Credentials for Privacy and Interoperation

Vicki E. Jones Neil Ching Marianne Winslett

Department of Computer Science
University of Illinois at Urbana-Champaign
1304 West Springfield Avenue
Urbana, IL 61801 USA
{vjones,ching,winslett}@cs.uiuc.edu

Abstract

We consider the problem of providing secure, private access to applications and data in a world-wide distributed client-server environment such as the Internet of the future. In such a system, the set of potential users of a service may extend far beyond the local community knowable to the application providing the service. Applications will not generally have prior knowledge of the individual making a request upon which an access control decision can be based and furthermore, knowledge of an individual's identity may not be directly useful. We frame our discussion in the context of supporting credentials *which are submitted with a request, and propose a list of desiderata for such credentials. We evaluate several well-known proposals for credentials, focusing on issues related to privacy and scalability, and then point out the research issues that remain before such schemes can be deployed in a world-wide environment with strong privacy guarantees.*

1 Introduction

Recent years have seen radical increases in both the availability of networked information services and the number of users of these services: the Internet is growing exponentially, approximately doubling in size each year with estimates of users as high as 27.5 million as of October 1994 [11]. In 1994 the worldwide personal computer market grew by 20 percent [7], information services managers at large US companies spent 75 percent more on client/server systems than in 1993 [4], and the worldwide multimedia personal computer market was four times as large as it was in 1993 (an incredible 312 percent growth rate) [6]. Furthermore, it is predicted that home PC shipments will rise at a 21 percent compound annual growth rate (CAGR) and the rest of the US PC market will grow at an 8.9 percent CAGR [5].

This explosion of users (both as information producers and consumers) warrants re-examination of many aspects of large, open, distributed computing systems. We wish to encourage thought on the problem of providing *secure access* to information and resources on a wide scale. We anticipate that the development of efficient, flexible solutions to this problem — already a challenging issue for the entire range of software systems deployed in real-world enterprises — will be complicated by the increasing need and demand for interoperability between such diverse systems. New techniques for supporting the interoperation of heterogeneous environments are the focus of much current research, and we highlight potential extensions to existing security mechanisms in this emerging software technology.

Software systems such as database systems or distributed file systems, and higher-level user applications built on top of them, often rely upon the access control mechanisms provided by the underlying operating system, or implement limited mechanisms of their own patterned after those developed for operating systems. While the security features found in many modern operating systems (both research and in production) are quite sophisticated, we believe that they will not scale well for use in very large, open computing environments such as the World-Wide Web or a CORBA[1] environment.

Our belief is based upon the following observation: traditional protection schemes are based on the assumption that there is a known community of subjects

[1]The Common Object Request Broker Architecture (CORBA) specifies an architecture for interoperation between heterogeneous systems in a distributed environment. See [9] for more details.

0-8186-7318-4/95 $4.00

entitled to access the data managed by a system. This community is represented within the system by a set of identities, and enforcement of access control policies is based upon knowledge of these identities and of information keyed to these identities. However, as techniques for providing wide electronic access to information and services are developed and deployed in the move to a global computing environment, the set of potential users of an application may extend far beyond the community of users local to the system where the application resides. We contend that the traditional identity-based approach will not suffice in this situation for several reasons.

First, maintaining databases of identities will cause applications to suffer in efficiency. It would clearly be infeasible for an application to have prior knowledge about every individual who may submit a request for service, as this would place an unacceptable burden on the application server. More realistically, this information might be partitioned among multiple authorities (e.g., as suggested in existing authorization frameworks such as X.509 [18]), but then processing a request for service may require several extra messages in order to consult the relevant databases to determine the appropriate access. This will often result in unacceptable response times for an interactive session and significant extra network load.

Also, it would be cumbersome for users to ensure that they have registered themselves with each application (or with the authorities consulted by the application) prior to their first use of it. It is to the service provider's advantage to avoid imposing such a burden on users if possible, and may be essential in circumstances where the price per transaction is very low and optimal levels of both performance and ease of access are crucial to obtaining a profitable volume of traffic, such as is true for many applications on the World-Wide Web.

Although pre-registration as a condition of access will be appropriate for some situations, a shift away from *a priori* knowledge of identities as a basis for authorization and towards flexible, anonymous or semi-anonymous access mechanisms would help to avoid the difficulties just mentioned. In addition, the systematic wide-scale use of identity-based access control is likely to have other pragmatic drawbacks heretofore not severely felt in systems with a centralized administration: in particular, the global adoption of a uniform handling of identities would infringe on the *autonomy* of enterprises and the *privacy* of users.

Furthermore, identity information in and of itself is often not germane to the access control decision, but is used to determine that a subject belongs to a group or role authorized to obtain the requested access. In such situations it is preferable to avoid the nuisance of establishing and maintaining registries of identities in favor of mechanisms which would allow the user to supply information about himself or herself with anonymity.

Individuals often cross organizational and system boundaries in conducting their business, and organizations often find it useful to cooperate and share resources in a circumscribed manner. An enterprise may thus need to recognize individuals belonging to foreign realms, of whose organization it is unaware and over which it has little influence. Conversely, an enterprise may wish to export information pertaining to itself while retaining the freedom to manage this data in what it deems the most suitable manner. We believe that a traditional approach to authorization, in which a central authority (or hierarchy of authorities) ensures secure access based upon a system-wide understanding of identities, will not meet the demands of such autonomous interoperation without modification.

The issue of privacy for individuals complements that of autonomy for organizations. Personal privacy seems particularly vulnerable to compromise in the haste to reach the Information Age, perhaps because an architecture based upon a universal representation of identity is the most direct application of current technology. This is not an unreasonable approach. In many countries, individuals are required to possess a unique ID, which is used in claiming employment, medical and financial services, and so on. In the US it has been argued that a national identification system would aid in combating welfare fraud, credit card fraud, illegal immigration, tax evasion, and drug trafficking.

However, we believe that given the premium placed on privacy in the U.S., such measures will not be palatable to a large portion of the population, either in the "real world" or in cyberspace. Global identities, or identities which may be linked across system boundaries, make it possible and easy to track individuals and compile personal information — a prospect objectionable to a citizenry and marketplace accustomed to guarantees provided by such laws as the Privacy Act of 1974 and the Fair Credit Reporting Act of 1970.[2] Many have begun to question if se-

[2] One need not be an alarmist to seriously consider such possibilities. For example, providers of Internet services are under considerable internal pressure to collect information about every on-line action of their customers and sell this information to marketing organizations. The general public doesn't realize the

cure, privacy-preserving interoperation between systems will be achievable.

Pre-registration of an individual's identity as a condition of access will sometimes be appropriate depending on the service and service provider. However, adopting such an access control strategy universally may add substantial processing overhead and diminish ease of access to services intended for wide-scale use and increase the risk to privacy. We believe that a more flexible mechanism is needed to allow users to demonstrate that they are entitled to access certain data or services, without requiring pre-registration. The scheme should be based upon the notion of a user presenting credentials in support of a request. A credential contains information about the requester and originates with a well-known and trusted authority who is in a position to vouch for its accuracy. Although the credentials supplied with a request may contain enough information to allow the requester to be identified precisely, they need not do so, and could simply convey knowledge of some of the requester's attributes.

This data is used by the recipient to determine how the requester fits into the application's authorization scheme and so supplies the basis for the access control decision pertaining to the request.

The credentials supplied with a request generally will not all refer to the requester by a common name or global identifier. Just as paper credentials for an individual (e.g. medical history, voter registration card, or student transcript) are issued by independent agencies and organizations, each of which concerns itself with a particular aspect of the the individual, electronic credentials originate from separate systems, each of which maintains its own registry of subjects. An individual could thus possess several identities corresponding to distinct roles in distinct systems, and each credential which the individual may obtain will refer to him or her in terms of an identity known to the issuing system.

The chief advantages of a credential-based authorization scheme derive from its ability to support the large-scale use of such "multiple identities." From the user's point of view, this would mirror the way activities are conducted in real life. It also ensures a certain amount of privacy, since service providers and credential issuers would need to collude fairly extensively in order to deduce that different identities are facets of a single individual. Furthermore, it allows service providers and other systems autonomy in structuring and maintaining their local databases of information on individuals, since there is no requirement to use wide-scale naming schemes for individuals (although it may often be useful to refer to a individual by the value of some well-known, standard key attribute, such as a social security number or X.500 distinguished name).

In the remainder of the paper, we discuss the potential desirable properties of a credential-based authorization scheme and existing mechanisms to meet the needs described. While this discussion includes many aspects of a secure credential system, our primary emphasis is on protecting privacy. Section 2 lists desiderata for credential schemes, Section 3 discusses today's approaches for providing supporting credentials in light of the desiderata, and Section 4 concludes the paper with open questions for future research.

2 Desirable Properties for *Wide-Use* Credentials

A credential-based authorization scheme (a "credential system") relies on credentials to make decisions about authorized access to protected data and services. A credential represents a statement made by an authority having knowledge of some real-world enterprise. A variety of such statements may need to accompany a request for service in order to make meaningful decisions. A superset of desirable properties for a credential system is enumerated below. We use the terms "individual" and "user" interchangeably to refer to a person or organization as well as any agent acting on behalf of a person or organization.

The properties mentioned below are not all needed within a single system — some are even mutually exclusive. Those desirable for a particular system will depend on the type of service and level of protection desired.

- Interoperable: A service provider's choice of access authorization model should not preclude the ability to interpret and use credentials issued by other authorities.
- Expressive: Credentials should be able to indicate useful information about individuals other than their identity.

  For example, to obtain access to the on-line proceedings of the New Security Paradigms workshop (stored on the computer of the conference chair), an individual may need to provide a credential indicating that they are a member of SIGSAC and

extent of information maintained in marketing databases such as the Lotus CD. When the Lotus CD project was publicized, it was so upsetting that a write-in campaign was undertaken, forcing the cancellation of the project.

another credential indicating that all SIGSAC members are ACM members. The latter credential does not specify an attribute of the user presenting the credential, but is just as important as a direct user attribute for access authorization decisions.

- Extensible: The credential system should allow new individuals and organizations to be added easily and the format should allow the expression of new types of information.
- Spontaneous: Users should not be required to pre-register with a provider in order to receive services. For example, pre-registration would be an excessive burden for an ACM member who wanted to browse a paper abstract from the New Security Paradigms workshop.
- Anonymous: It should be possible for an individual to receive services without revealing a unique, universal identifier of the individual (such as a fingerprint, retinal scan, social security number, etc.).

  As noted earlier, many services will not depend on the identity of the user for authorization decisions, relying instead on the user's ability to demonstrate possession of some other characteristic, such as ACM membership. Digital cash is an example credential which is at least in theory completely independent of identity.[3] A user need only demonstrate an ability to pay in order to receive many services. A digital storefront's request for identifying information would often not be well received.
- Scalable: Credential systems should be robust as the number of users, service providers, and issuing authorities increases. For example, the organizers of the New Security Paradigms workshop should not be forced to store the identities of all ACM members on their workstation.
- (Not) subject to collusion: It should (not) be possible for organizations to collude and determine more about an individual than obtained directly from credentials presented by the individual.
- (Not) transferable: An individual should (not) be able to give a credential to someone else and have them use it as their own. A credential specifying membership in ACM would be intended to be non-transferable, while digital cash, under some schemes, is intended to be transferable.
- Inexpensive: The cost of obtaining and using credentials should be reasonable.
- Verifiable: It should be possible to determine that the issuer of a credential is indeed the source indicated in the credential, and that the credential has not been changed since it was issued.
- Unforgeable: It should not be possible to produce a credential purporting to come from another source, such that the credential may then be verified by that source.

The last two items above, verifiability and unforgeability, should set minimum standards for any credential system. A third desirable property, confidentiality, was omitted from the list because it is difficult to precisely specify what constitutes a reasonable attempt to preserve confidentiality. Beyond these three, the desiderata of most import to protecting privacy are anonymity and non-collusion. Individuals' confidence in their own control over the information about themselves will determine their willingness to provide the information to others. If they are assured that the information about themselves they choose to divulge is not subject to wilful and unauthorized distribution or correlation, they will be more willing to release sensitive information to those who may need it.

In the next section we discuss some existing systems and the tradeoffs they have made with respect to the above desiderata.

3 Candidate Credential Systems

A variety of existing mechanisms provide some of the desiderata specified above. Due to space limitations we are unable to analyze all of them, but this section reviews some of the better known systems in light of the desiderata, with an eye to extensions to their originally intended use.

3.1 Capabilities

Since 1962 when the concept of a capability was first introduced [8] and 1966 when the term was coined [10], the world has seen a variety of capability-based systems with slightly different properties. In most capability-based systems today objects are represented by a physical name or address known as a capability. This capability is effectively a ticket whose

[3] In some systems [1], digital cash is anonymous until it is used improperly, at which point it is possible to determine the illegal user.

possession authorizes the holder to access the specified object in a specific way. Capabilities are protected objects—they are not allowed to migrate into any address space directly accessible by a user process. They are unforgeable and, in most systems, can be transferred to other users. Because capabilities are relatively inexpensive to create and use, many systems have a hybrid implementation—access control lists are used to obtain a capability to an object which is used for subsequent access to the object. Therefore, whether or not the capability-based system is spontaneous, anonymous, extensible, or facilitates collusion depends on the requirements of the authentication system which issues the capabilities. Scalability of capability-based systems is limited by the centralized distribution of capabilities in these hybrid systems.

Since capabilities only specify the allowed access for an individual to an object, they can only express a limited amount of information. The language (and meaning) of capabilities could be extended to allow expression of, say, ACM membership. Using capabilities to facilitate interoperation between autonomous capability-based systems, however, would require capabilities to be available outside the system processes, thus subverting their inherent security or requiring a completely different implementation approach. Implementation of a capability system in a network environment would require facilities for the secure transfer of capabilities between systems and would allow access control decisions to be made outside of the domain of the target object, thus requiring a more general means of referring to an object than its local address.

3.2 Taos

The security mechanisms in the Taos operating system [17, 15] provide facilities for sending and receiving messages over secure channels and for identifying the source of a request in support of an access control decision. In typical fashion, this decision is made by consulting an access control list once the requester has been identified. Among the notable features of Taos, however, is a sophisticated treatment of identities at the the operating system level. Several varieties of *principals* are recognized, which are the entities that can be authenticated as the source of a request. Authentication of a principal is accomplished through cryptographic means. The association between the name of a principal and a key which is used to identify and secure that principal's activities is managed by the Taos security service, which may provide a certificate of identity attesting to the binding. Since the processing of a request often involves interaction of the requester with other principals (e.g., in a delegator–delegate arrangement) and elements of the Taos system (e.g., processes and communication channels), an authentication may require evidence of proper transfer of authority between such entities, also provided in the form of credentials that are obtained from the security service.

In a simple configuration, users in a distributed Taos system reside at particular sites, and the security service at each site maintains a local database of principals and groups which it consults in order to issue certificates of identity and group membership. Since creation, provision and use of credentials is controlled by Taos and these functions are made available through a well-defined API, forgery or unauthorized disclosure at the application level is unlikely. A degree of anonymity is possible as several named principals might correspond to a single real-world user (one could use Taos identities as pseudonyms), but requests must be presented under *some* name so it may yet be feasible to discern patterns of access for particular individuals, by examining authentication logs, say. Many widely-dispersed users could be handled by imposing a hierarchical organization among Taos certification authorities; however, there would remain the assumption that each site uses Taos, constraining the possibilities for interoperation. Also, Taos realizes a principled design for access control which takes into account various properties of the request and of the requester (e.g., the physical origin of the request, or attributes of the requester such as group membership or the adoption of privilege-modifying roles), thus offering an increase in expressivity over simpler schemes [15], and use of delegation certificates provides a controlled transfer of authority. A fixed set of credential types is used to encode such data; a similar design which is extensible with respect to the statements contained in credentials may be possible.

3.3 Kerberos

Kerberos [16] is an authentication system providing evidence of a principal's identity. In Kerberos tickets are used to securely pass the identity of an individual between the authentication server and a specific end server. They are encrypted to ensure they are verifiable, unforgeable, and confidential. These tickets contain information from a Kerberos database of data about registered individuals. Pre-registration with the Kerberos authentication server which maintains the database is required but registration with individual system services is not. As with capabilities, the user identification requirements of the authentica-

tion server determine the spontaneity and anonymity properties of a particular Kerberos system. To receive a service a ticket and an authenticator must be presented with the request. The authenticator allows verification that the individual presenting the ticket is the individual to whom it was issued—thus rendering tickets non-transferable. Unlike tickets, an authenticator can only be used once; a new one must be generated for each request for service. Since authenticators are created by the client, and multiple authentication servers can exist to create tickets, scalability is not an issue within a particular Kerberos system.

Tickets are valid only within the realm of the issuing ticket-granting server for a limited period of time and cannot be used in other Kerberos systems. Tickets can be used to authenticate a principal to another Kerberos system, but the principal must obtain tickets from the local system for that system's services. In version 5 of Kerberos, tickets can be forwarded but the identity of the original source of the ticket is not forwarded so the local service must decide independently of the source whether or not to accept the ticket. Kerberos tickets can be generated by multiple services but are not useful outside the realm of issue; thus they offer the potential for interoperation, but not a solution. The potential for extensions to increase expressivity and extensibility also exists even though the current design is limited to identity.

3.4 X.509 Authentication Framework

The certificates specified in CCITT X.509 [18] are a kind of identity credential, one which binds a name to a cryptographic key (a public key). X.509 defines a structure with fields specifying the name, the associated key, the issuer of the certificate, and auxilliary information which includes the parameters used by the issuer to generate its digital signature for the certificate. Knowledge of these parameters and the public key of the certificate issuer allows a recipient to verify the affixed signature; getting the public key of the issuer may entail obtaining and verifying a chain of one or more additional certificates until a public key is found whose trustworthiness has been independently established.

The specification defines a concrete syntax for certificates intended to serve as identity credentials and suggests a broad framework for authentication based upon their use. It makes few prescriptions on the significance attached to the receipt of a credential from a certification authority. What follows from the knowledge that authority A vouches that public key k corresponds to the individual with distinguished name DN depends upon the policies observed by the authority in issuing certificates and the nature of the trust placed in the authority by the community and enterprises which it serves — issues on which the specification is largely silent. The possibilities range from authorities who issue "high-assurance" certificates in which can be placed a degree of confidence comparable to that for physical credentials, to authorities who issue semi-anonymous "persona" certificates which establish identities not associated with a permanent holder and primarily suitable for correlating messages within a transaction or session. ([12] contains an extended discussion.)

Guarantees of anonymity and the possibility of collusion to track individuals by the certificates they present are determined by what information authorities require from individuals before issuing certificates and the policies that govern their handling of this information; cost of use and scalability rely upon the details of certificate processing within a particular framework and system implementation. Although X.509 certificates have a rigid format tailored to use for authenticating individuals as described above, there is nothing preventing (ab)use of the format to provide information other than identity in a certificate,[4] so there is the potential for expressivity with the use of such credentials.

3.5 Chaum's Credentials

Chaum and colleagues [3, 2] describe a credential mechanism which allows individuals to control the transfer of information about themselves between organizations. An individual is known to an organization by a digital pseudonym which not only identifies the individual to the organization, but provides the medium for issuing credentials. Unlike other proposals, the individual can then transfer the credential to any other pseudonym by which he or she is known and use the credential elsewhere. The issuing of pseudonyms is tightly controlled by a special authority (or hierarchy of authorities) which assigns a unique identifier to the individual. As in other proposals, anonymity and spontaneity depend on the identification and pre-registration requirements of this authority. An individual can increase anonymity within the system by changing his or her pseudonym with a

[4] Specifically, a distinguished name is formally a set of attribute–value pairs; the intended interpretation is that they specify a path to an individual in a directory information tree, but other interpretations are possible. See also PKCS#7 [14], a specification similar to X.509, and extensions to X.509 in PKCS #6 [13].

particular organization in a fashion that prevents association between the old and new pseudonyms.

Even though transfer of credentials between pseudonyms is possible in order to facilitate sharing of credentials between organizations, transfer of credentials between individuals is prevented. Like most, Chaum's credentials are unforgeable and verifiable. However, Chaum's credential mechanism is the only proposal discussed here that can prevent collusion. Chaum's language is expressive since a credential (a number) can mean anything but the system has limited extensibility: once the set of credentials, individuals, and organizations is decided upon it is very difficult to change. Another drawback is expense: these credentials are relatively expensive due to multiple encryption and exponentiation operations required to provide the other nice properties.

4 Research Agenda

With the work described above available to draw upon, where do we stand with respect to our goal of a credential mechanism suitable for wide use and offering reasonable guarantees of privacy? We find many open questions with respect to credential management, credential transfer, credential cost, and support for multiple identities of an individual.

In the systems surveyed, the degree of privacy obtainable is not entirely implicit in the system, but is also determined by how the system is used. For example, consider the question of what information must be presented in order to obtain credentials. If credential issuers require that clients provide unique, universal identifiers for themselves before credentials will be issued, then many opportunities arise for collusion between issuers. Similarly, the existence of credential issuers that do not require such identification of their clients does not in itself guarantee privacy, since any service may decide that it will accept only credentials that contain a universal identifier. (Since the set of credential issuers will probably be much smaller than the set of services that accept credentials, the opportunities for collusion will be different at the two different levels.) Also, services that do not require presentation of universal identifiers might still extract enough information from their clients (e.g., home telephone numbers) to make invasion of privacy quite possible. We cannot hope to control a provider's terms of service but we can refine our original concern about privacy, asking instead what technical barriers remain that might prevent services from accepting credentials that do not contain universal identification.

Alone among the schemes described, Chaum's approach provides strong guarantees of anonymity and non-collusion by allowing the free use of an unlimited set of pseudonyms, and also guarantees of untraceability. Such powerful guarantees will sometimes be appropriate, but more often, these guarantees will prove too strong for the comfort of many services. For example, a merchant dealing in controlled goods (e.g. firearms, pharmaceuticals) must comply with federal laws prohibiting sale to certain kinds of individuals (e.g. convicted felons).

The next step down from Chaum's approach, with respect to anonymity and non-collusion, is an environment in which the individual has a variety of different security identities for use with different protection domains, but cannot freely create and use new pseudonyms. This guarantees services a certain level of knowledge about their clients (whatever information was required to obtain the particular identity currently being used), but makes it somewhat difficult for service providers to collude and infer more information about their clients than they each already know. To some degree, this mirrors the current situation in the real world. However, in a world-wide environment, individuals would find it very useful to be able to transfer the information from a credential issued to one of their identities, to another of their identities, as Chaum's scheme offers. For example, an individual may wish to demonstrate a positive bank account balance to potential creditors without revealing the account number. Unfortunately, other than having the transfer done by a trusted mediator, no techniques are known for accomplishing the transfer; Chaum's approach to transfer only works when the language used to describe the properties of individuals is fixed in advance, which violates our need for scalability, interoperability, and autonomy. Thus one potential avenue for research is scalable mechanisms that allow transfer of credentials between identities administered by different protection domains.

Assuming that we lack the ability to transfer credentials between identities of the same individual, the next step down in terms of privacy guarantees will be a facility that allows an individual to obtain a service by presenting a set of credentials, some of which apply to one of that individual's identities, and some of which apply to others. For example, a student applying for employment may need to share his or her transcripts from several institutions with the potential employer. To do so, the student must convince the employer that each transcript indeed refers to him or her. The technical difficulty with this approach is that in

many cases it will be necessary for the individual to prove to the service provider that he or she is indeed the possessor of several of the mentioned identities. In other words, the individual will need to authenticate simultaneously to multiple identities. No authentication protocols currently in existence have this property. At first, one might think that to validate such a request, it suffices to send it through several iterations of the current authentication protocols, once for each identity in the request. However, not all authentication protocols can be composed in this manner (e.g., one cannot wrap an interactive retina scan around a message), and such composition may not provide the necessary guarantee of validity, unless it can be verified that each iteration or step is performed by the same subject.

The discussion about transfer of credentials has so far considered only transfer of credentials between identities of an individual. Transfer of privilege between individuals, or delegation, has been investigated in the context of Taos, but not under the other systems surveyed here. The utility of modelling real-world delegation such as assigning power of attorney increases as more transactions are conducted electronically. We would like to see additional research on mechanisms for delegation that allow interoperation between systems operating under widely different authorization models.

If individuals obtain a wide variety of credentials, and access a variety of services with differing credential requirements, then credential management becomes an issue. Individuals will probably require automated assistance to determine what credentials to present with a request for service while still being assured that their credentials are distributed appropriately. Research is also needed on the question of how best to explain to a user or his or her agent exactly what credentials are needed for a particular service request. Such explanations should be short, universally intelligible, and preferably sufficiently general that they apply to more than just the current request for service.

The credential mechanisms surveyed in this paper fell short on expressiveness, with two exceptions. Chaum's approach is expressive, but unfortunately at the cost of a fixed language, which is not acceptable in a system for world-wide use. X.509 is also expressive, in the sense that some of its fields can be used to hold any string. The drawback of X.509 in this regard is that certain information is required by the X.509 format (e.g., public key) that may not be relevant for a particular credential. We believe that a modified version of X.509, incorporating a more flexible format and also perhaps including type information for the fields that may hold arbitrary strings, would be a good basis for a more general-purpose wide-use credential mechanism. We also believe that standard, widely-understood languages and lexicons should be established to facilitate parsing and interpretation of incoming credentials by service providers and mediators.

Credentials intended for world-wide use must have reasonable cost. What cost may be considered reasonable depends to a large degree on the service being obtained. For example, a reasonable price for obtaining authorization to read an abstract of the New Security Paradigms workshop will be different from the reasonable price for applying for admission to college, i.e., proving that one is the person to whom certain SAT scores apply and who has a certain high school transcript. In addition, some credentials may be expensive to obtain, but cheap to use, and vice versa. Furthermore, the cost of verifying a credential is often not intrinsic in the credentialing system: it may require traversal of a hierarchy of domain authorities, for example. The traditional assumption that most traffic in a hierarchy will be local traffic will not be a reasonable assumption for some of the new world-wide applications, expecially the World-Wide Web. As different credential mechanisms present different cost/assurances–provided tradeoffs, we anticipate that there will be niches for each of them, from Chaum's approach at the high end, down to capability-like mechanisms. Many open implementation problems remain before we will know how to set up relationships between credential authorities so that credential verification is efficient and explanations of required credentials can be simple.

References

[1] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In S. Goldwasser, editor, *CRYPTO '88*, pages 319–327. Springer-Verlag, 1988. Lecture Notes in Computer Science No. 403.

[2] David Chaum. Showing credentials without identification: Transferring sign atures between unconditionally unlinkable pseudonyms. In *Auscrypt '90*, pages 246–264, Berlin, 1990. Springer-Verlag.

[3] David Chaum and Jan-Hendrik Evertse. A secure and privacy-protecting protocol for trans-

mitting personal information between organizations. In *Advances in Cryptology—CRYPTO '86 Proceedings*, pages 118–167, Berlin, 1986. Springer-Verlag.

[4] Dataquest. *Corporate Client/Server Purchases Up 75 Percent.* Press release December 20, 1994. http://www.dataquest.com/press/pr9423.html.

[5] Dataquest. *Dataquest Predicts Home PC Boom.* Press release July 19, 1994. http://www.dataquest.com/press/previous/pr9402.html.

[6] Dataquest. *Worldwide Multimedia PC Market Quadrupled in 1994.* Press release March 13, 1995. http://www.dataquest.com/press/pr9507.html.

[7] Dataquest. *Worldwide PC Market Grew by 20 Percent in 1994.* Press release January 24, 1995. http://www.dataquest.com/press/ pr9503.html.

[8] J. B. Dennis and E. C. Van Horn. Programming semantics for multiprogrammed computations. *Communications of the ACM*, 9(3):143–155, March 1966.

[9] Object Management Group. *Common Object Request Broker Architecture and Specification*, 1991.

[10] J. K. Iliffe and J. G. Jodeit. A dynamic storage allocation system. *Computer Journal*, 5(3):58–78, October 1962.

[11] Texas Internet Consulting/Matrix Information and Directory Services. *New Data on the Size of the Internet and the Matrix*, January 1995. MIDS Press Release. ftp://tic.com/matrix/news/v5/press.501.

[12] S. Kent. Internet privacy-enhanced mail. *Communications of the ACM*, 36(8):48–60, August 1993.

[13] RSA Laboratories. *PKCS #6: Extended-Certificate Syntax Standard*, version 1.5 edition, November 1993.

[14] RSA Laboratories. *PKCS #7: Cryptographic Message Syntax Standard*, version 1.5 edition, November 1993.

[15] Butler Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, November 1992.

[16] Jennifer G. Steiner, Clifford Neuman, and Jeffrey I. Schiller. Kerberos: An authentication service for open network systems. In *USENIX Conference Proceedings*, pages 191–202, Dallas, TX, Winter 1988. USENIX.

[17] Edward Wobber, Martin Abadi, Michael Burrows, and Butler Lampson. Authentication in the Taos operating system. *ACM Transactions on Computer Systems*, 12(1):3–32, February 1994.

[18] CCITT Recommendation X.509. *The Directory - Authentication framework*, 1988.