Access Management for Networked Information Resources

by Clifford Lynch
Executive Director, Coalition for Networked Information

Introduction

Over the past year the Coalition for Networked Information (CNI) has been leading a broad-based exploration of access management and authentication issues that are emerging as institutions share resources and license access to published materials through the Web. Scores of institutions and content providers participated in the discussion and contributed to the development of a lengthy white paper that is now being cast into final form; both the draft from the spring of 1998 and, when completed, the final version, can be found at CNI's website <http://www.cni.org/>. The interested reader is urged to consult this white paper for a much more in-depth discussion of the issues summarized here.

This article is intended to briefly describe the crux of the access management problem and the two major architectural approaches that are now coming into use to address it. It will also devote considerable emphasis to important policy issues that emerge as part of the consideration of approaches to access management; these issues are real, immanent, and potentially controversial. Research libraries have both an opportunity and an obligation to provide their patron communities with leadership in understanding the issues, and to work with these communities to formulate appropriate policies proactively and thoughtfully, rather than in response to a crisis or scandal. Many universities are now implementing institutional authentication and access management systems to support specific application and security requirements; typically this work is being overseen by the institution's information technologists and it is often motivated by policy imperatives and priorities that are very different from those of the library. Because authentication and access management represents a common infrastructure that will ultimately need to serve many purposes, it is essential that libraries and information technology groups establish a dialog now to ensure that the systems which are ultimately deployed honor the complete set of campus community requirements.

This article does not discuss the technical details involved in authentication and access management—they are complex and the state of commercially available technology is continually changing. Suffice it to say that, as of this writing (November 1998), there does not seem to be any simple, inexpensive, ready-to-deploy comprehensive solution for authentication and access management; while there are many promising components available, management complexity, user training and acceptance, system integration, and cost are still major issues. Organizational readiness within the content provider community is also a barrier to the use of advanced technology-based approaches; content site operators are no more able than libraries to obtain commercial turn-key solutions. But as libraries offer a growing portfolio of network-based information resources, they face immediate and pragmatic needs for access management. In developing strategies, it is important to recognize that technical choices about access management have policy implications.

The Access Management Problem Defined

A library negotiates a license agreement on behalf of its patron community—for example, the faculty, staff, and students of a university. This license agreement gives members of the community the right to use some network-based resource—perhaps a commercially offered electronic encyclopedia or scholarly journal at a publisher website, or a specialized research database at another university that is part of a resource-sharing agreement. Users connect to the site using web browsers running on their personal workstations or on public workstations. As the site hosting the licensed resource processes these requests for web pages, it must determine whether the requesting user is in fact a member of the appropriate user community. If so, he or she is given access; if not, access is refused.

Obviously, this basic scenario can be elaborated upon—for example, there is a need for finer-grained systems that can be used to allow only registrants in a specific multi-institutional course access to electronic reserve materials provided by the library at one of the institutions. These more elaborate situations raise additional problems both in terms of scale and system design and are useful to have in mind to understand the extent of the requirements; however, we will concentrate on the basic scenario here.

Access management needs to be routine and easy to implement; once a contract is signed, lengthy technical negotiations between institution and content supplier should not be necessary before users can have access. In a world of networked information resources, access management needs to be a basic part of the infrastructure, and must not become a barrier to institutional decisions to change or add resource providers.

While it may be a complex administrative question to determine the membership of the patron community in a world of extension students, affiliated medical professionals, visiting scholars, and others who may blur the boundaries, libraries have a great deal of experience in devising these policies. The problem here is a technical one between two computers: how does the resource host machine actually determine that the user at the requesting workstation is really a member of the relevant community? In other words, given a policy for determining community membership, how is this policy implemented technically? In the physical world of circulating materials, the “technical” problem is dealt with through the presentation of a library card indicating that the holder is indeed an authorized borrower; this library card is issued upon the presentation of policy-specified documents that identify the user and prove that he or she meets the policy criteria for authorized community membership (a faculty ID, for example, plus a picture ID).

Some Approaches That Don’t Work Well

Historically, many computer systems employed user IDs and passwords. When a user ID was first defined, it would have certain rights and privileges associated with it. The user would then supply a password upon demand to demonstrate his or her rights to use a given ID. Leaving aside technical problems associated with the use of passwords on an insecure network and the possible technical remedies, there is a more basic architectural issue. We are in a world where users may routinely want access to many different network resources controlled by many different organizations, such as publishers; but with the ability to follow links on the Web it may not even be clear to the user which publisher owns what resource. Users will not be able to remember and manage a large number of different user IDs and passwords issued by different publishers; further, each publisher will need some cumbersome procedure to validate the user as a community member prior to issuing a user ID and initial password (for example, involving mediation by local library staff). This does not scale up in a practical way to a world filled with electronic information resources.

The library might issue the user a single ID and password for all licensed external network resources, and then transmit lists of these IDs and passwords to each content supplier so that the supplier can use the list to validate users. However, there are architectural problems here, as well: a very large number of publishers would need to be notified every time the list of user IDs and passwords changes, with inevitable timeliness and synchronization problems. Also, each publisher takes on an enormous responsibility for protecting the security of the ID/password list; a security breach at any publisher will mean a security breach at every publisher doing business with the library for networked resource access, an unacceptable liability. Here the networked information scale-up means that too many independent parties must rely on each other to maintain security, and that the cost of accurately maintaining a synchronized common-access management database will be very high (to say nothing of the standards that would need to be put in place to make such a shared database a reality).

As a stopgap measure, many institutions are currently using electronic “place"—the user’s source IP network address—as a substitute for other methods of demonstrating proof of community membership. If the user’s connection request came from a network that belonged to the university, it is assumed to be from a community member. Network ownership does not change too frequently, so maintaining a list of valid network numbers and making this available to publishers is a tractable administrative burden. Further, since the list of network numbers, unlike IDs and passwords, doesn’t need to be kept secret, there is little interdependence. This approach works well and, indeed, has the virtues of simplicity and transparency to the user, as long as all users come to the resource providers through the campus or library network. However, in an era when many institutions are discontinuing dial-up modem pools in favor of commercial Internet service providers, and when new access technologies to the home, such as cable television modems for Internet access, are beginning to deploy, a great deal of legitimate user access now takes place from sources other than the campus or library network. With growing needs to support distance learners, part-time students who may do academic work from their place of business, and people who want to exploit the “anytime, anywhere” promise of networked information resources from their homes, limiting access by source IP network address disenfranchises more users every day. For most universities, for example, it is clearly no longer acceptable to tell community members they can only access networked information resources from on-campus workstations.

The Two Emerging Architectures for Access Management

Two general approaches are emerging to address the access management problem. The first is the use of proxies. Here, an institution develops an internal authentication system and uses it to validate user access to a special machine on the institutional network, called a proxy. Once a user is validated, he or she can ask the proxy to contact an external resource; in some variations, the proxy mediates the user’s entire interaction with the external resource, while in other (perhaps slightly less secure and user-transparent) variations the proxy drops out of the interaction after making the introduction. From the resource operator’s point of view, it is only necessary to know that the institution’s proxy server can be trusted to pre-validate all users before contacting the external resource host. Proxies have a number of advantages: authentication is an internal matter to the institution and external service providers need not be concerned with the details of how this is accomplished; at least in theory, the institution has complete flexibility in deciding what resources a user can have access to through the proxy and when; and the proxy can act as a central control point for the institution in access management. On the negative side, the proxy is a high-impact central point of vulnerability for outages or capacity problems (though, of course, it’s possible to have multiple proxy machines), and configuration management in the proxy can become extremely complex and labor-intensive, particularly if not every valid proxy user has access to all resources. Further, proxies do not eliminate the need for an authentication system; they only isolate its scope to the proxy and the members of the institutional community.

The other approach is based on credentials. The basic idea here is that the institution issues each user the electronic analog of a community ID card. The user, or the user’s browser, presents these credentials upon demand to any resource provider that requests them; the resource provider can then, through a fast electronic transaction, validate the credentials with the issuing institution. The validation process shares with proxies the vulnerability to outages or capacity problems on the part of institutional systems that verify credentials, though these vulnerabilities are more circumscribed. A compromise of the credential-verification system may be more serious than compromise of a proxy: proxy compromise usually means that unauthorized users get access to resources for the period during which the proxy is compromised, while a compromise in the credentialing system may well mean that new credentials need to be issued to the entire authorized user community. The major practical difficulties with the credentials-based approaches, however, involve technical problems, standards, cost, and software integration.

The simplest credentials-based approach would be to have the institution just issue the user a user ID and password for external resources, and to have external resource providers validate the ID/password pair with the institution (as opposed to having the institution distribute lists, as discussed earlier). This reduces, but doesn't eliminate, the interdependence among external resource providers in the maintenance of security; further, standards don't exist for such a validation process and no off-the-shelf software supports it. The industry is moving towards a technology based on public key cryptographic certificates (X.509) in off-the-shelf software, and, at least in theory, this should work well for personal machines (but not for shared or public workstations, which would have to be handled some other way, such as IP source address checking, at least until a much more complex hardware-based infrastructure becomes widely deployed). X.509 removes the interdependent security issue because credentials are computed for each use from information that is held only by the user and never directly transferred to the content provider.

The problems with this approach include integration with browsers, the mechanics of issuing and distributing these electronic cryptographic certificates to users, the cost and complexity of acquiring and operating the infrastructure for managing and validating public key certificates, and government regulation of the export and use of cryptography in various countries, which causes problems in an increasingly international world of scholarly resources. All of these problems with cryptographic certificates are slowly but steadily getting better (except, perhaps, for the government regulation issues), but, at least today, implementation of such an approach is an enormous challenge. Finally, it is important to note that X.509 was really designed to support applications such as electronic commerce, and there are some significant problems in relation to the policy problems discussed in the next section.

Policy Issues in Authentication and Access Management

The development of any access management strategy raises policy issues in areas such as privacy, accountability, and the collection of management data. It is important to recognize that libraries must decide whether to address these issues through legal means (that is, by negotiating contractual obligations on the resource supplier as part of the license agreement), through technical means (for example, by making it impossible to collect personal data), or by a combination of the two.

Library experience in other contexts offers some insights. For example, libraries have aggressively championed and defended the privacy of their patrons. They have done this through both legal means—by developing privacy policies and by requiring a legal subpoena as a condition for divulging records—and also by technical means by not keeping historical circulation data on an individual basis, which then limits the amount of information that they can be compelled to disclose under any circumstances. Patron privacy has been such an important value to libraries that they have, by and large, used a dual technical/legal strategy to provide their patrons with the strongest possible protection.

In the electronic environment, it is easy for a publisher to track the use of content in great detail—what material is being viewed, for how long, and how often. Depending on how the access management system is structured, particularly when credential-based approaches (and specifically, routine implementations of X.509 certificates) are employed, the publisher may be able to correlate this usage information to a specific individual by name, to a long-term "pseudonymous" identity that the publisher can link to an institution but not to an actual individual within that institution's community, or simply to a transient and anonymous member of the institution's user community. Clearly, the license contract between institution and publisher can speak to the collection, retention, use, and disclosure of such usage data on a policy basis, but libraries and patrons may find it desirable to limit the ability of the publisher to collect information by the design of the access management system on a technical basis, as well.

One important point needs to be made about user privacy that underscores the need for contractual constraints even in conjunction with an access management system that provides some level of anonymity: often users will make their actual identities known to content providers for other reasons, independent of the access management system, such as to take advantage of email-based current awareness services or personalization options in a user interface. The access management system is not the only way in which privacy can be compromised, or bargained away for increased function or convenience.

A license agreement represents a commitment on the part of the licensee institution to honor the terms of the license, and to educate members of its community about their obligations under the license. The publisher and library share a need for some level of accountability by community members: if a single user accesses publisher content hundreds of times a day from three continents, it’s likely that something is wrong; perhaps that user doesn’t understand his or her obligations, or perhaps credentials have been compromised. There is a need for the publisher and the library, acting together, to be able to investigate such situations effectively, and, if need be, to block access by specific individuals who seem to be violating the terms of a license agreement. But, in order to do this effectively, a publisher needs to be able to at least provide an institution with enough information (such as a pseudonym) to permit the individual in question to be identified; note that this does not necessarily mean that the publisher can directly identify the individual in question, but only that the publisher can provide the institution with enough information to identify the user. The need for accountability contradicts, to some extent, the mandate to design anonymity into an access management system, and argues for a pseudonymous approach. This can be achieved with credential-based approaches (though it is not the standard model for X.509 certificates, for example), but is more difficult with proxy-based architectures.

Finally, there is the issue of management data. Electronic information resources promise libraries much more accurate and detailed data about what content is actually being used and how often—though even at this level libraries may want to make contractual stipulations to protect patron privacy; for instance, it is not clear how many universities would be comfortable having a list of all the articles read by members of their community (with frequency counts) posted on the Web every week. But greater problems arise when libraries want to have these usage statistics, at whatever level of aggregation, demographically faceted—for example, to drive internal cost allocation processes within the library’s institution. The simplest solutions are often to pass demographic attributes to the publisher along with identities or pseudonyms and to get the publisher to do the work of generating management data for the library—but this path can rapidly compromise the privacy of pseudonymous users by making them more identifiable and, if actual identities are used, it makes the privacy problem even more acute by raising the stakes on the amount of information disclosed. In discussing this issue as part of the CNI authentication work, it appears that many libraries, because of the privacy considerations, are backing away from the continuous collection of demographically faceted usage data, recognizing that they can always do specific studies in order to obtain “snapshot” information, much as has been done with journal usage historically.

Demographically faceted usage data offers a good illustration of the scaling issues that we face in the move to network-based electronic information. If a library were dealing with just a single publisher, it would be reasonable to have the publisher return detailed (transactional) usage data with pseudonyms to the library. The library could then look up the pseudonyms to obtain demographic data and summarize the transaction data into management information. Here privacy is protected by library policies and does not depend on the publisher, who knows only pseudonyms. But this is intractable with hundreds of publishers, each supplying transaction data at different time intervals and in different formats (and even based on different models of the transactional events being tracked and logged). Without very detailed standards, which don’t exist today, and very rigorous adherence to such standards by the publisher community, libraries needing demographically faceted usage data will be forced to export demographic information to each publisher.

Conclusion: Towards Practical Short Term Solutions

Credential-based solutions require substantial work by publishers; at least a few of the larger publishers have a sophisticated understanding of the access management issues and are, I believe, prepared to work with libraries and universities to put credential-based solutions in place. However, many content providers offer only IP source, network-based authentication and have a limited understanding of the broader issues. One of the great advantages of proxy-based approaches is that, from the content provider perspective, they appear to be a simple extension of IP source network authentication: they do not require any additional work by the content provider and all of the complexity of an authentication system is hidden within the institution's proxy servers. This suggests to me that most institutions will be forced to continue to support a proxy-based approach for the foreseeable future, if they are to manage access to a wide range of publishers. Credential-based approaches may be useful in authenticating users to the proxy, and can also be exposed to those content providers that understand and support them where it is useful, but user needs for consistency of mechanism may well mandate that all access be managed by proxy except for specialized inter-institutional resource sharing arrangements (such as electronic reserves).

Finally, it is worth noting that institutional authentication systems are not driven solely by demands to access network-based information resources; they are needed as a means of managing access to institutional resources and services, and, in a time when many networks within higher education institutions are under constant attack, they are also being viewed by some institutions as a means of improving security. Some institutions are now seeking to establish both policy and technical infrastructure that forbids unauthenticated access to the institutional network and the services available on it. To the extent that these local authentication mechanisms are exposed externally and reduce the privacy of library patrons in interacting with external resources, I believe that they require very careful examination and discussion. At many institutions, the focus of the information technology managers is primarily on the management of local resources, while the focus of the library is on access to resources beyond the boundaries of the institution, and it is essential that these two perspectives be balanced in the development of institutional infrastructure and policy.

CNI's 1998-99 Program Plan

CNI's 1998-99 Program Plan has recently been issued. Specific projects for the year are detailed under the broad themes of:

- Developing Networked Information Content
- Transforming Organizations, Professions, and Individuals
- Building Technology, Standards, and Infrastructure

The full text of the plan is available at: <http://www.cni.org/program/>.

Table of Contents for Issue 201 | Other Networked Information Articles | Access and Delivery Services Articles

ARL Newsletter Home ARL Home

Maintained by: ARL Web Administrator Site Design Consultant: Chris Webster **Last Modified:** December 15, 1998