Authorization and Attribute Certificates for Widely Distributed Access Control[1]

William Johnston, Srilekha Mudumbai, and Mary Thompson
Information and Computing Sciences Division, Lawrence Berkeley National Laboratory
Berkeley, CA, 94720

Abstract

We describe a system whose purpose is to explore the use of certificates for the distributed management of access rights for resources that have multiple, independent, and geographically dispersed stakeholders. The stakeholders assert their use-conditions in authorization certificates and designate those trusted to attest to the corresponding attributes. These use-conditions implicitly define access groups through their requirement for certain attributes. All use-conditions must be satisfied simultaneously, so the actual access group is the intersection of all of the groups. A policy engine collects the use-condition certificates and attribute certificates when a user attempts to access a particular resource. If all of the use-conditions are met, a capability is generated for the resource. The policy engine can provide several different policy models depending on whether any relationship is established among the use-conditions. The system architecture and implementation is described, together with some of the identified strengths, weaknesses, and vulnerabilities.

1. Policy-Based Access Control in Widely Distributed Environments

The general problem that we are addressing is that of providing a mechanism for independent and widely distributed stakeholders to assert their authority over a resource in a flexible and automated fashion. Our immediate motivation is, in an open network, to enable the sharing of valued resources within the scientific community generally, and for distributed collaboratories in the DOE2000 project[1], in particular.

Distributed scientific systems and collaborative environments that are geographically dispersed over wide-area networks to support:

- multi-user instruments at national facilities,
- distributed supercomputers and large-scale storage systems,
- data sharing in restricted collaborations,
- network-based multimedia collaboration channels,

give rise to a range of requirements for distributed control of access. Among other things administration of such resources as, e.g., network quality-of-service will need to be handled by an automated authorization infrastructure so that management of resource availability and enforcement of use-conditions (e.g. allocation), can be done automatically.

In all of these scenarios, the resource (data, instrument, computational and storage capacity, communication channel) has multiple stakeholders (typically the intellectual principals and policy makers), and each stakeholder will impose use-conditions on the resource. All of the use-conditions must be met simultaneously in order to satisfy the requirements for access. This model is common in society, and is illustrated in Figure 1, where a hypothetical research medicine facility (the "beamline") has multiple stakeholders, some of which are unrelated: The Dept. of Energy sets broad policy for use of the National Laboratories; the Laboratory (LBNL) sets site access rules; the administration of the Advanced Light Source (an ultra-high intensity X-ray source) sets the safety rules; the University has management oversight of the

[1]The work described in this paper is supported by the U. S. Dept. of Energy, Office of Energy Research, Office of Computational and Technology Research, Mathematical, Information, and Computational Sciences under contract DE-AC03-76SF00098 with the University of California. This is report no. LBNL-41349. wejohnston@lbl.gov, www-itg.lbl.gov/~johnston

1. "The fusion of computers and electronic communications has the potential to dramatically enhance the output and productivity of U.S. researchers. A major step toward realizing that potential can come from combining the interests of the scientific community at large with those of the computer science and engineering community to create integrated, tool-oriented computing and communication systems to support scientific collaboration. Such systems can be called ***collaboratories***." From "National Collaboratories — Applying Information Technology for Scientific Research," Committee on a National Collaboratory, National Research Council. National Academy Press, Washington, D. C., 1993.

DOE2000 is a U. S. Dept. of Energy initiative is to bring innovation to, and accelerate the development of, communication systems, computational capabilities, and collaboration strategies that current and emerging technology make possible. Emphases here are two-fold — ACTS, in which is focused the development of the essential infrastructure and simulation tools, and the National Collaboratory, in which the ACTS tools, networks, research facility instrumentation, and collaboration tools and strategies coalesce into collaboration environments that will revolutionize major components of scientific inquiry. See http://www.mcs.anl.gov/Projects/doe2000 .

Laboratory and, among other things, must review and agree to all treatment protocols involving humans; and finally the principal investigator (the "owner" of the resource) establishes who participates in the experiments. Each of these stakeholder use-conditions has one or more attribute certifiers, as indicated on the right hand side of Figure 1. The abstraction of this societal model is what we wish to achieve in our computer-based access control system.

Further, it is common that scientific collaborations are diffuse, with the principals and stakeholders being geographically distributed and multi-organizational. Therefore, without reliable, widely deployed, and easily administered access control mechanisms it will not be possible to build the collaborative environments that will enable broad sharing of scientific and intellectual resources within the scientific community. The access control mechanism must accommodate these circumstances by providing;

- distributed management of policy-based access control for resources,
- origin authentication, integrity, confidentiality, etc., of resource related information,
- transparent access for authorized users and strong exclusion of unauthorized users,

in an operating environment where stakeholders, users, and system/resource administrators may never meet face-to-face.

2. Goals

The overall goals for access control mechanisms in distributed environments that we wish to reach in a computing and communication based working environment, are the general principles that have been established in society for policy-based resource access control. Each responsible entity – principals/stakeholders – should be able to make their assertions (as they do now by signing, e.g., a policy statement) without reference to a mediator, and especially without reference to a centralized mediator (e.g. a system administrator) who must act on their behalf. The mechanism must be dynamic and easily used by all concerned – stakeholders and users – while maintaining strong assurances. Policy should be specifiable and enforceable completely within the scope of the stakeholder and user communities. An access policy might involve just

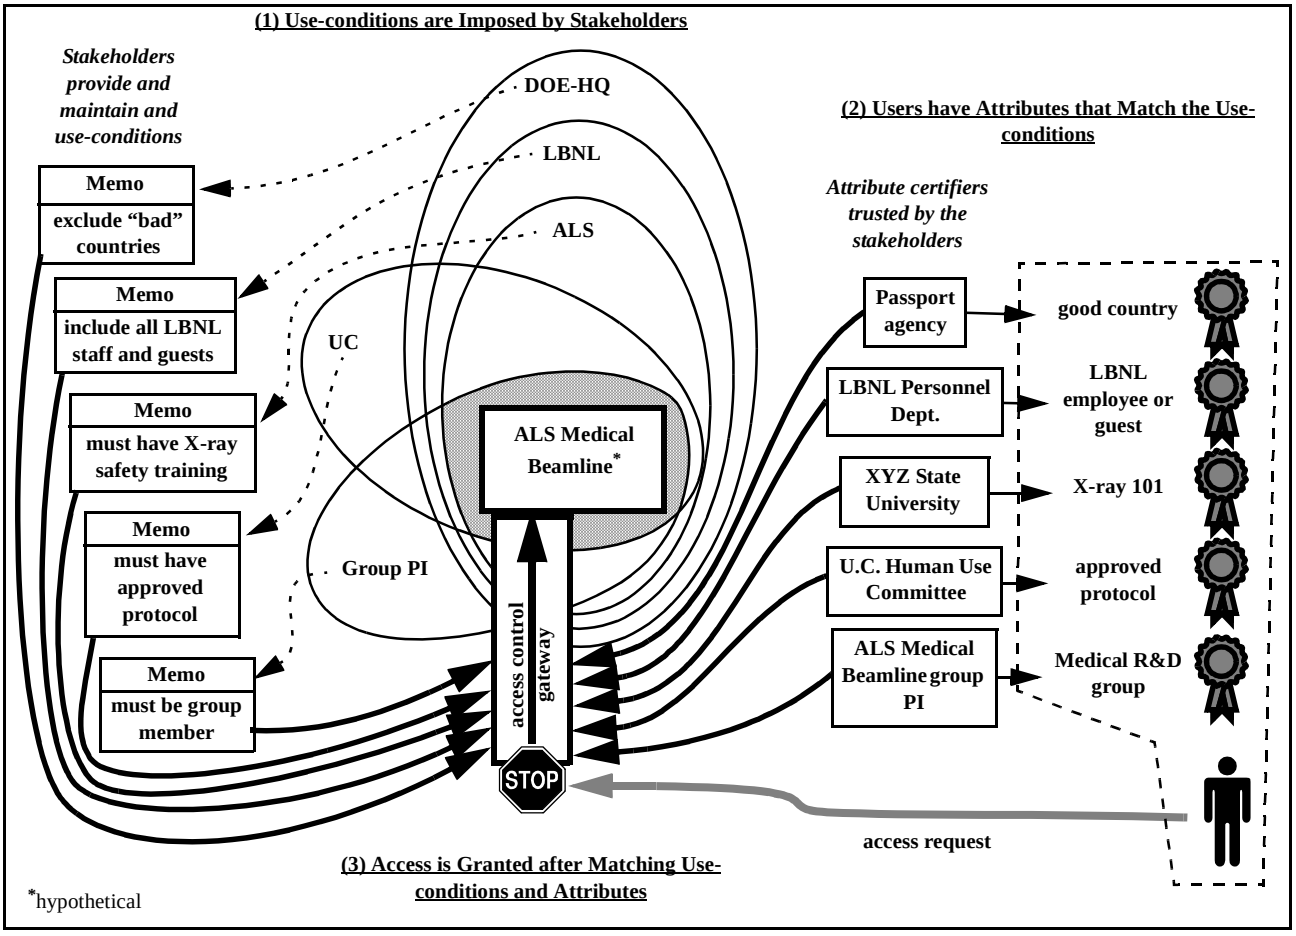

Figure 1 **Societal Access Control Model**

a few people, or thousands. The scope of specification and enforcement should be commensurate with the affected community and resource. Only in this way will computer-based security systems achieve the decentralization and utility needed for the scalability to support large distributed environments.

Specific goals are that the computer based resource access control mechanisms should be able to collect all of the relevant assertions (identity, stakeholder use-conditions, and corresponding attributes), and make an unambiguous access decision without requiring entity-specific or resource-specific, static configuration information that must be centrally administered. (This does not imply that such specific configuration is precluded, only that it should not be required.) Having made the policy-based decision, it should be able to ensure compliance both on the part of the users and unrelated parties. It is also a goal that the mechanism should also be based on, and evolve with, the emerging, commercially supplied, public-key certificate infrastructure components.

3. Expected Benefits

In order for distributed environments dealing with anything of value or importance to be successful, both protection and policy enforcement are essential. Without these characteristics our on-line environments will fade away as repeated (successful) attacks and misuse will cause the confidence of stakeholders and users to decline. On the other hand, successful deployment of large-scale, distributed, on-line authorization systems will benefit a broad spectrum of scientific, commercial, and governmental activities whose outreach and efficiency are dependent on reaching a large and geographically dispersed community.

For security to be successful in distributed environments – providing both protection and policy enforcement – each principal entity should have neither more nor less involvement than it does in the currently established procedure that operates in the absence of computer security. That is, those who have the authority to set access conditions or use-conditions by, e.g., holographically signing statements in a paper environment, will digitally sign functionally equivalent statements in a distributed environment. The use of these credentials will be automatic, and the functions of checking credentials, auditing, etc. will be performed by appropriate entities, as is the societal model. The expected advantages of computer-based systems are in maintaining access control policy with greatly increased independence from temporal and spatial factors (e.g. time zone differences and geographic separation), and in automating redundant tasks such as credential checking and auditing.

A further expected benefit is that this sort of a security infrastructure should provide the basis for automated brokering of resources that precede the construction of dynamically, and just-in-time configured systems to support, e.g., scientific experiments with transient computing, communication, or storage requirements.

4. Approach: Authorization Based Distributed Security

Our approach to the general goals noted above is based on cryptographically (“digitally”) signed documents (“certificates”) that convey identity, authorization, and attributes. A digital signature can assert document validity without physical presence of the signer or physical possession of holographically signed documents. The result is that the digitally signed documents that provide the assertions of the stakeholders, attribute authorities, etc., may be generated, represented, used, and verified independent of time or location.

Assertions are implemented through the use of signing “authorities” that provide assured information as digitally signed documents: Identity authorities connect human entities and systems to digital signatures; stakeholder authorities provide use-conditions; attribute authorities attest to user characteristics, etc. Additional components include reliable mechanisms for generating, distributing, and verifying the digitally signed documents; mechanisms that match use-conditions and attributes; and resource access control mechanisms that use the resulting credentials to enforce policy for the specific resource. All of these mechanisms rely on public-key cryptography for digital signatures and public-key infrastructure for certificate management. (For a general introduction to public-key technologies see [3] or [7].)

5. Architecture and Implementation

The general architecture is illustrated in Figure 2. The elements of the architecture include document signing mechanisms (the “authorities”), certificate distribution servers, a policy engine, and an access control gateway.

The signing authorities generate certificates that assert identity, use-conditions, and user attributes. The certificates are made available through trusted agents of the authorities – Web servers and LDAP directory servers. The user-client presents identity credentials to the resource access control gateway, together with a request to access a particular resource. The access control gateway authenticates the user and then passes the access request to the policy engine. The policy engine identifies all of the stakeholders and then searches their certificate servers for use-conditions related to this resource. The use-condition certificates are validated and the identity of the trusted attribute certifiers extracted. The attribute certifier servers are searched for certificates

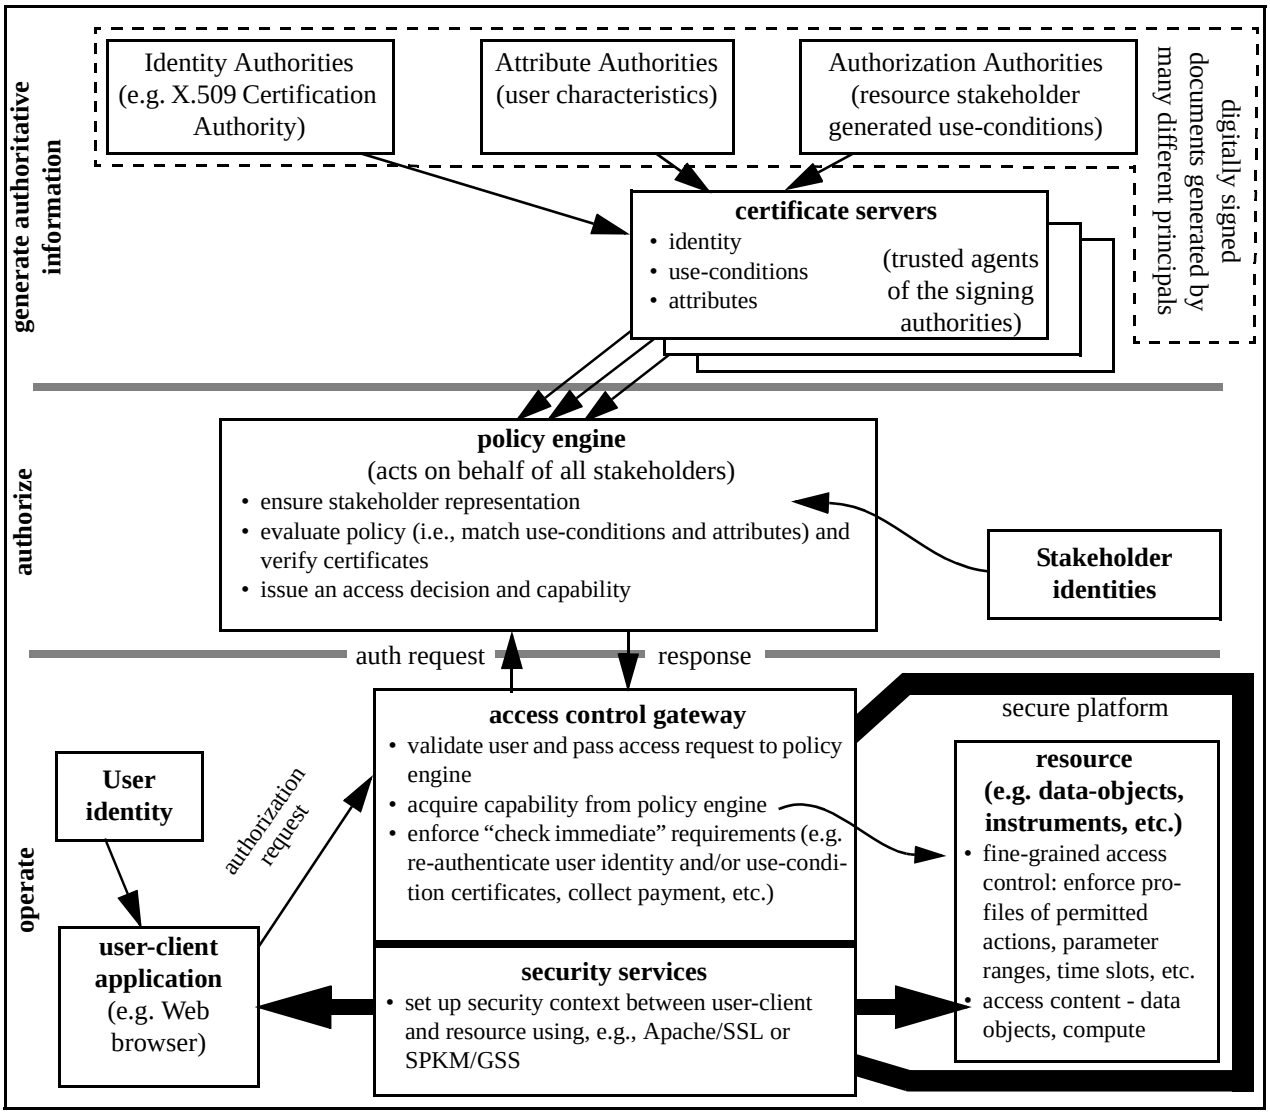

Figure 2 **An authorization and attribute based access control architecture.**

that provide the required attributes for the user requesting access. If all of the use-conditions are satisfied by locating and verifying the corresponding attributes, then a "capability" is returned to the access control gateway (or directly to the resource server). The access control gateway may also invoke various "check-immediate" mechanisms. For example, all certificates related to the capability might be revalidated (potentially important if the policy engine is remote from the access control gateway or if the capability is cached by either the user or the access control gateway).

When this is complete, all of the information is available to establish a security context for the service that provides a secure communication channel between the user-client and the resource. This security service is not part of the credential and policy mechanism, but is a required part of the overall system. The secure channel might be set up between the user-client and the gateway (which acts as a proxy for the resource) or directly with the resource. The user-client now has a valid and secure access to the resource.

The implementation does not provide for establishing *ab initio* policy (i.e. policy directly encoded in the certificates), but rather provides for enforcing policy agreements that are established in the usual societal manner. That is, the system does not do any semantic analysis of use-conditions or attributes, but rather treats them as tokens whose relationships are governed by a simple set of rules. The semantic content of the tokens is represented by out-of-band agreements among the principals.

Identity certificates are required for all parties in order to

validate the various digital signatures. We currently use X.509 identity certificates and these may be provided by any "trusted" certification authority ("CA").

There are several trust relationships involved in our approach: 1) CAs that verify the identity of the principals must be agreed upon since counterfeiting a CA could permit validating a counterfeit use-condition or attribute; 2) there must be a way to ensure the representation of all of the stakeholders since ignoring any of them could weaken the access control requirements and permit invalid access; 3) the stakeholders that establish use-conditions must specify whom they trust to certify that users have the attributes to match their use-conditions. For trust relationship (3), trusted attribute certifiers are directly specified in the use-condition certificates. We are still experimenting with ways to establish the trust relationships (1) and (2). All of the variations that we have considered involve a mutually trusted third-party to maintain the required information once it has been agreed on by the principals. (The agreement establishes who the stakeholders are, not necessarily what their use-conditions are.) Specification of a CA provides the "true" identity of the principals, and this may involve many otherwise unrelated CAs. The CAs are only related by the agreement among the stakeholders that they have adequate identity policies for the purpose at hand. Specifying the stakeholders defines who "owns" the resource and who has authority to set policy for its use. In principle, we could have a trusted third-party operate the policy engine and supply this information to the policy engine (e.g. as yet another type of certificate). In practice, the policy engine and access control gateway read this information from a configuration file that is maintained by a trusted third-party. Although this "smacks" of the central administration that we want to avoid, we don't believe that this type of administration violates the spirit of the original goals because, in practice, this information (definition of the stakeholders, etc.) tends to be static, changing only when fundamental policy or ownership changes, or when new resources are set up.

6. Status

All of the elements illustrated in Figure 2 are operational in a system called "Akenti". At this point the primary prototype is an Apache Web server (with Akenti replacing the standard access control module) that is used as an access control gateway for a variety of Web-based resources. The initial experimental operating environment involves three independent CAs, about six government and commercial organizations scattered around the country, 10s of stakeholders, and something less than a hundred users.

In addition to the technology issues of integrity and management of the access control system and associated computing platforms, useful security is as much (or more) a deployment and user-ergonomics issue. That is, the problem is as much trying to find out how to integrate good security into the end-user (e.g. scientific) environment so that it will be used, trusted to provide the protection that it claims, administered easily, and genuinely useful in the sense of "providing distributed enterprise capabilities" (i.e., providing new functionality that supports distributed organizations and operation), as it is trying to address the more traditional security issues.

The Akenti prototype (see [5]) provides a policy engine that implements both flat and hierarchical multiple-use-condition policy models, uses X.509 identity certificates and ad hoc attribute and use-condition certificates obtained from Web and LDAP servers, and provides a policy evaluation service to the Apache Web server using the Secure Sockets Layer [6] and to an implementation of SPKM/GSS [4].

While the security architecture provides the basic technology, in order to accomplish a useful service the architecture must be applied in such a way that the resources are protected as intended by the principals. This involves understanding the information / resource use and structure model, and developing a policy model that will support the intended access control.

Although we are just starting to evaluate the Akenti system, the strengths of the approach appear to be that it can provide the sort of distributed management of use-conditions from multiple stakeholders that are described in the goals. The most obvious current weakness is that the various user interfaces must evolve to provide very simple mechanisms for the stakeholders and attribute certifiers to deal with only the information that is required to accomplish their task. (Steady progress is being made in this area.) For the user who has the correct credentials the access control is almost transparent. When problems occur it is our intent to provide specific information on what credential is missing, expired, etc., and while designed, this feature is not yet implemented. The two obvious vulnerabilities are that while stakeholders are named, their specific use-conditions are maintained on servers that must be "trusted" by the stakeholders. If those certificate servers are not secure, then use-conditions could be deleted, resulting in weakened access control. There are also many opportunities for denial-of-service attacks, since if a required certificate is unavailable, then a legitimate user will be denied access. As we gain some operational experience, we will better be able to assess the importance of each of these and the difficulty of addressing the associated issues.

7. Acknowledgments

Figure 1 is reprinted with permission from Morgan Kaufmann Publishers, Copyright 1998. Case Larsen wrote many of the certificate handling and cryptographic libraries used in the Akenti implementation. Bob Aiken and Mary

Anne Scott of DOE/ER/MICS have been consistent supporters of this approach to security and access control.

8.0 References

[1] "The Akenti Approach", http://www-itg.lbl.gov/security/Akenti

[2] "About the Apache HTTP Server Project," http://www.apache.org

[3] W. Ford, *Computer Communications Security: Principles, Standards, Protocols, and Techniques*. Prentice-Hall, Englewood Cliffs, New Jersey, 07632, 1995.

[4] J. Linn, "Generic Security Service Application Program Interface", Sep 1993. Available at http://ds.internic.net/rfc/rfc1508.txt. Also see more recent and related drafts at the IETF Common Authentication Technology home page (http://www.ietf.cnri.reston.va.us/html.charters/cat-charter.html) and at http://www.ietf.cnri.reston.va.us/ids.by.wg/cat.html.

[5] S. Mudumbai, W. Johnston, M. Thompson, and A. Essiari, "Akenti -A Distributed Access Control System",draft available at http://www-itg.lbl.gov/security/publications.html

[6] Netscape Corporation, "The SSL Protocol" http://live.netscape.com/newsref/std/SSL.html

[7] B. Schneier, *Applied Cryptography*, Second Edition. John Wiley & Sons, 1996