MANAGING ACCESS TO ELECTRONIC INFORMATION: PROGRESS AND PROSPECTS

Alan Robiette

Based on a paper given at the UKSG 24th Annual Conference, April 2001 at Heriot-Watt University

Managing access to electronic materials, especially where delivery is via a web interface, is surprisingly difficult to effect for large user populations. Traditional approaches, using IP address validation or username/password methods, are far from satisfactory. The Athens system developed for the UK academic community is, perhaps, the first successful demonstration that these problems can be overcome, with consequent gains both for suppliers and for the users and their institutional managers. A new generation of access management projects is just beginning to emerge and the review discusses how these are likely to influence the design of access management regimes in the near future.

Alan Robiette, Director,
Joint Information Systems Committee's (JISC) programme on authentication and security
E-mail: agr@westgate.force9.co.uk

Introduction

Managing access to electronic publications continues to be a headache for suppliers and customers alike. This is particularly true, where the customer is a large educational institution (university or college), where the numbers of individual users entitled under the licensing terms to view the materials may be thousands or even tens of thousands, and where the rate of turnover in the student population is at least 30% per annum and in some cases considerably higher.

The problem is also particularly acute where the access mechanism is via the World Wide Web, as is now almost universal. The web was designed by its creators to make access to information as easy and unfettered as possible. Mechanisms to control and restrict access are few, and many of these are relatively unsophisticated, making the design of a fine-grained security system for the web a difficult task.

The aim of this article is to review briefly the methods in most common use today, and to suggest how various developments, currently under way, may lead to major improvements for all concerned.

Terminology

In much of what follows a distinction is made between *authentication* and *authorisation*.

Authentication may be defined as the process of establishing, with as much certainty as possible (or, perhaps more precisely, with as much certainty as the context requires), the identity of the individual using a given electronic identifier. The authentication step typically involves the individual presenting some sort of credential, such as a password or a smart card and PIN, as evidence that he or she is the person to whom the electronic identifier was issued. Authorisation is a second process. Once the authentication step has established with whom the system is dealing, the authorisation step establishes what functions that person is allowed to carry out. Authorisation typically involves examining various attributes associated with the identifier in question, and from those deciding what actions should be permitted. In a university, for instance, one user might have an attribute "undergraduate student" and a second "academic staff": another attribute, for either, might be "main subject department". If the institution so wished, an authorisation scheme could be drawn up determining which electronic publications any member of the university was allowed to view, based entirely on attributes of this kind.

Such authorisation schemes are often, in the literature, referred to as "rôle-based access control"[1,2].

Note that many older computing environments do not make explicit distinction between authentication and authorisation. This is because the process of authentication (or login) typically placed the user directly into a context in which all their privileges were pre-defined; in other words, authorisation was tacitly amalgamated with authentication. However, most modern thinking treats the two concepts separately.

Traditional approaches to web access control

A seminal white paper, edited by Clifford Lynch[3], describes so well the main methods of access control – and their pros and cons – that only the essentials need to be repeated here. The two principal technologies, both still in widespread use by suppliers, are

- IP address validation, and
- access control via username and password (also described as htaccess)

Where IP address validation is in use, the supplier's web server simply captures the IP address from which the request originates; it then checks this against a table of IP address ranges, which the two parties have agreed to recognise as specifying the customer's site (or sites) and machines. This method is simple and relatively secure, since although IP addresses can be counterfeited this requires some effort on the part of an impostor. It is also efficient in performance terms.

The disadvantages are: that the IP address ranges need to be constantly maintained (not always a trivial task, for large and complex customer sites); that IP addresses only identify machines, not people (a factor that obscures usage statistics in many cases); and that in the simple case, staff and students will only have access to the resource when on university or college premises. Off-campus access can be provided, if the institution runs a web proxy service[4], but to do so is itself a considerable overhead and, in any case, proxying does not work well with all web content.

Username/password authentication involves the web service provider in creating and maintaining a list of usernames for all those entitled to use the resource, a formidable task for a large user population. The individual staff and students have the problem of remembering a plethora of usernames and passwords for the range of resources that they have permission to access. In addition, this method is not at all secure, since the web basic authentication protocol passes the usernames and passwords across the network unencrypted (the alternative, more secure, digest authentication protocol is not widely implemented).

Thus, both of the traditional approaches have real drawbacks, and continue in use principally because better solutions have been slow to arrive. Both also fail to distinguish between authentication and authorisation: the act of authenticating automatically gives access to the resource in question.

The Athens system

Athens[5] is an access management system developed in the UK for the higher education environment, which has now been extended to further education and the National Health Service. The aims of Athens are, firstly, to provide a much reduced administrative load for suppliers, and, secondly, to remove from users the burden of having a separate username and password for each resource. Athens also allows some separation of authentication and authorisation.

These goals are achieved in Athens by holding username/password data for all access-controlled resources in a single centralised database. (The technology used allows this database to be physically replicated, for better performance and resilience against server or network failure, but logically it is a single data structure.) In this is stored the list of Athens usernames and passwords for all user accounts, from all participating institutions: for a given individual the username and password are the same whatever the resource accessed. Each user account has associated with it a list of permission fields, indicating whether or not that user is allowed access to each of the Athens-controlled resources: these permission fields effectively represent an authorisation matrix, connecting users with resources.

An important aspect of the Athens database is its distributed administration. The user records for a given institution are administered by that institution (by one or more local administrators, since the Athens design allows administration to be delegated to sub-administrators, if desired). The local Athens administrator has two tasks: to create and subsequently maintain the required usernames for the institution, and to populate and maintain the permission fields, which determine whether or not a particular user is allowed to access a particular resource. Space precludes describing the system in greater detail, but this is the essence of the design.

From the supplier's point of view, both the authentication and authorisation decisions become the responsibility of the subscribing institution; in other words, the university or college is trusted to manage local usernames and authorisation permissions in accordance with the contract between the two parties. Subject to this implicit agreement, user administration becomes extremely simple for the supplier. if the user presents a username and password which check with entries in the central Athens database, and, if the entry in the Athens authorisation matrix for the supplier's resource also validates correctly, the user is permitted access. Suppliers have to run some special software (known as the "Athens agent") on their web servers, but in most cases this is a small price to pay for the subsequent ease of administration.

Athens has proved to scale remarkably well. At the time of writing, there are well over 1 million user accounts for the further and higher education communities, belonging to hundreds of institutions, and enjoying access to a substantial range of Athens-controlled resources. There is no comparable system currently operating on a national scale known to the author.

What, if any, are the limitations of Athens? For the user it does not, at the time of writing, provide a genuine "single sign-on" service, i.e. access to all resources that the user is entitled to view with just a single entry of the Athens username and password. Although only one username/password pair is involved, this does have to be re-entered each time a new resource is accessed. A single sign-on capability is projected as a future enhancement.

For the institution the drawback is that Athens usernames are distinct from any usernames issued by the campus for its own resources, so that the administrative load is higher than if these two sets of electronic credentials could be integrated. This is one limitation which some of the new systems under development aim to eliminate.

Current developments

A new generation of access management projects is currently emerging. These projects, by and large, all share many of the same design goals. The basic framework is characterised by the following scheme:

- authentication and authorisation are logically separated;
- authentication is the responsibility of the user's institution;
- the authentication process may, optionally, create a token of some form indicating that the user has authenticated at the "home" institution, which will enable other co-operating resources to by-pass the authentication step (single sign-on);
- authorisation is typically rôle- or attribute-based;
- the authorisation process involves the home institution disclosing user attributes to the external resource to allow the access decision to be made;
- attribute disclosure may be selective, in order to protect the user's privacy.

Perhaps the most ambitious of the new projects is Shibboleth [6], which is being developed under the auspices of the Internet2 collaboration in the United States. Shibboleth describes its aims as "to develop architectures, frameworks and practical technologies to support inter-institutional sharing of resources which are subject to access controls." The architecture is intended to serve a significant range of applications, which include the electronic library, but are not restricted to that scenario Other potential uses are the teaching and learning environment (for controlled access to learning materials), and peer to peer collaboration between different institutions, e.g. in a research context. Shibboleth also places special emphasis on user privacy issues.

Shibboleth's internal mechanisms will make use of the evolving standard SAML (Security Attribute Mark-Up Language), which is being developed by an open consortium (OASIS-Open[7] representing a strong consensus of commercial companies operating in the security and e-commerce sectors. The project plans to deliver an open source reference implementation of the Shibboleth architecture, some time in 2002.

So far as it is possible to judge, Shibboleth seems likely to be very influential in setting the direction for future access management systems. Another project to note is PAPI[8], developed within the Spanish academic and research community, which is more specifically focused on access management for electronic materials than Shibboleth and, partly for this reason, has achieved operational status sooner. A particularly flexible authorisation model, designed for controlling access to scientific resources but of more general applicability, is found in the Akenti project[9]; this makes use of digital certificates and related concepts, which may become the prevalent authentication technology of the future.

The UK academic community has its own project code-named Sparta[10], the aim of which is to procure and deploy a successor to Athens, building on the success of Athens and at the same time incorporating the new approaches, exemplified by such projects as PAPI and Shibboleth.

Summary and conclusions

This is a formative period for access management technology. The limitations of the traditional methods – IP address validation, and server-based username/password authentication – are all too clear, yet these technologies persist for lack of well-developed alternatives. The Athens system was a pioneer in showing how some of the problems of the traditional procedures could be overcome, and showing too that its more sophisticated approach could be scaled up to large user populations.

The scene is now set for a new generation of systems, making use of modern developments such as XML and digital signature technologies and in many cases incorporating rôle-based authorisation as the most cost-effective approach to the authorisation problem. Several such projects are at an advanced stage or entering production. It will be of great interest to see how they are received by commercial suppliers. It is an exciting prospect.

References

1. Bacon, J., Hayton, R., and Moody, K., Middleware for Digital Libraries, DLib Magazine, October 1998, http://www.dlib.org/dlib/october98/bacon/10bacon.html
2. Lupu, E., Milosevic, Z., and Sloman, M. Use of Roles and Policies for Specifying and Managing a Virtual Enterprise, 1999, Ninth IEEE International Workshop on Research Issues on Data Engineering: Information Technology for Virtual Enterprises (RIDE-VE'99), http://www.doc.ic.ac.uk/~ecl1/papers/rideve99.pdf
3. Lynch, C. (ed.), *A White Paper on Authentication and Access Management Issues in Cross-Organizational Use of Information Resources*, 1998, http://www.cni.org/projects/authentication/authenticationwp.html
4. See for example Hunt, S., *Remote User Authentication in Libraries: Proxy Servers and Authentication*, http://library.smc.edu/rpa.htm#proxy
5. EduServ, Athens: *The Access Management Solution*, 2001, http://www.athensams.net/
6. Internet2, Shibboleth, 2001, http://www.internet2.edu/middleware/shibboleth/
7. OASIS, *OASIS: Accelerating Electronic Business*, 2001, http://www.oasis-open.org/
8. Lopez, D., and Castro, C., *The PAPI System (Point of Access to Information Providers)*, 2001, http://www.rediris.es/app/papi/index.en.html
9. Thompson, M., Johnston, W., Mudumbai, S., Hoo, G., Jackson, K. and Essiari, A., *Certificate-based Access Control for Widely Distributed Resources*, 1999, Proceedings of the 8th Usenix Security Symposium, http://www-itg.lbl.gov/security/akenti/papers/usenixsec99.html
10. Robiette, A., Sparta: T*he Second-Generation Access Management System for UK Further and Higher Education*, 2000, http://www.jisc.ac.uk/pub00/sparta_disc.html