BY DAVID L. WASLEY

TECHNICAL BUILDING BLOCKS TO ENABLE E-BUSINESS

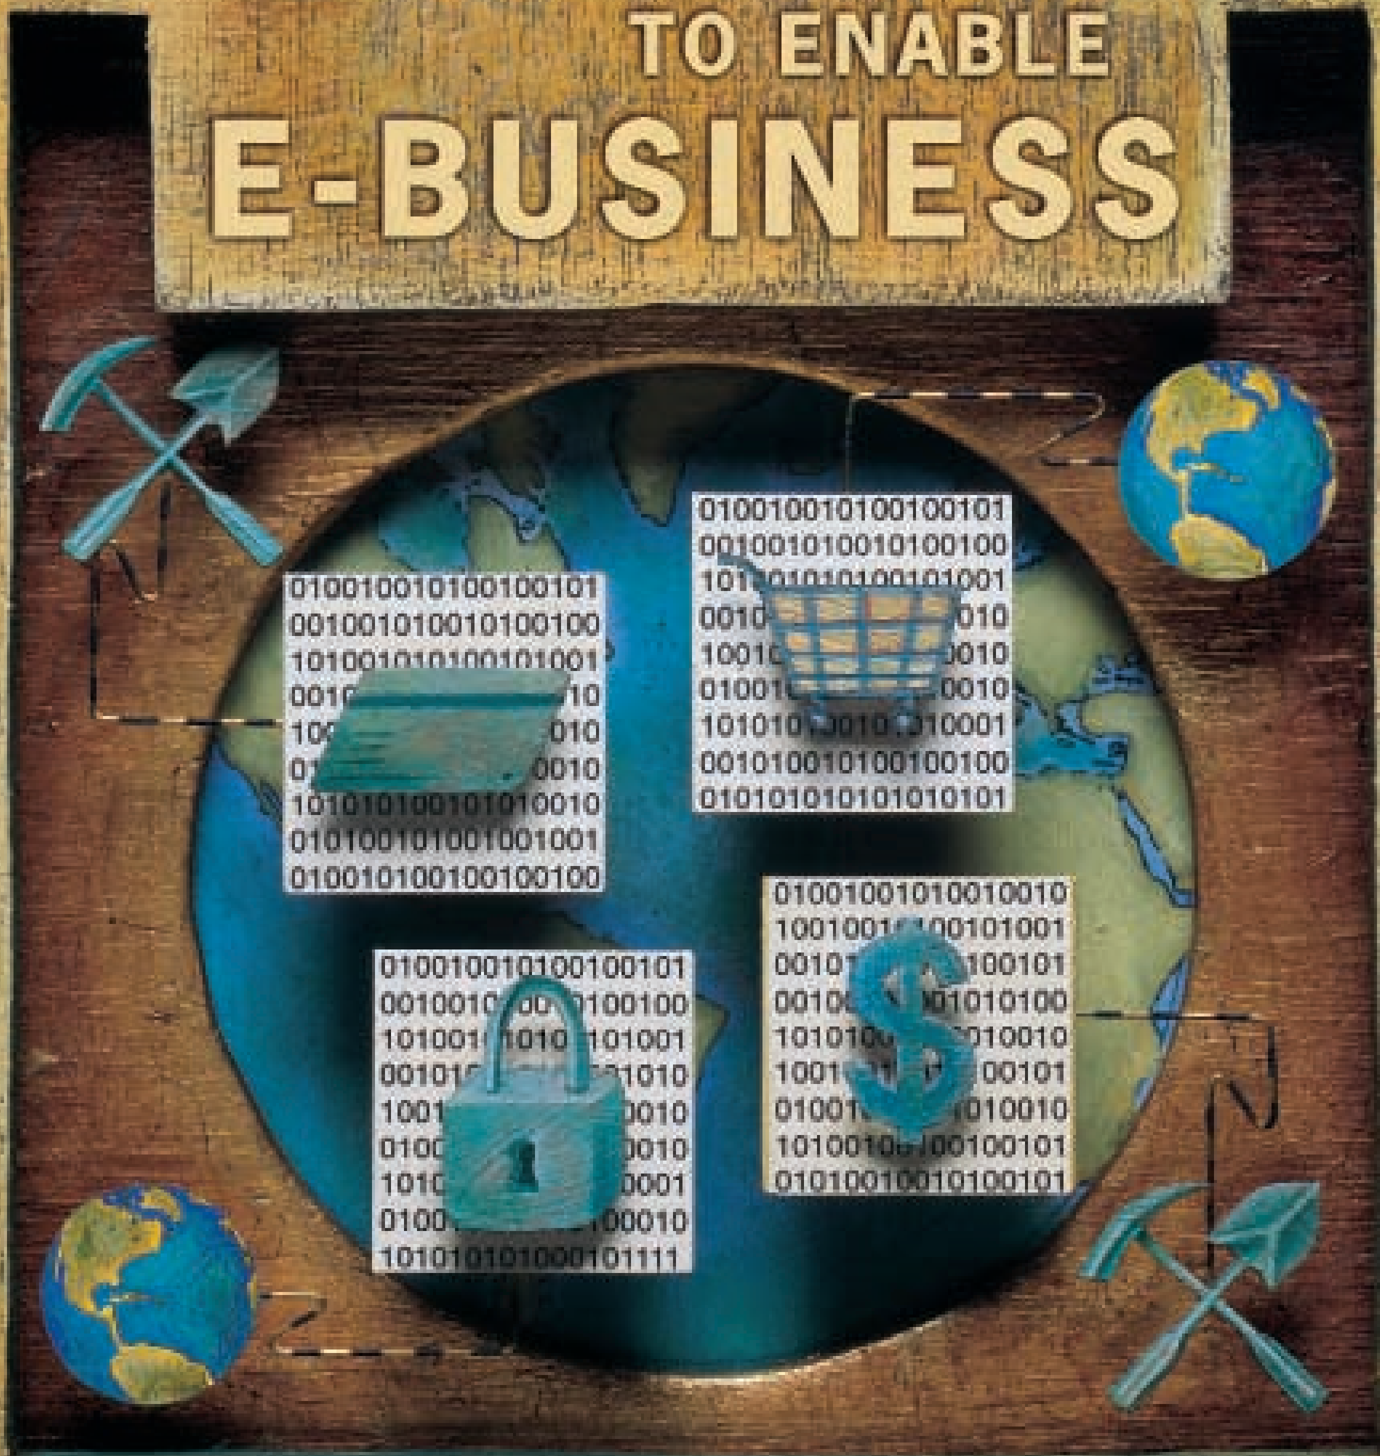

ILLUSTRATION BY MARIA RENDON

Access to on-line information resources and the use of electronic transactions increasingly augment the operation of modem colleges and universities. These resources and functions are important both internally and in collaboration with external partners. To reap the full benefits that these resources and processes can offer, all members of the campus community must be able to function easily and securely throughout this rapidly growing electronic information environment. Although campuses have made great strides toward developing robust and ubiquitous communications network infrastructures, the rich information environment envisioned in *Sustaining Excellence in the 21st Century* (Katz and West, 1992) will not be realized unless equally robust and ubiquitous enabling services are developed within the infrastructure of our campuses and throughout the higher education community.

I define the electronic information environment as that set of electronic information services, on-line resources, communications services, applications software, and workstations that enable us to teach and learn and work more effectively and without the constraints of time or place. Within this environment, we need directories and other finding aids, credentials that can establish identity and roles for both consumer and supplier, and a myriad of other supporting services.

Infrastructure support services will enable easy and secure access to information resources, support authorized and verifiable transactions over the network, and make possible appropriate management of licensed materials and other intellectual property. With a coordinated approach, we will be able to leverage investments being made already in new applications on our campuses as well as across the higher education community.

As the electronic information environment grows and expands, it becomes necessary to acknowledge and implement certain constraints on users of these resources—the "network citizens" who navigate that environment. We must develop analogues of the administrative controls with which we are familiar in the traditional environment. These controls should be supported by generalized infrastructure services put in place and managed by the institution.

Existing access control and authentication mechanisms are largely a legacy of older centralized technology. Traditional institutional applications have been developed from the ground up, providing for all aspects of the process in an idiosyncratic way. Today we understand that many of the components of these applications have a great deal of commonality. These include a secure and reliable way to affirm that users of our resources are who they purport to be and that they are authorized to use the resources to which they seek access. There must be a source of definitive information regarding their affiliation with the campus as well as other business-related data. There must be directory services to help them find the resources they need. There must be standards and supporting services for encryption to secure data transmission and to create the digital equivalent of a personal signature. There must be efficient mechanisms to support accounting for the use of a wide variety of network-based resources as well as services that can be supported by network communications.

Lack of common general solutions for authentication, authorization, directory services, and encryption will impede the development of a broad range of information resources, from client-server financial systems to digital libraries. If properly designed, these technical building blocks can be combined and extended to form a set of common services, or "middleware," that will enable a wide variety of complex capabilities that are available almost transparently to the end user (see Figure 1). These new and critical middleware services will enable members of the academic community to become true universal network citizens and roam freely and securely, without undue let or hindrance, throughout the emerging electronic information environment.

The availability and use of these and other building blocks also can form the basis for more robust and efficient management of institutional resources. Rule-based authorization processes can support distributed operation and responsibility. Well-designed and reliable support services can help the campus protect intellectual property as well as institutional data. Data encryption standards and services will enable the reliable use of e-commerce within the institution and with external partners.

Differences among campuses in their information environment middleware services tend to inhibit sharing of resources, a tendency that will become exacerbated as critical new shared resources such as virtual digital libraries begin to emerge. Currently a faculty member or student at one campus who is visiting another campus must make special arrangements in order to make use of many of the services at the second campus. If the services require access control, the individual must acquire yet another password and account name. Coordinated middleware services can allow the identity and affiliations of any member of the higher education community to serve them wherever they may be working at the moment.

I will continue to use the metaphor of a network citizen traveling throughout the electronic information environment to illustrate the basic middleware building blocks. It is with the goal of catalyzing the development and implementation of the best technologies to implement these services within the higher education community that we begin this management tour of the electronic information environment.

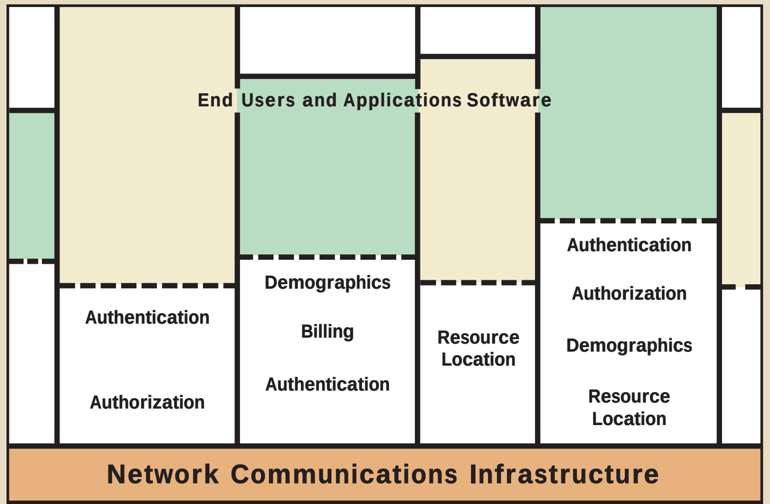

Figure 1. Existing Applications with Independent Solutions to the Same Set of Problems

The Terrain: The Electronic Information Environment

The electronic information environment is an increasingly complex territory in which valuable resources can be found, but it is foreign to our usual senses. Many of the familiar problems and processes are found here but take a different form. What is our identity? What can we do here? How do we find resources? How do we traverse the environment with safety? How do we account for what we use? Do we share a common reference for time or place?

The foundation of the electronic information environment is the complex of interconnected networks (including our campus networks) known as the global Internet. Accessible from desktops both within our campuses and via the Internet are an ever growing variety of databases, information servers, vendors doing business via e-commerce, computational servers, and a panoply of applications programs from e-mail to collaborative systems based on interactive virtual reality. E-commerce is developing rapidly in many forms and will enable reliable transaction of institutional business as well as new opportunities for individuals.

Our helpers in traveling the information environment are client software and programs that interact with other electronic entities by means of well-defined protocols. Just as in the analogue world, these protocols provide for successful interoperation among widely differing entities. The World Wide Web is an excellent example of this concept: the same set of complex text, graphics, audio, and video can be found and viewed on almost any modern computer regardless of what kind of computer might hold them. One significant advantage of this layered approach to complex systems is that individual pieces of the environment can be built or replaced without affecting the rest of the infrastructure. Another advantage is that end-user platforms can be tailored to different user communities and application needs. This layered approach must guide us in the development of essential information environment support services.

Passports: The Authentication Service

The fundamental support service that will empower the network citizen is a universal, distributed, reliable, and robust digital credential system. Just as travelers must carry their passport when entering a different country, network citizens should carry a recognized credential when entering the electronic information environment. This credential, which can be validated by any service encountered, will ensure the authentic identity of the individual holding it.

Passports are recognized as authoritative around the world because the issuing authority is recognized by international treaty. So too should the network credential be recognized as authoritative so that network citizens can visit various sites and resources without having to resort to site-specific identification. Instead of an international treaty, there must be agreements among cooperating administrative domains to trust each other's credentials. Such agreements are based on understanding the methodology and management of the process for issuing the credentials.

The concept of a digital credential that refers to a single electronic identity is powerful because it can enable easy access to resources for the traveler. The credential itself may not carry much information about the individual, much as the traditional passport carries not much more than a picture, an address, and a passport number. However, that passport number can be the key to determining additional attributes about the holder. In an analogous way, the digital credential should have an identifier (e-ID) that is unique to the credential holder and can be used as a key to discovering further identity information.

Relying on a single credential is fraught with the potential for abuse if it is mishandled. Since e-IDs will let travelers into many places, they must be issued only with strong assurance that individuals are who they claim to be. Once the e-ID is issued, the holder must recognize that it is an extension of his or her personal identity, protect it carefully, and never lend it to another person.

The fundamental support service that will empower the network citizen is a universal, distributed, reliable, and robust digital credential system.

Even with reliable credentials, it will be desirable in many cases for an individual to have several e-IDs to be used for different purposes. All individuals and their organizations must consider the potential impact of a compromise of the security of any e-ID. For example, the manager of an administrative computing system might have an e-ID that carries with it special privileges. That e-ID should be used sparingly and only in conjunction with that system so that compromise of the more powerful e-ID would have consequences of a more limited scope and could be dealt with more readily. The system manager would use a less powerful e-ID when doing routine work, such as checking e-mail or writing reports.

An e-ID represents an assertion on the part of a registration authority that a known individual or entity is represented by that e-ID. If the registration authority is well designed and reliable, then services in cooperating administrative domains, or realms, can be comfortable accepting those externally registered e-IDs. Thus, in the general case, the e-ID must indicate not only the individual but also the registration authority that issued the e-ID. Ultimately this concept could be extended to commercial registration authorities so that, for example, high school students registered with a local electronic notary might be able to use their own credential's e-IDs in submitting electronic applications for admission to a college or university.

Once a campus has a robust digital credential issuance and authentication service in place with reliable operational support, all of the important campus server applications should be adapted to use this system. Parts of this process could take a long time since many applications are vendor supported and vendors are not yet responding to this broader vision. Until there is a well-established generalized authentication service, each server or application has little choice but to implement its own idiosyncratic method of authentication, and the network citizen must deal individually with each one encountered.

Tickets and Passes: Attribute and Eligibility Databases

Eligibility, as distinct from authentication (which warrants the identity of a particular user), is an equally important support service in the electronic information environment. Eligibility helps define what network citizens may do or that to which they may gain access. It may be based on a person's e-ID or on any set of attributes associated with that e-ID, such as the person's affiliation with or within the institution, role, or current status. Often it is desirable to separate authentication service from eligibility and attribute functions, and it may be desirable to support them on different servers.

Eligibility does not necessarily imply authorization. The network citizen may be eligible to gain access to a resource, but at the time access is requested, the service might be oversubscribed or otherwise unavailable. For example, holding an airplane ticket may not result in authorization to board that flight. Authorization is a function of the service or application based on specific business rules. These rules might take into account a variety of factors in addition to the individual's identity, such as time of day, location of the individual, or the availability of resources.

In the historical model of mainframe computing, authentication (often referred to as access control) and authorization were often closely linked. Authorization usually was determined by loose association of attributes with the user account identifier (for example, "root" or "user, group, and other" in Unix systems) or by tables of authorized users. Since each system kept its own set of authorization data, management of these data in a consistent way across a large number of servers and client platforms was problematic, although a number of projects attempted to address aspects of this dilemma. Little thought was given to generalized information environment management mechanisms until the number of different systems and services began to grow dramatically.

Today we need generalized and scalable network-based mechanisms not only for authentication but also for eligibility and attribute information. As an example of how this might be provided, suppose the network citizen requests access to a restricted database. The database server can simply query a specified attribute server, supplying the network citizen's e-ID, and receive information on the roles, affiliations, or other attributes defined for that individual. The database server then can apply a prescribed set of business rules to determine whether that particular network citizen is to be allowed access to the restricted data.

Today we need generalized and scalable network-based mechanisms not only for authentication but also for eligibility and attribute information.

Suppose the network citizen's affiliation with the institution is "full professor and dean of the college." This might imply eligibility to view college budgetary data and approve spending plans and personnel actions, as well as to gain access to academic records and information. The same person might also be appointed to a multi-campus task force on student diversity and because of that affiliation be eligible to retrieve sensitive student ethnicity and gender data. A database server that returned roles and affiliations for any given institutional e-ID would make management of this type of generalized eligibility much easier.

Such eligibility or affiliation data might be useful in transactions external to the institution as well. Appropriate individuals in any department could be given authority to submit electronic data interchange (EDI) purchase orders over the network, for example. A different set of individuals could have authority to approve payment of EDI-based invoices electronically. Workflow systems could identify individuals who need to be informed of such transactions for possible post-transaction audit.

Whereas administration of an attribute service should be hierarchical and coordinated centrally, responsibility for actual eligibility data with respect to any given service should be distributed to conform with campus management structure. This is primarily an administrative rather than a technical issue. Suffice it to say that developing and managing a database service that combines all important attributes, roles, and affiliations and allows them to be managed by the office of record would provide important generality while maintaining the institution's established administrative structure.

Who's There? The Demographics Database Service

One particularly vexing problem is the maintenance of accurate personal data, such as home address, campus office address, or preferred e-mail address. Today there are far too many different databases wherein the same data are entered, usually by different individuals, from separate forms that must be filled out repetitively by the person who actually "owns" the information. For example, in many cases the same individual is both a student and an employee, which means that the same personal data often are maintained by entirely separate offices. Clearly it would be desirable to have a single comprehensive database of record that would hold information about all members of the campus community and in which personal data could be maintained by the relevant individual or appropriately designated staff.

With strong authentication and a well-designed attribute service, it would be possible to build such a demographics database system. It might well be combined with a basic eligibility service so that a single database system supports both. Fields within a record would have associated rules for access or updating. All institutional applications that require use of personal data would use this comprehensive demographics database by either periodic downloading or indirect relational reference. In particular, on-line directory services for locating campus community members would use this authoritative database as their data source.

Campus network citizens could be responsible for maintaining all of their own personal data and also could check on the completeness or accuracy of other attribute data, such as payroll title or salary, status toward a degree, or the parameters of their employee benefits.

Maps and Guidebooks: Directory Services

Information services and resources abound on campuses and beyond. The community of users changes and moves about, and the topology of the network changes periodically. How does a network citizen find anyone else or any particular service or information?

The Domain Name Service was the first widespread directory service in support of the information environment. Other common directories exist today, such as directories of people and searchable databases of information resources. Many more kinds of maps and guidebooks are needed to serve the new and complex information resources we are deploying. For example, the network citizen might want to find an on-line copy of Van Gogh's *Child with Orange* and the nearest color printer that she or he is eligible to use and that is capable of high-resolution, large-format printing. New resource location services must support more complex data and search strategies.

A well-managed set of directory servers and search engines for information objects will help the network citizen discover resources and navigate easily throughout the electronic information environment.

Safe Passage: Encryption and Digital Signatures

Authentication and access control alone are not sufficient to guarantee safe passage throughout the electronic information environment. Passwords can be guessed or stolen, data can be monitored in transit, and identities sometimes can be forged. The strongest defense against these challenges involves the use of modern encryption methods to ensure privacy of data transmission and create the digital equivalent of a pen-and-ink signature.

Our networks are truly open systems, which is one of their strengths as well as a source of many vulnerabilities. Any transmission might be intercepted, and any data received might be questionable. Attacking computers on the Internet has become a rampant obsession among certain curious and occasionally antisocial groups. Fortunately, modern encryption technology offers potential solutions for most of these concerns.

Encryption of data while in transit can protect privacy as well as the confidentiality and integrity of data. The technology required to implement data encryption has become commonplace and should be considered for all administrative or other institutional business applications. Encryption alone may not guarantee authenticity, however. We need the equivalent of a seal or at least a recognized signature associated with the contents of the document.

A digital signature must be some set of data that cannot be forged and that binds the contents of a digital document to a specific individual, role, or other entity. It must be something that only the signing entity could have created and must be verifiable by anyone in the electronic information environment.

Standards now exist for creating this type of digital signature using public key cryptography (PKC). Clearly digital signatures of this sort would enable a wide variety of institutional business to be transacted over the network with at least as reliable verification and auditability as we have now with paper forms and manual signatures.

A public key infrastructure (PKI) is critical to enabling the use of encryption and digital signatures. Fundamental to the PKI is a unique pair of very large prime numbers, generated for each credential holder, that are keys used for encryption and decryption of information. Software on the network citizen's workstation generates this pair of encryption keys and gives one of them (the "public" key) to the PKI certificate authority (CA). The other key (the "private" key) is closely guarded by the network citizen. The CA stores the public key in a directory along with the PKI certificate. Upon request, the CA or directory server provides any registered user's public key to any application that needs it.

In addition to support of the local community, a CA must have a way to find other trusted CAs on other campuses or anywhere else in the electronic information environment. This ability is part of an overall PKI and may be implemented in a number of different ways.

The basis for trusting a traditional signature is either direct knowledge or the ability to look it up in an archive maintained by a trusted authority, such as a bank. In the digital world, trust must be established through preestablished contracts based on mutual understanding of business practices, the basis for registering individuals with the CA, and the viability of the cooperating CAs.

Until and unless encryption mechanisms and support services are in use, no one should send anything of value or of a sensitive nature, such as a credit card number, over the data network.

The resulting so-called web of trust must be scalable to millions of users in thousands of locations. This can be achieved through a hierarchical model wherein a community-based CA registers subordinate CAs after verifying their viability. A consortium of college and university campuses, for example, could operate such a certificate authority and the associated services on behalf of its members. This CA in turn could register under national and international CAs to allow fully general and trustworthy access to verifiable public key directory servers anywhere in the world. A campus PKI with its root certificate authority also could register subordinate CAs for departments, the library, or special-purpose requirements.

Until and unless encryption mechanisms and support services are in use, no one should send anything of value or of a sensitive nature, such as a credit card number, over the data network.

Sharing Limited Resources: License Servers

The network citizen is now equipped with the basic tools for verifying identity, invoking authority, and concluding transactions safely. These capabilities enable easy and appropriate access to a wide variety of resources within our electronic information environment. Unfortunately, not all of those resources are without significant cost to the institution.

A traditional library might hold five copies of a popular book or journal. Ideally this would be enough to meet the peak demand at any one time for this resource. It would not be economical to purchase one copy for every registered patron, yet we often provide software and other resources in this cost-inefficient way.

It is quite possible for software or other digital resources to be purchased by subscription, much like printed documents today. The cost of such a subscription would be based at least in part on the size of the simultaneous user community. For example, a physics department might purchase a subscription for twenty simultaneous "users" of a virtual physics laboratory software package. During laboratory classes, the students make use of the software on computers located in the facility. In the evening, when doing homework, up to twenty students could make use of the same "subscription" to run the software on their own personal computer.

One way to enable this type of sharing of expensive resources is a network-based license server (NLS). The NLS serves as a clearinghouse to ensure conformance with the terms of the institution's subscription. An early implementation of the NLS concept was available with the Apollo domain computers. Macintosh and PC versions of license servers offer similar capabilities. Standardizing on an openly available NLS technology would enable publishers to develop products that could fit readily into our electronic information environment. A single NLS could moderate access to a wide variety of licensed resources, including databases and documents as well as software.

Paying the Bills: Automated Debit and Unified Invoicing

Much of the electronic information environment today is accessible without direct cost to the network citizen. However, as the real costs become significant, the campus may need to find efficient ways at least to account for the usage of expensive resources and possibly allocate some of the acquisition and support costs toward the end users. A building block that could help achieve this is an efficient network accounting server.

Ultimately every member of the campus community, as defined in the campus's demographics database and corresponding eligibility servers, could have one or more "virtual accounts." Transactions for services would be posted to a designated network accounting server using encrypted data flows. A wide variety of network citizen services, from print-on-demand syllabi to lunches and storehouse items, could be accounted for in this way.

The network citizen should expect a single statement each month describing all services used and any costs incurred anywhere within the institution. This might include transactions with external partners as well. Eventually this monthly accounting could result in an automated debit against an external financial service, much like debit cards are used today.

What Time Is It? The Network Time Server

It may not seem obvious but many of the services we are developing need to have a common frame of reference for time. Billing information, for example, must show accurately the date and approximate time of the transaction. Network management information often needs time stamps to be accurate to within a few milliseconds. Electronic postmark or notary services must have an auditable date and time guaranteed to be within a known degree of accuracy. Thus, an important element in the set of enabling services within our information environment is the network time server.

Technologies in support of network time services exist today and are deployed on most campuses. However, not all of these are synchronized with each other or with a universal time standard, such as the National Institute of Standards and Technology broadcast standard time service.

Even where synchronized network time servers exist, not all essential end systems can take advantage of them yet. Campus information technology managers must understand the importance of this element of distributed systems and take appropriate steps to ensure integration of this service.

Our network citizens may wish to set their own "electronic watch" from this service as well. Most modern workstations can be configured with automatic utilities to accomplish this.

Where to Now?

The building blocks I have identified are all part of a larger set of standards that comprise an information technology architecture. The basic enabling services include the following features:

- A coordinated set of authentication servers
- An attribute and demographic database server on

each campus that can include a wide variety of information, including affiliation and eligibility data

- Directory and resource location servers
- PKI servers that manage certificates and public encryption keys for individual users and make possible digital signature verification
- Electronic license servers in support of site-licensed software, library materials, and databases
- Billing transaction servers that can handle a large volume of small-value debit records extremely efficiently
- Time servers and digital notary servers that form the basis for reliable and verifiable on-line content and digital institutional archives

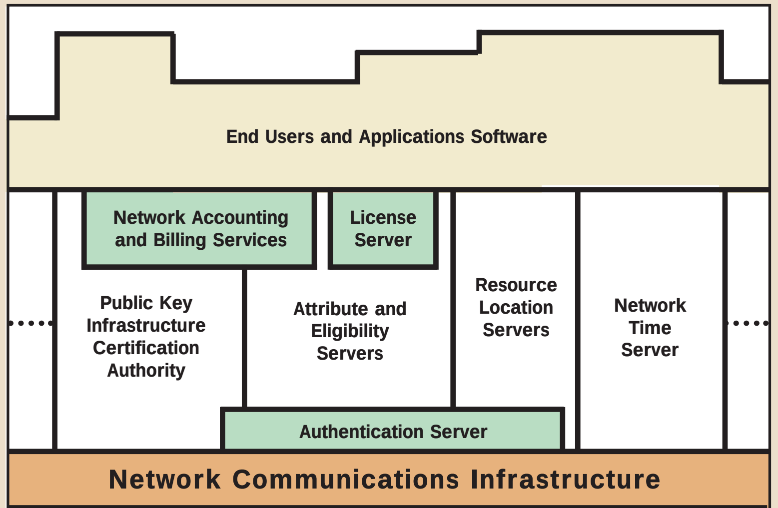

Figure 2. Information Environment Building Blocks with Common Interoperable Solutions for Basic Supporting Services

Figure 2 shows how the building blocks might relate to each other and to the communications system and applications programs. Other building blocks might include software version control servers to ensure that campus users have access to the latest version of critical application programs, and alias servers to support consistent mapping between e-IDs and e-mail addresses or traditional identifiers such as employee number or student ID number. Current and potential technologies behind each of the building blocks described above are in different states of development and deployment.

With persistent vision and cooperative efforts, we can refine and deploy appropriate versions of all these enabling services over the next few years. If we do not start now, it may become very difficult to develop and retrofit a coordinated set of these services later. There is much to be done before network citizens are fully empowered. e

Note

M. Stuart Lynn, a colleague and friend to most of us in the information technology community, provided ideas, advice, counsel, and encouragement in the creation of this chapter.

Reference

Katz, R. N., and West, R. P. *Sustaining Excellence in the 21st Century: A Vision and Strategies for College and University Administration.* Boulder, Colo.: CAUSE, 1992. [www.educause.edu/ir/library/pdf/PUB3008.pdf].

The "E" Is for Everything **is the second volume in the EDUCAUSE Leadership Strategies series, sponsored by PricewaterhouseCoopers. A complimentary copy of the book has been sent to the primary representative at each EDUCAUSE member institution and organization; additional copies may be purchased from Jossey-Bass (see http://www.josseybass.com) or from EDUCAUSE (see http://www.educause.edu/pub/pubs.html#books).**