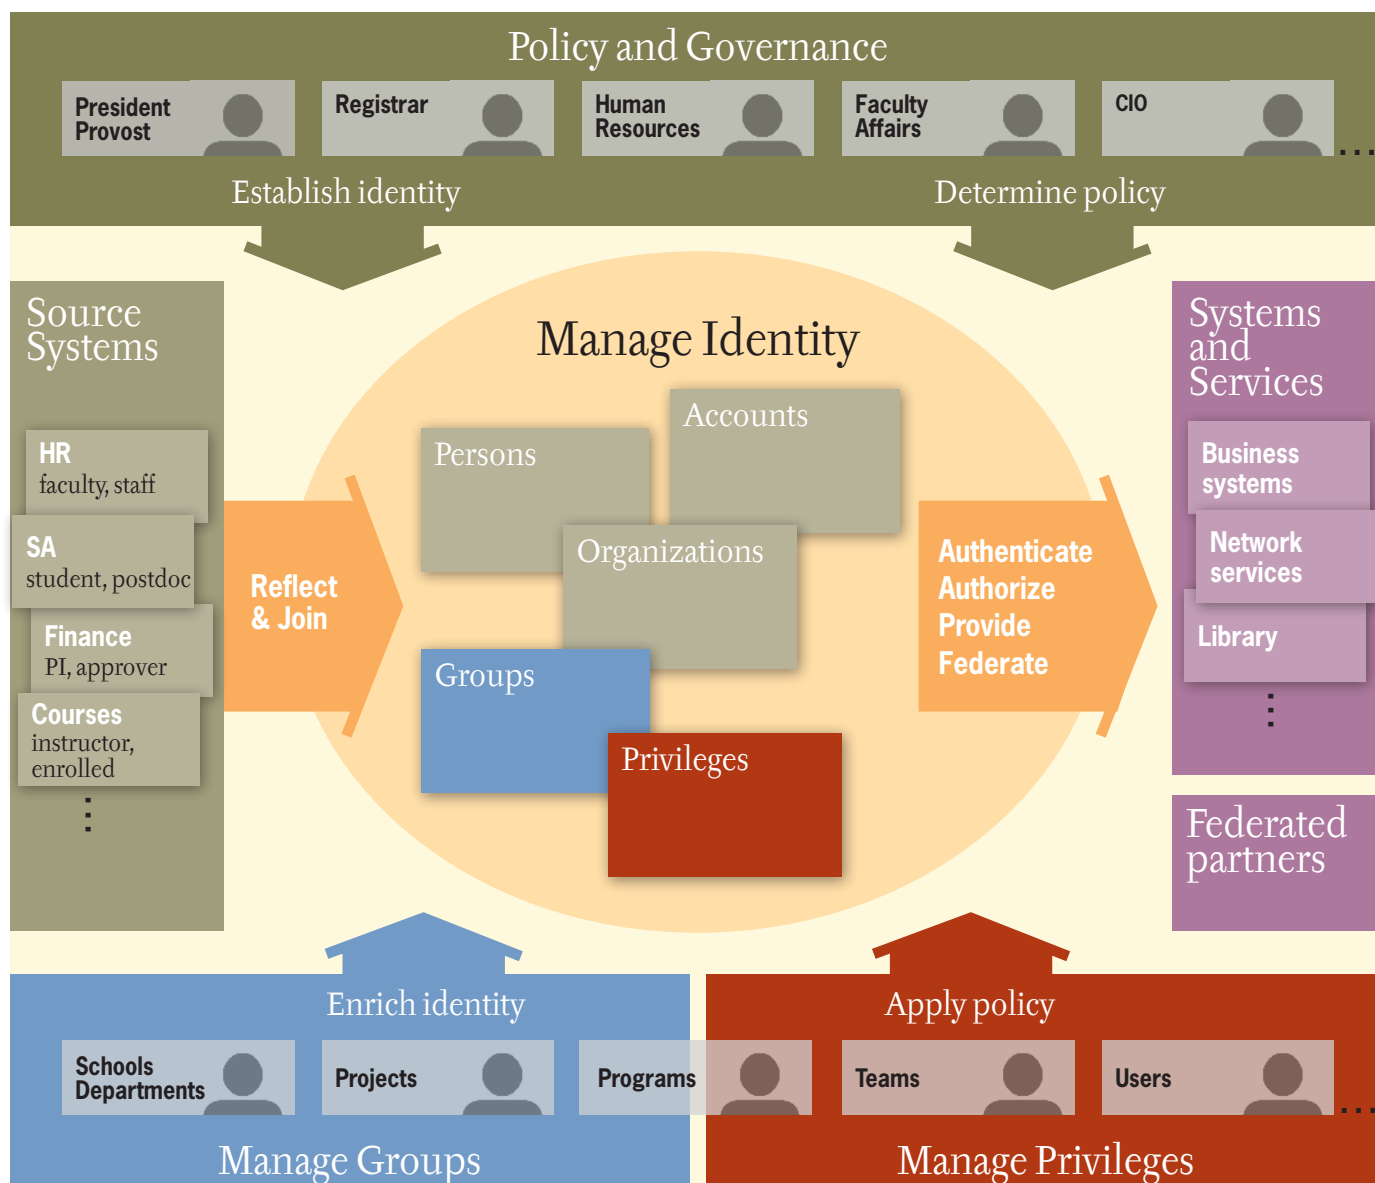
Figure 1 provides a model for a local identity management system. For this example, the yellow cloud in the middle should be considered the identity management system.

At the top, the dark-green area focuses on policy and governance. The first step is to develop a process that verifies and establishes the identity of a person who has been given college/university credentials. At most institutions, this requires developing a process to review an official government-issued


identity card (such as a driver's license) in order to validate that the account credentials are being assigned to the appropriate person. Because colleges and universities are complex communities with many different groups, it is important for the campus leadership team to develop policy on how members of the community will be added to the identity management system.

On the lefthand side, the light-green area identifies some of the many record-keeping systems that may be sources of information for the identity management system. These source systems look for

FIGURE 1. A Model Framework for an Identity Management System



Source: Developed by Ann West, of Internet2, and Lynn McRae, of Stanford, and used with permission from Internet2.




## One key benefit of an identity management system is the ability to manage the complexity of higher education.

specific kinds of transaction events and trigger when an update needs to occur in the identity management system. For example, someone was hired or terminated in the human resource system, or a student was added to a specific course. One key benefit of an identity management system is the ability to manage the complexity of higher education. As a case in point, if a campus hosts a number of summer conferences, the people attending these conferences need to have

an account to utilize the computer facilities or possibly to authenticate for access to the Internet. Therefore, a mechanism is necessary so that the conference organizers can add people to the identity management system as conference guests, with a date for when access will be automatically turned off. This delegation back to the department overseeing the business function provides better customer service and demands less from the IT organization.

On the far right, the top lavender box identifies some of the specific internal systems and services to be managed through the identity management system. From the set of defined systems and services provided, data elements can be identified that must be present in the identity management system. One key decision to make when developing the identity management system is how quickly to propagate the changes that occur in the various systems of records. For some situations, such as turning off access when an employee is terminated, changes may need to be updated in the identity management system very quickly. In many situations, a single daily update from the source system to the identity management system will be sufficient. It is important to work with the business process owners to understand the risks and requirements. Although there is a trend to move to real-time updates of the identity management system from the source systems, it is always possible to develop specific business processes that provide a workaround should that not be the case (for example, a human resources representative may call someone in the IT organization to have an account password reset when an employee is terminated).

The blue box at the bottom, “Manage Groups,” is at the heart of an identity management system. Groups provide the flexibility for managing access and offering collaborative services. Groups can be generated based on an attribute about a person, or groups can be formed ad hoc, when individuals need to collaborate and use shared resources. Identifying the groups that a person belongs to can be essential to defining which resources a person can access. In many cases, the group memberships are used by the campus portal to identify the services a person should be able to access. Identifying the groups to create requires consultation with business process owners and is often driven by the services being providing. When setting up groups, an institution should begin by reading a description of the eduPerson schema. The eduPerson schema has an affiliation attribute and lists some standard affiliations that have been identified to date. The permissible values include the



## With the advent of web services, the model of applications managing all aspects of the application security is beginning to prove unworkable.

following: *faculty, student, staff, alum, member, affiliate, employee, and library-walk-in.*<sup>6</sup>

The red box, “Manage Privileges,” is a key box to consider. Presently, many services (or applications) manage privileges within the specific logic of the application. With the advent of web services, the model of applications managing all aspects of the application security is beginning to prove unworkable. The emerging model of web services, also known as service-oriented architecture (SOA), utilizes functionality from outside

the application to perform business processes and is requiring industry groups to rethink their approach to application-level security. At present, the World Wide Web Consortium (W3C, <http://www.w3.org/>) is developing standards for this. On another level, Web 2.0 collaboration tools such as wikis and blogs require that access management be dynamic and not require the central IT organization to intervene. One interesting pilot project is COmanage (<http://middleware.internet2.edu/co/>), a tool that can use the identity

management system to allow users to set up collaboration spaces. These collaboration spaces work across institutions by leveraging federated identity management (discussed below).

Finally, the bottom lavender box on the right, “Federated Partners,” refers to externally hosted services or inter-institutional services being shared. The idea is to have a standard method, as part of the identity management system, for supporting these services. As a new external service is identified, use of this method would quickly set up a mechanism to provide access to that service. An example of such a method being used today is InCommon.

As noted earlier, InCommon works to link identity providers with service providers. Participants include more than 116 higher education institutions, 41 sponsored partners (commercial service providers), and six federal agencies and nonprofit groups. InCommon coordinates common definitions and guidelines for security and privacy and for data



## These three elements—InCommon, SAML, and Shibboleth—are changing the way institutions manage their relationships with externally hosted solutions.

interchange, and it validates that both parties are who they commit to be and are acting in good faith. This information is then encapsulated in meta-data that is included within certificates allowing the identity provider and the service provider to share information.

InCommon supports the sharing of information in real time using standards such as SAML, developed by Internet2 through the efforts of the higher education community and the Organization for the Advancement of Structured Information Standards (OASIS, <http://www.oasis-open.org/who/>), and through web-based community-source tools such as Shibboleth. One of the important aspects of this technology is that it was designed to support the academic value of privacy by encouraging service providers to request the minimal amount of information necessary to support a transaction. An example of how this works is accessing an electronic database that includes content licensed for faculty and student access. When someone tries to access this content, the system will validate if the person is an authorized faculty member or student and will share just that high-level information with the content provider.

These three elements—InCommon, SAML, and Shibboleth—are changing the way institutions manage their relationships with externally hosted solutions. Many campuses now encourage or require externally hosted service providers to become members of InCommon or to agree to join InCommon as part of their procurement requirements. In addition, these three elements facilitate the process so that external providers can become service providers to additional institutions. For example, in 2006, the Univer-

sity of California Office of the President implemented UCTrust (<http://www.ucop.edu/irc/itlc/uctrust/>). By having the individual institutions of the University of California system join InCommon, the University of California system office could deploy, across the entire system, an application that provided employees with access to retirement information. Also in 2006, the Virtual Library of Virginia (VIVA) licensed content from the Public Broadcasting System (PBS) and used these three elements to support distribution to close to 400,000 students.<sup>7</sup>

### Identity Management and Information Security

Initially, identity management was thought of separately from information security and was focused on directory services. The identity management infrastructure was developed and designed for provisioning services, especially centralized authentication. Using the lightweight directory access protocol (LDAP), institutions of higher education focused on developing a comprehensive directory service of all members in the community, pulling data from a variety of source systems. Using the authentication service built into LDAP, institutions were able to create a single authentication system for the institution. The single authentication system led to the development of web single sign-on tools that facilitated automatic sign-on across distinct web applications.

The relationship between identity management and information security is focused on three distinct areas: authentication, access management, and compliance. *Authentication*, the process of ensuring that users are who they say they are, is central to security. A major benefit of

linking identity management with information security programs is the process of validating identity. This process, often called “identity proofing,” requires that users establish their identity by providing credentials that confirm they are who they say they are. This identity proofing is done as part of the policies and procedures for the identity management system and may be overseen by the information security officer for the institution. Alternatively, the information security officer may work closely with other offices, such as user services, to make certain this is done according to policy.

*Access management* is focused on provisioning services based on group membership or specific attributes that qualify one for access to an application or service. For example, based on their role, faculty members may automatically be provisioned to have access to certain administrative functions, such as advising. Maintaining application security is an area that many institutions struggle with because it requires inter-office communication regarding the ebb and flow of people assuming different roles throughout the institution. Leveraging identity management to automatically provision application security is a major benefit.

*Compliance*, especially in support of legal mandates and auditing requirements, is a critical component of an information security program. A centralized identity management infrastructure that controls authentication and application security provides a single point for security compliance logging and auditing. Often, these audit reports and security logs are reviewed by the information security officer. As institutions seek to follow information security standards, such as ISO 27002, they will find that identity and access management is one of the core components of their information security program.

### Federated Identity Management and InCommon

Federated identity management is the practice of using identifying credentials in one domain or organization to access assets in a different domain or organization. Individual organizations identify employees, partners, customers, and so on, and they build internal processes around

## Strategic and Practical Steps

If your campus does not have an identity management system in place, you should leverage the resources of the EDUCAUSE Identity Management Working Group (<http://www.educause.edu/IDMworkinggroup>) to build a successful business case and a project plan for implementing one. In particular you should focus on the following three tasks:

1. *Establish a data governance process to oversee identity management.* If your institution does not have an IT governance or data governance process, you should convene senior leaders and critical stakeholders to ensure that the strategic needs are well understood and that the necessary resources are available.
2. *Conduct a risk assessment, including an inventory of applications or services.* It is sometimes difficult to make the case for a comprehensive identity management system or to assign resources until you know the extent of the institutional risk. A risk assessment will help the institution to engage business process owners and to inventory existing systems and may lead to the discovery of redundant efforts for authenticating access to services.
3. *Plan for the establishment of an enterprise identity management system.* The risk assessment is likely to surface findings that will require remediation and the allocation of resources. A plan that identifies necessary steps, priorities, resources, and timelines will help you keep the building of the identity management system on track. As part of this plan, campuses should implement the requirements associated with the eduPerson attribute definition.

Once the identity management system is in place or the project has begun, the campus should consider taking the following two steps:

1. *Join InCommon.* InCommon provides a standards-based method for connecting identity providers and service providers. By joining InCommon, members of the higher education community can demonstrate to service providers that this is a preferred method for service providers to adopt.
2. *Require, in procurement RFPs, that service providers support InCommon.* By selecting vendors based on their support for InCommon, colleges and universities will use their collective buying power to build a marketplace for federated trust services.

those identities and the degree of assurance to which they can attest the individual is who he or she purports to be. With federated identity management, various organizations agree on how they will trust other organizations' practices for identifying and assuring individuals. A simplistic example of a federation at work is an interstate highway system. If someone is issued a driver's license in one state in the United States, all other states have agreed to recognize both the person's certification and his or her knowledge of driving laws, enabling the driver to access all state highways and interstate highways. Organizations must have robust, trustworthy identity manage-

ment practices in place before they can develop and design services and partnerships that leverage federation. With such practices in place, an organization is in a position to consider both asserting its identity into other realms and accepting other identities into its own.

A trust federation like InCommon plays at least two important roles in an online trust environment. InCommon acts as a scaling factor in relationships, enabling organizations to manage their trust relationships in a scalable manner. If all federating organizations had to manage their trust relationships on a one-by-one basis, there would be tremendous dupli-

cation of effort. By becoming a member of a trust federation such as InCommon, organizations and institutions leverage the network effect of a community of trust. And this leads to a second role of InCommon: the establishment of standards for identity management practices among participating institutions.

More technically, federated identity management is focused on managing the exchange of attributes, called meta-data, between identity providers and service providers. The meta-data describes what attributes will be shared and what level of assurance must be in place for a connection to be allowed. This process leverages Shibboleth to share messages using SAML. Through a handshaking process, the service provider and the identity provider exchange information to validate that the person requesting this service is eligible.

InCommon functions as the lead in negotiating a common set of requirements between identity and service providers. The basic InCommon requirement for identity providers assumes that institutions have made some effort to validate identity but may not have done a formal review against a government-issued ID. InCommon is working with the National Science Foundation (NSF) and the National Institutes of Health (NIH) to support a higher level of assurance, called InCommon Silver, that will require validation of a government-issued ID. InCommon also organizes workgroups to determine the meta-data that will be shared and works to distribute the appropriate certificates necessary for this information exchange to occur. Without a central service such as InCommon, each institution would be left on its own to negotiate with service providers, and the process would not scale.

### The Challenges

As higher education institutions begin to fundamentally rethink the way they provide services, they face a number of new challenges, mostly around identity and access management.

The first and perhaps most arduous challenge is how to encourage pervasive adoption of best practices in identity management among campuses and institutions. A successful identity infrastructure



## A successful identity infrastructure requires thinking holistically about identities and the interdependencies that exist.

requires stepping outside of the departmental and divisional silos used for applications and thinking holistically about identities and the interdependencies that exist. It requires thinking about service provisioning and the end-user experience in the broadest sense and looking at transactions differently—taking into consideration the various interdependencies among groups providing services and developing processes, infrastructures, and policies that break down silos and sustain the new environment.

Another challenge is building support for the benefits of federation without a national mandate. In the United Kingdom and Switzerland, the national government has supported and mandated the use of federation. In the United States, this is not the case. The United States needs enlightened institutional leaders to embrace the federation, just as they have embraced the Internet. With these enlightened institutional leaders will come the service, transaction, and information providers.

Higher education is well positioned to be a leader in creating this marketplace and in establishing a set of best practices that support the expansion of federation-based trust services. At the same time, if the higher education community fails to collectively embrace and adopt these trust services, the resulting systems will undoubtedly focus on commercial transactions, at the expense of supporting research and academic needs.

### **The Opportunities**

The higher education community has been a pioneer in promoting the development of campus networks and in leveraging the Internet to support the academic mission. Those of us in the higher education community take great pride in our early role in shaping the Internet into what it is today and in continuing to push to advance the state-of-the-art. InCommon represents the community coming together to solve one of the most vexing issues faced today: how do we establish and manage trust services on the

Internet? Increasingly, our work is less about just passing data and more about building value. InCommon represents a way for the community to build shared services that we can leverage.

There is no debate about the explosive innovation that has occurred as a result of building the Internet. If the promise of federated identity can be realized, a similar explosion in innovation will occur. The Internet Society (ISOC) has determined that the issue of trust is both

important and crucial for the long-term growth and success of the Internet.<sup>8</sup> Just as the higher education community's adoption of the Internet in the 1980s helped to demonstrate the potential of the Internet, the community's adoption of federated trust services is needed now to demonstrate and unlock the potential of trust services. With budgets under pressure as they are today, leveraging new ways of providing services is essential to success.

Ubiquitous adoption of standardized trust-based approaches to federating identities will unlock the opportunities to be found in connecting people to people, and people to information. By joining InCommon and adopting federated trust services, colleges and universities will be able to shape this endeavor to meet their needs—gaining vendor support for providing new standards- and cloud-based applications that can be leveraged, allowing researchers to use trust-based services to easily traverse network defenses that now block their research, and providing new ways of sharing content across institutional boundaries for teaching and learning. Through the higher education community's collective support, information and people will be able to discover each other with trust, and new collaborative business models based on federating identities will emerge to meet the community's needs.

#### Notes

1. "The Gathering Cloud: Is This the End of the Middle?" in Richard N. Katz, ed., *The Tower and the Cloud: Higher Education in the Age of Cloud Computing* (Boulder, Colo.: EDUCAUSE, 2008), pp. 22–23, <<http://www.educause.edu/thetowerandthecloud>>.
2. Mark C. Sheehan and Judith A. Pirani, "Spreading the Word: Messaging and Communications in Higher Education," *EDUCAUSE Center for Applied Research (ECAR) Research Study*, vol. 2, no. 9 (2009), Key Findings: <<http://net.educause.edu/ir/library/pdf/EKF/EKF0902.pdf>>.
3. Data from the Quilt CIS (<http://www.thequilt.net/proj-cis/>) annual survey show the cost per megabit has dropped by over 400 percent from 2003 to 2008.
4. For example, see Brian L. Hawkins, "What Higher Ed Leaders Need to Know about IdM," *EDUCAUSE Review*, vol. 42, no. 5 (September/October 2007), pp. 84–85, <<http://www.educause.edu/library/ERMO7510>>, and see Norma B. Holland, Ann West, and Steve Worona, "A Report on the [EDUCAUSE] Identity Management Summit, November 2–3, 2006," <<http://www.educause.edu/Resources/AReportontheIdentityManagement/154436>>.
5. Prateek Mishra, ed., "Differences between OASIS Security Assertion Markup Language (SAML) V1.1 and V1.0," OASIS draft, May 21, 2003, <<http://www.oasis-open.org/committees/download.php/3412/sstc-saml-diff-1.1-draft-01.pdf>>.
6. Internet2 Middleware Architecture Committee for Education, Directory Working Group, "eduPerson Object Class Specification," June 30, 2008, <<http://middleware.internet2.edu/eduperson/docs/internet2-mace-dir-eduperson-200806.html#eduPersonAffiliation>>.
7. See "VIVA Virginia!" InCommon case study, April 8, 2008, <[http://www.incommonfederation.org/docs/eg/InC\\_CaseStudy\\_VIVA\\_2008.pdf](http://www.incommonfederation.org/docs/eg/InC_CaseStudy_VIVA_2008.pdf)>.
8. The Internet Society, "Trust and the Future of the Internet," August 2008, <<http://www.isoc.org/isoc/mission/initiative/docs/trust-report-2008.pdf>>.