

Authentication in the Taos Operating System

Edward Wobber, Martín Abadi, Michael Burrows, and Butler Lampson
Systems Research Center, Digital Equipment Corporation

Abstract

We describe a design and implementation of security for a distributed system. In our system, applications access security services through a narrow interface. This interface provides a notion of identity that includes simple principals, groups, roles, and delegations. A new operating system component manages principals, credentials, and secure channels. It checks credentials according to the formal rules of a logic of authentication. Our implementation is efficient enough to support a substantial user community.

1 Introduction

In this paper we present:

- a design for a general form of distributed system security, including both the external interface and the major internal interfaces;
- a careful explanation of how, in the implementation, an authentication corresponds to a proof in a formal theory [2, 9];
- a demonstration that a clean and sound design can have an efficient implementation.

We use the access control model [10] of security. In this model there are objects (files, printers, etc.), requests, and principals (users, machines, etc.) that utter requests. Each object has a guard or reference monitor that examines each request and decides whether or not to grant it. The request must first be authenticated to identify the principal that uttered

it, and then authorized only if the principal has the right to do the requested operation on the object. We do not address either denial of service or the kind of non-disclosure security policies that are based on an information flow model.

The notion of compound principals [6] makes the access control model more powerful by providing a uniform vocabulary to denote all the principals in a distributed system, including users, machines, programs, delegations, roles, and groups.

We assume that a distributed system is made up of *nodes* connected by an insecure network. A node is a shared-memory computer running an operating system trusted for security local to the node.

In our system there is a component in each node called the *authentication agent* which is responsible for managing principals, credentials, and secure channels. It implements all credential exchanges and validations, communicating with agents in other nodes when necessary. It uses a certification database for names, group memberships, and executable images.

The agent provides security services to application programs. From the underlying system it needs only a bidirectional secure channel to each application and global names for the channels between the application and the outside world.

Our design has several advantages:

- Credentials are managed by the authentication agent. Clients deal with authentication entirely in terms of principals.
- Compound principals denote the sources of requests precisely and uniformly.
- Our system is based on a logic in which the authentication of a request corresponds to a proof. The logic guided our design and aids in understanding the implementation.

The main disadvantage is lack of compatibility with other security mechanisms such as Kerberos [8] or OSF DCE Security [12].

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

SIGOPS '93/12/93/N.C., USA

© 1993 ACM 0-89791-632-8/93/0012...\$1.50

Many systems that offer distributed security do so entirely at the level of the application, either to avoid changing the kernel or because most operating systems do not support a coherent model of user identity throughout the network. Our basic design can be implemented in the same way, with the authentication agent linked as a library in each application.

In fact, however, our distributed security is part of the operating system. This has one major advantage: the notion of identity or principal is built in at a very low level and is represented consistently everywhere. There is no distinction between local and remote principals. Minor advantages are that it's easy to provide the necessary secure channels between the authentication agent and applications, and easy for a child process to inherit the authority of its parent. The trusted computing base doesn't get any bigger, because the operating system must be trusted anyway.

The setting for this work is Taos, the operating system for the Firefly shared-memory multiprocessor [17]. It is completely multi-threaded, yet also implements a protected address-space model close enough to that of Unix that it can run most Unix binaries. Remote procedure call is the primary means of interprocess communication. Although Taos has been a convenient test vehicle, our only real dependency on it is that we could adapt it to our needs.

The next section reviews the logic. Section 3 discusses the application programming interface (API) to Taos security. Section 4, the heart of the paper, describes the implementation in detail. Finally, Section 5 discusses system performance.

2 Background

We use shared key encryption to secure short-term node-to-node channels. All other encryption is public key and is done only for integrity, not for secrecy. We write K and K^{-1} for the public and secret keys of a key pair. We say that a message encrypted with K^{-1} is signed by K so that we need to mention only the public key.

Authentication often relies on the use of signed statements called certificates. These form the building blocks of credentials, which are proofs of authenticity. We view certificates and credentials both as logical formulas and, in the implementation, as data.

Time does not appear explicitly in the logic; formally, assumptions and proofs concern only a given, implicit instant. In our system, on the other hand, a time interval qualifies each certificate. A certificate is valid only for the specified interval. Therefore, the conclusion of a proof is valid only for the intersection

of the intervals of all the certificates used in the proof. Since these certificates typically originate at different nodes, it is important that nodes have loosely synchronized clocks. However, we can easily tolerate a one-minute skew because certificates are valid for at least a few minutes. The most obvious effect of a large skew is that authentication becomes impossible because the validity interval of a formula is empty or does not include the current time. If a certificate originates at a node whose clock is much later than real time, or is used at a node whose clock is much earlier, it is also possible that the certificate will be taken as valid even though it should have expired. For synchronization, we do not have a secure time server, and instead rely on the clocks of individual nodes.

2.1 Some notations and rules

We write A **says** S to mean that principal A supports the statement S (an assertion or a request). We write $A \Rightarrow B$ when A *speaks for* B , meaning that if A makes a statement then B makes it too:

if $(A \Rightarrow B)$ **and** $(A \text{ says } S)$
then $(B \text{ says } S)$

We think of A as being stronger than B . The \Rightarrow relation is a partial order. It obeys many of the same laws as implication, so we use the same symbol for it.

Principals include:

- *Simple principals.* Users, machines.
- *Principals in roles.* We write A **as** R for A in role R (for example, *Bob as Admin* for *Bob* acting as an administrator). A principal can adopt a role in order to reduce its rights [9, Section 6]. That is, $A \Rightarrow (A \text{ as } R)$.
- *Channels.* Network addresses, encryption keys. C **says** S if S appears on channel C . In particular, K **says** S represents a certificate containing S and signed by K . A channel is the only kind of principal that can directly make a statement, since a message can arrive only on a channel.
- *Groups.* Sets of principals. If A is a member of G and A **says** S , then G **says** S , so $A \Rightarrow G$.
- *Conjunctions of principals.* The principal $A \wedge B$ **says** S when both A and B do.
- *Principals quoting principals.* The principal $B|A$ **says** S when B **says** that A **says** S .
- *Principals acting on behalf of others.* The principal B **for** A **says** S when B **says** that A **says** S and in addition A has authorized B to act on its behalf (A has delegated to B).

The operations **as**, \wedge , $|$, and **for** are monotonic with respect to \Rightarrow , so for example if $B \Rightarrow B'$ then $(B \text{ for } A) \Rightarrow (B' \text{ for } A)$. The *handoff axiom* represents the transfer of authority:

if A says $B \Rightarrow A$
then $B \Rightarrow A$

In other words, we believe that B speaks for A when A says it. Similarly, we have a *delegation axiom*:

if A says $(B|A) \Rightarrow (B \text{ for } A)$
then $(B|A) \Rightarrow (B \text{ for } A)$

which means that we believe A when it says that $B|A$ speaks for B for A , that is, that B can act as A 's delegate.¹ Delegation differs from handoff in that both parties lend authority to the composite, and B quotes A in speaking for B for A so that it does not exercise this authority accidentally.

There is a natural role associated with many groups, for example the role of administrator with the group of administrators. Hence we use group names as roles, and adopt the general rule that if A is a principal, G a group, and $A \Rightarrow G$ then $(A \text{ as } G) \Rightarrow G$.

2.2 Logic and authentication

This section gives a simplified example of how logic can be used to reason about authenticating compound principals; there is more detail in later sections. In the example a machine $Vax4$ is booted with an operating system OS . Together, $Vax4$ and OS form a node WS . A user Bob logs in to WS . We consider the reasoning necessary to authenticate requests from this login session to a file server FS .

In order to establish credentials, $Vax4$ must possess a secret. For example, if $(K_{vax4}, K_{vax4}^{-1})$ is a RSA key pair, then K_{vax4}^{-1} is a suitable secret. Let K_{vax4}^{-1} be available only to $Vax4$'s boot firmware, not to any of the operating systems it can run. At boot time, K_{vax4}^{-1} is used to sign a *boot certificate* that transfers authority to a newly generated key K_{ws} ; in the logic, this certificate reads:

$$(K_{vax4} \text{ as } OS) \text{ says } K_{ws} \Rightarrow (K_{vax4} \text{ as } OS) \quad (1)$$

We call K_{ws} the *node key* for WS . It speaks not for K_{vax4} but for a weaker principal $WS = (K_{vax4} \text{ as } OS)$, that is, K_{vax4} in the role of the boot image. After booting, WS gets the boot certificate and K_{ws}^{-1} , but does not know K_{vax4}^{-1} .

¹This axiom is not included in [9], but is suggested in [2], we adopt it for simplicity.

We treat login as a specialized form of delegation. When Bob logs in, K_{bob}^{-1} is used to sign a *delegation certificate* that transfers authority to WS :

$$K_{bob} \text{ says } (K_{ws} | K_{bob}) \Rightarrow (K_{ws} \text{ for } K_{bob}) \quad (2)$$

Consider now a request from the login session to a file server FS . There must first exist a channel C_{bob} over which to issue requests. As observed by FS , a request appears as a statement on this channel:

$$C_{bob} \text{ says } RQ$$

In order to back this request, WS supplies (2) and writes a *channel certificate*:

$$(K_{ws} | K_{bob}) \text{ says } C_{bob} \Rightarrow (K_{ws} \text{ for } K_{bob}) \quad (3)$$

This represents a handoff from the node to the channel. Now, by applying the handoff axiom and the delegation axiom to (2) and (3), FS can deduce $(K_{ws} \text{ for } K_{bob}) \text{ says } RQ$. Given the boot certificate (1), FS can infer, by monotonicity:

$$((K_{vax4} \text{ as } OS) \text{ for } K_{bob}) \text{ says } RQ$$

We still must prove that K_{vax4} and K_{bob} correspond to $Vax4$ and Bob . To do this we must trust some *certification authority* or CA. Trusting a CA with known key K_{ca} means believing that K_{ca} speaks for any principal. Thus, FS can use the certificates

$$\begin{aligned} K_{ca} \text{ says } K_{vax4} &\Rightarrow Vax4 \\ K_{ca} \text{ says } K_{bob} &\Rightarrow Bob \end{aligned}$$

the handoff axiom, and monotonicity, to deduce:

$$((Vax4 \text{ as } OS) \text{ for } Bob) \text{ says } RQ$$

That is, FS knows that $Vax4$ running OS requests RQ on behalf of Bob . The access control algorithm given in [9, Section 9] can now determine whether the request should be granted.

The remainder of the paper describes how this authentication logic is implemented in Taos.

3 An API for Authentication

The logic is rather complex to be presented directly through a programming interface. Instead, Taos defines a simple and consistent set of security services. They are based on an abstract datatype `Prin` that represents principals, and a subtype `Auth` that represents principals that processes can speak for.

Section 3.1 gives the interfaces for sending and receiving authenticated messages; that is, it explains how a process that can speak for a principal P can make another process believe P says S . Section 3.2

gives the interface for managing Auths; that is, it explains how a process can change the set of principals that it can speak for.

For brevity we omit procedure exceptions.

3.1 Authenticating messages

We begin with a simplified version of the interface for sending and receiving authenticated messages:

```
PROC Send(dest:Address; p:Auth; m:Msg);
PROC Receive(): (Prin, Msg);
```

```
PROC Authenticate(p:Prin): TEXT;
PROC Check(acl:ACL; p:Prin): BOOL;
```

`Send` transmits the statement `p` says `m` to the process at address `dest`. Symmetrically, if `Receive` returns `(p, m)`, some process that speaks for `p` has invoked `Send(p, m)`; in other words, the receiver can believe that `p` says `m`. The receiver can then use `Authenticate` to turn `p` into a string name. This string can represent a compound principal such as *Bob as admin*, or it can be a simple name. Simple names are convenient for existing applications. Section 4.5 describes the somewhat ad hoc rules Taos uses to reduce compound principals to simple names.

The ultimate goal of authentication is to tell the authorization service the source of a request. We therefore introduce another abstract datatype `ACL` to represent access control lists, and the authorization operation `Check` to determine whether `acl` grants access to `p`. `Check` both hides the details of naming and allows a convenient and efficient cache of recent successful authorizations. There are also operations for constructing and examining ACLs, but they are beyond the scope of this paper.

The `Authenticate` and `Check` procedures deal only in `Prins` and do not depend on how a `Prin` is transmitted. They remain unchanged as we improve the rest of the interface in this section.

This interface has no notion of a principal that a process speaks for by default. Instead, the `Auth` argument to `Send` requires the process to specify explicitly the principal that is uttering each message. Often a process has only one `Auth`, and we could have added a “working authority” to the process state and a `SetWorkingAuth` procedure (by analogy with the working directory), and dropped the `Auth` argument to `Send` or made it optional. This is similar to what Unix does with the effective uid. Or, to accommodate multi-threaded programs, we could have made the working authority part of the thread state.

This simple design is unsatisfactory because it ties authentication and communication too closely. To

separate them, we make explicit the relation between a channel `c` and the principal `p` that it speaks for.

The communication system makes secure channels on which a process can send. A channel is secure if every message received on it is sent by that process. We might also want messages on the channel to be secret, that is, received only by certain processes; this is a simple extension that we will not discuss further. An abstract datatype `Chan` represents secure channels.

To transmit an authenticated message, a process sends it on a secure channel, the receiver gets the channel `c` on which the message arrives, and a new operation `GetPrin` returns the `p` that the channel speaks for. In other words, `c` names the principal `p`.

For this to work a given channel must speak for at most one principal, so we need a cheap way to make channels. Our way is to take a single channel `c` on which a process can send securely, and then to multiplex many subchannels onto `c`, one for each principal that the process speaks for. Sending and receiving is done on these subchannels.

Our second try at an interface is thus:

```
PROC GetChan(dest:Address): Chan;
PROC GetSubChan(c:Chan; p:Auth): SubChan;
PROC Send(dest:SubChan; m:Msg);
PROC Receive(): (SubChan, Msg);
PROC GetPrin(c:SubChan): Prin;
```

The sending process calls first `GetChan` to get a secure channel `c` to the process at `dest` and then `GetSubChan(c, p)` to get a subchannel that speaks for `p`. The receiver calls `GetPrin` to recover a `Prin`.

The actual Taos interface has a further refinement: a process can utter many statements, perhaps made by different principals, in a single message. For example, one call could pass an array of names of files to delete and a parallel array of principals that are authorized to do the deletions. To make this work we must reveal the addressing mechanism for subchannels: it’s an integer called an *authentication identifier* or AID. The sender calls `GetAID` to learn the AID `aid` for a principal and sends it as an ordinary data value in the message. The receiver pairs the channel on which the message arrives with `aid` to recover the speaking principal. So the actual Taos interface is:

```
PROC GetChan(dest:Address): Chan;
PROC GetAID(p:Auth): AID;
PROC Send(dest:Chan; m:Msg);
PROC Receive(): (Chan, Msg);
PROC GetPrin(c:Chan; aid:AID): Prin;
```

In Taos the messages exchanged in this way are normally the call and return messages of remote procedure calls. RPC marshals an `Auth` parameter `p` by

sending the result of `GetAID(p)`, and unmarshals `aid` from channel `c` as the result of `GetPrin(c, aid)`. It also gets the channel from the RPC binding, and of course it encapsulates the `Send` and `Receive`. The result is that the RPC client can simply use `Prins` and `Auths` as arguments and results, and does not have to call any of the procedures in this interface except `Authenticate` and `Check`. This works for both calls and returns, so mutual authentication is possible.

3.2 Managing principals

A Taos process can obtain an `Auth` in five ways:

- by inheritance from a parent process,
- by presenting a login secret,
- by adopting a role,
- by delegating rights, or
- by claiming delegated rights.

All but the first of these produce a new and unique `Auth`. In particular, each user session on a machine is represented by a different `Auth`.

The interface for managing `Auths` is:

```
PROC Self(): Auth;
PROC Inheritance(): ARRAY OF Auth;
PROC New(name, password: TEXT): Auth;
PROC AddRole(a:Auth; role:TEXT): Auth;
PROC Delegate(a:Auth; b:Prin): Auth;
PROC Claim(b:Auth; delegation:Prin): Auth;

PROC Discard(a:Auth; all:BOOL);
```

`Self` returns a default `Auth` for the current process. The default is specified when the process is created. `Inheritance` returns all the `Auths` the process inherits from its parent.

`New` is used to generate entirely new credentials. The parameters describe a user name and a user-specific secret sufficient to generate the cryptographic credentials described in Section 4.3. The result is an `Auth` that represents *node for name*, where *node* is the local node. This result reflects the fact that the user cannot make a request without involving the machine and the operating system.

`AddRole` creates a weaker authority than a by applying a role. If `a` represents *A*, then the result represents *A as role*.

Roles are used in two ways in Taos. First, a process can restrict its rights to what is necessary to fulfill a particular function by calling `AddRole` on one of its existing `Auths`. Second, a Taos node can give some




Figure 1: An example of delegation

of its rights to a trusted process. Taos uses *secure loading* to determine whether an executable image is certified (see Section 4.4). After loading a certified image, Taos calls `AddRole` to create an `Auth` weaker than its own which it hands off to the new process (for example, `AddRole(Self(), "telnet-server")` for a login daemon). This mechanism bears some resemblance to Unix *setuid* execution. However, in Taos there is a stronger guarantee about the loaded program, and the program need not receive all the rights of the node. Further, the resulting rights, like those of the node, can be exercised over the network.

The procedures `Delegate` and `Claim` are used in tandem to implement delegation; Figure 1 shows an example. Suppose process *X* has an `Auth` `a` that represents *A*, process *Y* has an `Auth` `b` that represents *B*, and *X* wants to give to *Y* an authority that represents *B for A* by delegation. First, *X* gets from *Y* a `Prin` `p_b` that represents *B*. Then *X* calls `Delegate(a, p_b)` to make a new `Auth` `c` that represents *A* but also carries the property that *A says (B|A) => (B for A)*. Now *X* sends `c` to *Y*, which receives it as the `Prin` `p_c`. Finally *Y* calls `Claim(b, p_c)` to get an `Auth` `d` that represents *B|A*, and hence *B for A* by the delegation axiom. Before doing this, *Y* may wish to call `Authenticate(p_c)` to find out what principal `d` will represent.

A process can make an `Auth` `a` *invalid* by calling `Discard`. The effect is that the process can no longer use `a` to speak for `a`'s principal (although receivers might not find this out until their caches time out). If `all` is `TRUE`, `a` also becomes invalid in all the processes that have inherited it; this allows a process to take an authority away from its children.

If `a` was the result of `Delegate`, invalidating it has another effect: any `Auth` derived from `a` by `Claim` will also become invalid within a fairly short time (at most 30 minutes). The same thing happens if the process that called `Delegate` terminates.

The API provides no direct access to the logical operators `|` and `^` or to the handoff rule.




Figure 2: Structure of the authentication agent

4 The Authentication Agent

The authentication agent is responsible for handling most of the complexity of authenticating requests from compound principals. It has four parts. The secure channel manager creates process-to-process secure channels. The authority manager associates Auths with processes and handles authentication requests. The credentials manager maintains credentials on behalf of local processes and validates certificates authored on other nodes. Finally, the certification library establishes a trusted mapping between principal names and cryptographic keys, and between groups and their members. Figure 2 shows the structure of the authentication agent; arrows indicate call dependencies.

Only a few changes were needed to the rest of Taos to support the authentication agent: implementing authority inheritance in the process manager; supporting secure loading; and adding Auth parameters to all security-sensitive kernel calls.

4.1 The secure channel manager

The secure channel manager implements the `Chan` datatype described in Section 3.1. It controls the construction of node-to-node channels, and then uses them to provide process-level channels to its clients. Since the purpose of authentication is to prove that a channel utters a request on behalf of a principal, the secure channel manager must be able to attribute channels to processes and thereby link channels to the principals for which they speak. Our design does not mandate any one technique for implementing secure channels; such techniques are well documented [11].

4.1.1 Node-to-node channels

Given two nodes WS and WS' , it is easy to establish a shared-key channel C between them; we use the protocol described in [9, Section 4] to exchange shared keys. Let WS have node key K_{ws} ; then C speaks for K_{ws} from the point of view of WS' . When the protocol is reexecuted, C still speaks for K_{ws} , so C can be rekeyed without invalidating existing authentication state based on K_{ws} . Each node is responsible for caching and timing out the keys it shares with other nodes, and either end of a secure channel can trigger the generation of a new shared key.

The secure channel manager maintains a cache of keys shared with other nodes, indexed by node address and used to implement `GetChan` and `Send`. Another cached mapping from shared keys to node keys is used by `Receive` to get from the shared key that successfully decrypts a message to a node key.

Both of these caches can be flushed as necessary. In fact, both are flushed periodically in order to invalidate old keys. The key-to-node-key mapping is flushed with a period that is twice that of the address-to-key mapping so as to prevent misses caused by partners using older keys.

Taos does not implement hardware secure channels. The key exchange mechanism it implements is, however, suitable for constructing them. Herbison [7] discusses the use of encrypting network controllers to build efficient secure channels. Our system design is intended to operate best with encryption-capable controllers. DES hardware for such controllers has been shown to operate at speeds of 1 Gbit/sec [5], so performance should not be a problem.

In our implementation, software DES is used to sign channel certificates (see Section 4.3.4), but requests are made without signature to avoid the overhead of software encryption.

4.1.2 Process-to-process channels

The channels offered to clients of the API are always between two processes. These channels are formed by multiplexing process-level data across the node-to-node channels discussed in the previous section. The concrete form of the `Chan` datatype differs depending on the secure channel implementation. However, all channel implementations must support naming of channels by ChanIDs:

```
TYPE ChanID = { nk:NodeKeyDigest;
                pr:INTEGER; addr:Address };
```

The `nk` field of `ChanID` names the node key of the partner, `pr` identifies the partner process, and `addr`

indicates the address of the partner authentication agent. A message digest function is applied to node keys in order to produce small values for the `nk` field.

In Taos we exploit the fact that most communication employs a transport protocol under our control. We identify each process with a 32-bit integer *process tag* (PTag).² The operating system ensures that all transmissions are tagged with valid PTags. Thus, the receiver of a message can name its origin with a Chan containing the sender's node key plus a source `pr` equal to the sender's PTag.

The secure channel manager exports the primitives:

```
PROC GetChanID(ch:Chan): ChanID;
PROC PTagFromChan(c:ChanID): PTag;
```

`GetChanID` returns a ChanID given an abstract Chan. If `PTagFromChan(c)` is invoked at `c`'s source—where `c.nk` is the local node key—it returns the PTag for the process that controls `c`; otherwise it fails.

Process-level multiplexing can also be done with standard protocol implementations such as TCP/IP and UDP/IP that have small integer *port numbers* that identify the origin and destination of messages within a node. Port numbers would be perfect process identifiers if they were not reuseable. One possible workaround is to place restrictions on the reuse of port numbers. Another is to treat process channels as secure connections that must be explicitly opened and closed. This requires considerable care.

4.2 The authority manager

The authority manager implements the operations on Auths and Prins discussed in Section 3.2. The internal interface to the authority manager parallels the API quite closely. However, for each Auth supplied as an argument, the kernel call dispatcher appends the PTag of the caller. This PTag argument is used to ascertain that the caller *owns* each supplied Auth. We say that a process owns an Auth if the authority manager has given that process the right to use it. Whenever an Auth is explicitly returned to a process through the API, the calling process owns it.

Each new Auth is assigned a unique AID by the authority manager. In our implementation, AIDs are 96 bits wide, so there is no need to reuse one. The authority manager maintains a table of the Auths it creates, indexed by AID. Each entry contains:

- credentials for an Auth,
- a list of PTags of processes that own this Auth,

²Do not confuse PTags with 16-bit Unix process IDs, which can wrap around. Our PTags are never reused; a consequence is that Taos can create only 2^{32} processes per boot.

- credentials for unclaimed delegations (only if this Auth resulted from a call to `Delegate`), and
- a source from which to refresh delegation credentials (only if this Auth resulted from a call to `Claim`).

Much like Unix file descriptors, Auths can be passed by inheritance to child processes. The process manager implements this inheritance by calling:

```
PROC Handoff(a:Auth; ptag:PTag);
PROC PurgePTag(ptag: PTag);
```

`Handoff` adds `ptag` to the list of PTags of processes that own `a`. It is invoked when an Auth is inherited from a parent process. `PurgePTag` eliminates all instances of `ptag` in the credentials table. It is called when the process identified by `ptag` terminates.

4.2.1 Callbacks

As we have seen, AIDs and channels can be used to represent principals in network protocols. In order to allow this, the authority manager must be prepared to produce credentials on behalf of any Auth it manages. These credentials are obtained with callbacks. This saves the cost of passing complex credentials repeatedly. In fact, credentials are generated lazily, only when needed, and AIDs may be passed before the corresponding credentials exist. Although credentials could easily be bundled with requests, they are large enough (> 1 kbyte) to affect communications performance. Since the results of authentication are cached extensively, callbacks improve performance for nearly all applications, even in high-latency networks.

Suppose a user-level process receives a request on a channel `ch`. In this case, the API function `GetPrin` returns a Prin `p` constructed from `GetChanID(ch)` and the AID accompanying the request. Now the process can ask its authentication agent to resolve `p` into a principal name, for example with a call to `Authenticate(p)`. We use the PrinID datatype to represent Prins when passed across address-space boundaries (for instance between user space and the authentication agent):

```
TYPE PrinID = { ch:ChanID; aid: AID };
```

The implementation of `Authenticate(p)` asks the requester's agent (at `p.ch.addr`) to provide credentials for `p`. This agent looks up `p.aid` in its credentials table and determines whether `PTagFromChan(p.ch)` specifies a process that owns the corresponding Auth. If it does, the requester's agent returns a *channel certificate* as proof that the channel speaks for the principal that `p` represents. This proof consists of the

credentials found in `p.aid`'s hash-table entry and a statement that `p.ch` quoting `p.aid` speaks for the principal (see Section 4.3.1).

It is critical for performance that the results of `Authenticate` be cached. Caching can be implemented in user space, in the operating system, or both. Our implementation caches the results of authentication callback in user space with a timeout equal to the lesser of 30 minutes and the validity interval of the supplied channel certificate.

A callback also occurs when a call `Claim(me, p)` activates a delegation. The delegate's authentication agent passes `p` in a callback to the delegator's agent, which uses `p.aid` to find credentials suitable for signing a delegation certificate and returns a signed certificate to the delegate's agent. The delegation certificate need not be concealed. Any agent may request a copy, since it is useful only to the delegate's agent. That agent must remember `p` so that it can repeat the callback to refresh the delegation in case it expires.

4.3 The credentials manager

The credentials manager is the heart of the Taos authentication system. The primary functions of the credentials manager are to build, to check, and to store credentials. We explain the form of credentials and their logical meaning in the first two subsections. Then we give the interface to the credentials manager and discuss techniques for avoiding signatures.

4.3.1 Credentials

We understand credentials as having logical meanings. A credential is evidence that one principal Q speaks for another principal P . If the credential were written as a formula M , its recipient would want to check that M implies $Q \Rightarrow P$.

Taos encodes credentials as S-expressions. The encoding is designed to make straightforward the proof of the theorem that M implies $Q \Rightarrow P$. If an S-expression is a well-formed credential, then the corresponding proof that M implies $Q \Rightarrow P$ can easily be constructed from the S-expression. If in addition all signatures in the S-expression are recent and correct, then the S-expression is said to be valid; the S-expression is interpreted as M only if it is valid. Thus, deriving $Q \Rightarrow P$ is reduced to parsing a credential and checking signatures.

In this section we define our S-expression grammar for credentials. In Section 4.3.2 we give a table of correspondences between S-expressions and logical formulas, effectively recovering the logical form of a credential from the S-expression encoding. This

```

channel    = ( 'channel' prin prinID sig )
boot       = ( 'boot' k_as key sig )
login      = ( 'login' k_as session sig )
session    = ( 'session' key boot sig )
delegation = ( 'for' delegator delegate sig )

p_as       = ( 'as' prin role )
k_as       = ( 'as' k_as role ) | primary
primary    = ( key name )

prin       = boot | login | delegation | p_as
delegator  = delegate = prin
role       = name

```

Table 1: Grammar for credentials

logical form is used only in explaining our implementation; the implementation does not manipulate formulas. We also describe how to check whether a credential is valid.

Table 1 gives the grammar for credentials. Names, keys, PrinIDs, and signatures are terminals. The main production is the one for `channel`, because requests always arrive on channels. Primary names are only hints, used to simplify the mapping of keys into names. We say that a credential y is embedded in a credential x if y is a subexpression of x .

The signature in a certificate includes the interval of time for which it is valid plus an unforgeable value identifying the signer. This value is a digest of the certificate, encrypted by a RSA secret key. The digest is computed over the entire certificate, excluding embedded signatures, by a one-way function that reduces its input to a size small enough to sign conveniently; it is one-way in the sense that it is computationally hard to find a different input with the same digest. Taos uses MD4 [14] as a digest function.

We now discuss specific credentials in some detail. For each type of signed credential, we discuss an example, borrowing context from Section 2.2.

Boot certificates. A boot certificate describes a handoff from a machine key to a node key. In our example, the meaning M of the boot certificate is:

$$(K_{vax4} \text{ as } OS) \text{ says } K_{ws} \Rightarrow (K_{vax4} \text{ as } OS)$$

From M and the handoff axiom, we obtain:

$$K_{ws} \Rightarrow (K_{vax4} \text{ as } OS)$$

which is the formula $Q \Rightarrow P$ in this case. The boot certificate is encoded as:

$$(\text{boot } (\text{as } (K_{vax4} \text{ } Vax4) OS) K_{ws} sig_1)$$

Login and session certificates. A *login certificate* is a special form of delegation certificate. It denotes a delegation from a user's key to the conjunction of a node key with a temporary *session key*. The user's key should be in memory for the shortest possible time, to reduce the chance that the key will be discovered by an attacker. In Taos, it is present just long enough to sign the login certificate. This certificate is of long duration, on the order of days. More sophisticated login protocols could be used, since they can produce an equivalent certificate [1].

The temporary session key ensures that the login certificate becomes invalid when the user logs off, to make up for the longevity of the login certificate. The node key and the session key are combined in a *session certificate*, which represents a handoff from the temporary session key to the node key. A session certificate has a short timeout and is refreshed as needed until the end of the session. When the session ends, the temporary key is discarded so that the session certificate can no longer be refreshed.

In our example, *Bob*, with key K_{bob} , logs in to *WS*. We still have the boot certificate, so:

$$(K_{vax4} \text{ as } OS) \text{ says } K_{ws} \Rightarrow (K_{vax4} \text{ as } OS)$$

Let K_s be the session key; the session certificate adds:

$$K_s \text{ says } (K_{vax4} \text{ as } OS) \Rightarrow K_s$$

and the login certificate adds:

$$K_{bob} \text{ says } (P_1 | K_{bob}) \Rightarrow (P_1 \text{ for } K_{bob})$$

where P_1 is $((K_{vax4} \text{ as } OS) \wedge K_s)$. From the conjunction of these formulas we can derive:

$$(K_{ws} | K_{bob}) \Rightarrow (P_1 \text{ for } K_{bob})$$

In the notation introduced above, the conjunction is M , and the principals $K_{ws} | K_{bob}$ and $P_1 \text{ for } K_{bob}$ are Q and P , respectively.

In our encoding the session certificate is embedded inside the login certificate, and the boot certificate inside the session certificate:

```
(login
  (Kbob Bob)
  (session
    Ks
    (boot (as (Kvax4 Vax4) OS) Kws sig1)
    sig2)
  sig3)
```

The embedded certificates identify the machine, the node, and the session key, and give credentials for them.

General delegation certificates. The general form of delegation involves transfer of rights between principals. Suppose that *Bob* on *WS* delegates to a node *Vax5* as *OS* with key $K_{ws'}$. In our example, the formula that corresponds to this delegation is:

$$(K_{ws} | K_{bob}) \text{ says } (P_3 | P_2) \Rightarrow (P_3 \text{ for } P_2)$$

where P_2 is $(P_1 \text{ for } K_{bob})$ and P_3 is $(K_{vax5} \text{ as } OS)$. Conjoining this formula with those for *Bob's* login, we can prove:

$$(K_{ws'} | K_{ws} | K_{bob}) \Rightarrow (P_3 \text{ for } P_2)$$

In our encoding the entire delegation certificate is:

```
(for
  (login... sig3)
  (boot (as (Kvax5 Vax5) OS) Kws' sig4)
  sig5)
```

The login certificate given above is nested here in its entirety (abbreviated with an ellipsis) and used as the source of a delegation. The delegate is the boot certificate for *Vax5* as *OS*.

Channel certificates. Ultimately, channels are the only principals that make requests directly. A request on a channel is attributed to a principal that has handed off some of its rights to the channel. A channel certificate represents a handoff to a channel. In our system, each certificate authenticates a channel multiplexed on a node-to-node key. More precisely, the channel is a node-to-node channel quoting a process quoting an AID. Its encoding is a textual representation of the PrinID datatype from Section 4.2.

In our example, a channel certificate means:

$$(K_{ws} | K_{bob}) \text{ says } C_{bob} \Rightarrow (((K_{vax4} \text{ as } OS) \wedge K_s) \text{ for } K_{bob})$$

Conjoining this formula with those for *Bob's* login, we can prove:

$$C_{bob} \Rightarrow (((K_{vax4} \text{ as } OS) \wedge K_s) \text{ for } K_{bob})$$

When C_{bob} is the channel $key47|ptag13|aid42$, this certificate is encoded as:

```
(channel
  (login... sig3)
  key47 ptag13 aid42
  sig4)
```

Because channels are typically short-lived, a channel certificate normally has a short validity interval.

x	$S(x)$	$T(x)$	$P(x)$	$Q(x)$
boot	$Q(x.k_as)$	$x.key$	$P(x.k_as)$	$x.key$
session (s)	$x.key$	$P(x.boot)$	$x.key$	$Q(x.boot)$
login	$Q(x.k_as)$	$(P(x.s.boot) \wedge P(x.s))$ $ P(x.k_as)$	$(P(x.s.boot) \wedge P(x.s))$ for $P(x.k_as)$	$Q(x.s.boot)$ $ Q(x.k_as)$
delegation	$Q(x.delegator)$	$P(x.delegate)$ $ P(x.delegator)$	$P(x.delegate)$ for $P(x.delegator)$	$Q(x.delegate)$ $ Q(x.delegator)$
channel	$Q(x.prin)$	$x.prinID$	$P(x.prin)$	$x.prinID$
p_as			$P(x.prin)$ as $x.role$	$Q(x.prin)$ as $x.role$
k_as			$P(x.k_as)$ as $x.role$ or $P(x.primary)$	$Q(x.k_as)$ as $x.role$ or $Q(x.primary)$
primary			$x.key$	$x.key$

The immediate meaning $I(x)$ of a credential x is $S(x)$ says $T(x) \Rightarrow P(x)$ when $S(x)$ is defined, and *true* otherwise. The meaning $M(x)$ of a credential x is the conjunction of $I(x)$ with the immediate meanings of any credentials embedded in x . In all cases, $M(x)$ implies $Q(x) \Rightarrow P(x)$. We abbreviate **session** by **s**.

Table 2: The logical meaning of credentials

4.3.2 The meanings of credentials

As the previous examples suggest, each valid credential x in the grammar has a logical meaning $M(x)$. Now we define M in general. Since M is a function, the mapping from S-expressions to formulas is clearly unambiguous. We define validity later in this section.

It is convenient to use several auxiliary functions. A function I gives us the immediate meaning of a credential. Then $M(x)$ is defined to be $I(x)$ conjoined with $I(y)$ for every credential y embedded in x . Thus, the interpretation of a credential is its immediate meaning, plus the meaning of any embedded credentials. In the cases of **primary**, **p_as**, and **k_as** credentials, which bear no signatures, $I(x)$ is simply *true*. In the other cases, $I(x)$ is the assertion made by the top-level signature; it does not refer to other signatures or their timestamps, and has the form

$$S(x) \text{ says } T(x) \Rightarrow P(x)$$

where $P(x)$ and $T(x)$ are principals and $S(x)$ is the speaker, the principal that issues the credential. In particular, when $S(x)$ is a key, it is the key that should be used in the credential's signature.

In each case, the purpose of a credential x is to establish that $Q(x)$ speaks for $P(x)$. More precisely, the formula $M(x)$ should imply $Q(x) \Rightarrow P(x)$. For example, a boot certificate x of the form

$$(\text{boot } (K_{vax4} \text{ } \forall ax4) K_{ws} \text{ sig})$$

is intended to mean K_{vax4} says $K_{ws} \Rightarrow K_{vax4}$; this formula is $M(x)$. Let $Q(x)$ be K_{vax4} and $P(x)$ be

K_{ws} ; by the handoff axiom, $M(x)$ implies $Q(x) \Rightarrow P(x)$. In general, in order to derive $Q(x) \Rightarrow P(x)$ from $M(x)$ for a certificate x , it suffices to obtain:

1. $Q(x) \Rightarrow T(x)$, and
2. if $S(x)$ says $T(x) \Rightarrow P(x)$ then $T(x) \Rightarrow P(x)$.

In all cases (1) will follow from the meanings of embedded credentials. To obtain (2), we may use either

- $S(x) \Rightarrow P(x)$, and then the handoff axiom applies; or
- $P(x)$ is B **for** A and $T(x)$ is $B|A$ for some A and B such that $S(x) \Rightarrow A$, and then the delegation axiom applies.

The definitions listed in Table 2 satisfy these properties. We obtain that $M(x)$ implies $Q(x) \Rightarrow P(x)$ for every x , by induction on the structure of credentials. This theorem guarantees that validating x suffices to show that $Q(x) \Rightarrow P(x)$: if x is valid, then it is interpreted as $M(x)$ and then the theorem applies.

A credential is valid if all the signatures it contains are well-formed, timely, and were performed with the proper key. The proper key K for signing a certificate x is defined from $S(x)$, with a clause for each of the possible forms of $S(x)$:

- The proper key for a principal of the form A **as** R or $A|A'$ is the proper key for A , since it is A that must apply the signature.

- The proper key for a key is the key itself.

In general, K is the key that the principal $S(x)$ uses. If x is valid, then it has recently been signed with $S(x)$'s key K , so we can interpret x as a formula $I(x)$ of the form $S(x)$ says $T(x) \Rightarrow P(x)$. This is the justification for our logical reading of valid credentials.

An obvious generalization of this definition of validity allows the key that signs to be any key that speaks for K . This generalization is used in Section 4.3.4 to allow channel certificates to be signed with DES keys.

4.3.3 The Credentials interface

The credentials manager exports the `Credentials` interface to the authority manager. This interface defines an abstract type `CredT` that represents sets of credentials, as well as procedures for constructing `CredT`s and for signing and validating channel certificates. A `CredT` defines a principal P that can make requests, and contains an expression in the credentials grammar sufficient to prove that some other principal can speak for P .

The credentials manager holds a `CredT` representing the credentials for the node. Although the Firefly lacks the firmware necessary to generate a node key securely, Taos imitates secure booting by generating a boot certificate and node key at system-startup time. The node's `CredT` contains this certificate and key.

The operations on credentials are:

```

TYPE Cred = TEXT;
PROC New(name, password: TEXT): CredT;
PROC AddRole(t:CredT; role:TEXT): CredT;
PROC Sign(t:CredT; p:PrinID): Cred;
PROC Validate(cr:Cred; p:PrinID): TEXT;
PROC Extract(cr:Cred): Cred;
PROC SignDlg(t:CredT; cr:Cred): Cred;
PROC ClaimDlg(t:CredT; cr:Cred): CredT;

```

`New` produces a `CredT` containing a login certificate and a session key. The `CredT` returned by `AddRole` contains credentials for `t` as `role`.

Each value of the `Cred` datatype contains a textual representation of credentials according to the grammar of Table 1. The authority manager uses `Sign` to produce channel certificates in response to authentication callbacks. Similarly, it uses `Validate` to check the results of authentication callbacks and return principal names. `Extract` strips off an outer-level channel certificate, and returns the credentials of the principal for which the channel speaks.

The delegator's authority manager implements `Delegate` by finding and validating a channel certificate for the delegate. It then calls `Extract` to get the delegate's credentials, and stores the result.

The delegate's authority manager implements `Claim` by asking the delegator's agent for a delegation certificate (produced with `SignDlg`) and using it to call `ClaimDlg`. The result is a `CredT` representing *delegate for delegator*.

4.3.4 Signature techniques

We use three techniques to minimize the number of public key encryptions required to sign certificates:

- As described in Section 4.1, we can establish a shared key K between two nodes WS (with key K_{ws}) and WS' so that WS' believes that K speaks for K_{ws} . Therefore, WS can sign certificates about channels to WS' by encrypting with K instead of K_{ws} . Only WS' need believe these certificates. DES encryption (under K) is much faster than RSA encryption (under K_{ws}).
- When one process delegates to another on the same node, it is possible to avoid one signature. The delegation certificate structure remains the same, but no cryptographic signature is needed. If an off-node delegation follows, the signature of the outer certificate implies validity for the inner one. Both use the same key.
- When refreshing nested certificates, care must be taken not to invalidate higher-level signatures. It is sufficient to omit nested signatures from the certificate digests. For example, when a session certificate is refreshed, its validity times are changed. An enclosing login certificate can avoid refresh only if its digest omits the nested signature. This omission is safe since there is no mention of nested signatures in the immediate meaning of credentials.

4.4 The certification library

If ACLs contained public keys instead of human-sensible names, network security would be considerably less complex. Unfortunately, keys are big numbers that are too unwieldy for human users to manipulate. Moreover, at the highest level, computer security applies to names for people and resources. At some point there needs to be a mapping from keys to the principal names they represent.

The task of the certification library is to produce human-sensible names from keys. We also use it to recover keys from stable storage given passwords short enough for people to remember. Our certification authority (see Section 2.2) is a simple program that manages the database underlying these services. This

CA is off-line in the sense that clients need not communicate with it in order to trust its statements. A CA that could function without any network connections might be an interesting addition to our work. We could easily extend our system to incorporate a hierarchy of CAs [9, Section 5].

Bootstrapping trust. A practical system of any size must base trust on shared knowledge of a trusted CA. In Taos, this information takes the form of a CA public key. Certificates signed with this key are trusted. It is crucial to protect the corresponding secret key.

A user learns its own secret key and the public key of its trusted CA by decrypting a user-specific string stored in the name server. This string contains the user's private data encrypted under a DES secret derived from the user's password. We keep analogous strings for nodes. These strings would not be necessary if users carried public key smart-cards [1, 13].

Name certificates. Our system uses *name certificates* to describe the mapping from keys to names. These certificates are signed by a CA trusted for this purpose, much like CCITT X.509 certificates [4]. The logical form of a certificate that maps K_u to U is:

$$K_{ca} \text{ says } K_u \Rightarrow U$$

A simple extension of the grammar described in the previous section is used to express these statements.

Since certificates are statements signed off-line, they can be believed even if retrieved from untrusted storage. In Taos, we use a replicated, highly available name service [3] to store name certificates. Certificates are indexed by name in this store. The replication makes a denial of service attack more difficult.

We may now continue the example of Section 4.3.1. Given valid name certificates that map K_{bob} to Bob and K_{vax4} to $Vax4$, we obtain:

$$C_{bob} \Rightarrow ((Vax4 \text{ as } OS) \text{ for } Bob)$$

Therefore, when a request appears on the channel C_{bob} , it is attributed to $(Vax4 \text{ as } OS) \text{ for } Bob$.

Membership certificates. The certification library also manages and validates *membership certificates* stating that a principal U speaks for (is a member of) a group G :

$$K_{ca} \text{ says } U \Rightarrow G$$

Membership certificates are used in Taos ACL checking, and also in role processing and secure loading.

Image certificates. A third form of certificate managed by the certification library is the *image certificate*, used in secure loading to verify the executable image of a recently loaded program and to name the role under which that program should run. The certificate says that a given image digest I speaks for a role name R :

$$(K_u | R\text{-owner}) \text{ says } I \Rightarrow R$$

Image certificates are not signed by the CA. Instead, the CA permits a user U to write an image certificate for R :

$$K_{ca} \text{ says } (U | R\text{-owner}) \Rightarrow R$$

This says that U can quote $R\text{-owner}$ in order to speak for R , and hence U can release new versions of R by issuing the given image certificate.

Image digests can be computed using any secure one-way function. Taos stores image certificates as a file property on certain executable files.

4.4.1 The CertLib interface

The certification library exports the procedures:

```
PROC CheckKey(name:TEXT; k:Key): BOOL;
PROC IsMember(name, group: TEXT): BOOL;
PROC CheckImage(d:Digest; pg, cert: TEXT);
```

The credentials manager calls `CheckKey` to find and validate a name certificate that states that k speaks for $name$. The `IsMember` procedure ascertains whether $name$ is a member of $group$. `CheckImage` supports secure loading. It checks that the certificate $cert$ states that the image digest d speaks for the program pg , and that $cert$ is signed by a principal with control of images for pg .

4.5 Simplifying compound names

An authentication result in Taos is more often than not a compound principal. The principals that result from credential validations have the form:

```
principal = name
            | (principal for principal)
            | (principal as role)
```

where $name$ and $role$ are strings. Existing applications often deal only with simple names. The following function reduces a principal to a simple name:

- If the principal has a simple name, return it.
- If the principal is $B \text{ for } A$, apply this function recursively to A . (Checks can easily be added to guarantee that B is trustworthy.)

- If the principal is A as R , then apply this function recursively to A . Take the resulting simple name, and find a membership certificate stating that it speaks for R . If successful, then return R , otherwise fail.

For example, WS for Bob evaluates to Bob , and WS for $(Bob$ as $Admin)$ evaluates to $Admin$ if $Bob \Rightarrow Admin$ (that is, if Bob is a member of $Admin$).

4.6 Gateways

We built a gateway G between ordinary NFS clients and our system. It uses standard methods to determine the principal p making an NFS request and then forwards the request to the Taos file system as coming from $G|p$. Either this principal appears on ACLs, or there is a certificate K_{ca} says $(G|p) \Rightarrow q$.

This approach can be applied to accept messages authenticated by any other protocol. The tricky part is finding a place to put the gateway where it can intercept and translate the authentication protocol, which is often application-specific.

To go in the other direction and translate one of our authenticated messages p says m into another protocol, say Kerberos, the gateway would have to be able to authenticate itself as p in Kerberos. To achieve this, it would either need to have the user's password for long enough to obtain a Kerberos ticket-granting ticket, or to act itself as a Kerberos authentication server. We have not tried to implement this.

5 Performance

The authentication system described in this paper was in daily use for a year by a community of nearly 80 researchers and administrative personnel. The most commonly used authenticated application was Echo [3], a distributed file system used extensively within Taos. The Echo environment exercised all the Taos security features described in this paper except general delegation. In addition to authenticating the file system operations of normal users, Echo made use of role adoption and secure loading, both of which proved useful for system administration.

The performance of our system depends on the costs of the cryptographic operations:

RSA sign	RSA verify	DES	MD4
248 ms	16 ms	15 ms	6 ms/kbyte

Our RSA implementation [16] is coded in C and assembler. We use a 512-bit modulus and a public key

	<i>Auth login</i>	<i>Delegate</i>	<i>Auth delegation</i>
RSA sign	—	1×248 ms	—
RSA verify	3×16 ms	10×16 ms	7×16 ms
DES	2×15 ms	2×15 ms	2×15 ms
MD4	6 ms	18 ms	12 ms
S-expr	46 ms	165 ms	91 ms
RPC	2×5 ms	3×5 ms	2×5 ms
Total	140 ms	636 ms	255 ms
Measured	143 ms	671 ms	276 ms

Table 3: Authentication test timings

exponent of 3. The Firefly has 4 CVax processors, each running at about 2.5 MIPS. Our multiprocessor implementation of RSA signatures gains nearly a factor of two in speed. With only a single processor, it takes 472 ms to compute a RSA signature; this compares with 68 ms on a DECstation 5000, which runs at 20 MIPS. We use public-domain implementations of MD4 and DES (in C); much faster ones are possible [9, Section 4]. The DES number assumes that a different key is used for each 8-byte encryption.

In Table 3 we show the results of measuring three basic authentication operations. The numbers assume an existing node-to-node secure channel and a loaded name certificate cache. We show how time is divided between cryptographic functions and other parts of the system. We estimate that RPC with non-trivial arguments takes on the order of 5 ms [15]. The line labelled “S-expr” indicates the cost of parsing and writing S-expressions. This cost was about one-third of the total, but it could easily be reduced.

The first column of the table (*Auth-login*) shows the time required for the first authenticated RPC—subsequent calls to the same server using the same credentials will get cache hits. The caller's credentials are those for a simple login session. This test includes a callback to the caller's agent and a subsequent channel-certificate validation. We expect this cost to be incurred infrequently: for example, when the user's machine first contacts a file server, and whenever the credentials need refreshing thereafter (every 30 minutes).

The second test (*Delegate*) measures the time taken for a logged-in user to delegate to a logged-in user on another node. Delegation requires a hidden authentication, and hence three RPCs rather than two.

The final test (*Auth-delegation*) is similar to the first, except that the caller's credentials involve an additional delegation. Once again, the costs shown are incurred only on the first use of the credentials

and each time the cache is refreshed.

There are two important facts to be gleaned from Table 3. First, the cost of using credentials to make requests is considerably less than that of delegation. This is good, since delegations occur much less frequently than requests. Second, almost all of the component costs of authentication are compute-intensive. Moving to a faster processor should improve the actual performance linearly. The *Auth-login* test should take less than 25 ms on a DECstation 5000.

Even with faster processors, it is clear that caching at several levels is essential to system performance. The cache used to implement `Authenticate` prevents repeated authentication callbacks. It has a timeout of roughly 30 minutes, so there are at most two authentication callbacks to an Echo client in a 30 minute interval, regardless of the number of file system operations performed. The shared key cache in the secure channel manager prevents unnecessary key exchanges. The keys stored there expire with a much longer period (6 hours). The certification library also maintains a cache that saves the results of name certificate validations. There a cached result can remain valid until the certificate expires, although we flush results more frequently to speed up revocation. Further caching is clearly possible. For example, the meanings of common embedded credentials (such as boot certificates) might be cached.

6 Conclusion

We have described a framework for security in distributed systems that is based on logic. The logic takes shape in an operating system that was in daily use by a substantial community. Our system employs compound credentials to express the complex relationships among users, machines, and programs, yet little of this complexity shows through to users and programmers. Moreover, the careful optimizations that surround our use of public key cryptography ensure that it does not hurt performance. Although our implementation was not used on a large scale, the technique of off-line certification with minimal reliance on on-line services is well suited to large naming hierarchies [9]. Our design can accommodate fast revocation of name certificates and auditing along the lines discussed elsewhere [2, 9], but we have not implemented these features.

We have explained our system in logical terms, and in particular obtained a theorem that relates concrete credentials and their logical meanings. It would be interesting to obtain further theorems to prove the correctness of our implementation. Even stating the

proper results remains a challenge.

The need for well-founded and expressive distributed security systems will grow with the speed of processors and networks, the number of interconnected entities, and the complexity of applications. Our work shows how to design practical systems that meet this need and offers evidence that such systems can be built and can perform well.

References

- [1] M. Abadi, M. Burrows, C. Kaufman, and B. Lampson. Authentication and delegation with smart-cards. To appear in *Science of Computer Programming*, 1993.
- [2] M. Abadi, M. Burrows, B. Lampson, and G. Plotkin. A calculus for access control in distributed systems. *ACM Trans. Prog. Lang. and Sys.* 15, 4, Oct 1993.
- [3] A. Birrell, A. Hisgen, C. Jerian, T. Mann, and G. Swart. The Echo distributed file system. Report 111, Systems Research Center, Digital Equipment Corp., Aug. 1993.
- [4] CCITT. Information processing systems - Open systems interconnection - The directory authentication framework. CCITT 1988 Recommendation X.509.
- [5] H. Eberle and C. Thacker. A 1 Gbit/second GaAs DES chip. *Proc. IEEE Custom Integrated Circuit Conf.*, 1992.
- [6] M. Gasser, A. Goldstein, C. Kaufman, and B. Lampson. The Digital distributed system security architecture. *Proc. 12th National Computer Security Conference*, NIST/NCSC, 1989, 305-319.
- [7] B. Herbison. Low cost outboard cryptographic support for SILS and SP4. *Proc. 13th National Computer Security Conference*, NIST/NCSC, 1990, 286-295.
- [8] J. Kohl, C. Neuman, and J. Steiner. The Kerberos network authentication service. Version 5, draft 3, Project Athena, MIT, Oct. 1990.
- [9] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in distributed systems. Theory and practice. *ACM Trans. Comp. Sys.* 10, 4, Nov. 1992, 265-310.
- [10] B. Lampson. Protection. *ACM Operating Systems Review* 8, 1, Jan. 1974, 18-24.
- [11] Roger Needham. Cryptography and Secure Channels. *Distributed Systems, 2nd Ed.*, S. Mullender (editor), ACM Press, 1993, 231-241.
- [12] Open Software Foundation. *Introduction to OSF DCE*, Revision 1.0, Dec 1992.
- [13] J.-J. Quisquater, D. de Waleffe, J.-P. Bournas Corsair: a chip card with fast RSA capability. *Smart Card 2000*, D. Chaum (editor), Elsevier, 1991, 199-206.
- [14] R. Rivest. The MD4 message digest algorithm. *Advances in Cryptology: Crypto '90*, Springer-Verlag LNCS, 1991, 303-311.
- [15] M. Schroeder and M. Burrows. Performance of Firefly RPC. *ACM Trans. Comp. Sys.* 8, 1, Feb. 1990, 1-17.
- [16] M. Shand and J. Vullemin. Fast implementations of RSA cryptography. *11th Symposium on Computer Arithmetic*, IEEE Computer Society, June 1993.
- [17] C. Thacker, L. Stewart, and E. Satterthwaite. Firefly: A multiprocessor workstation. *IEEE Trans. Computers* 37, 8, Aug. 1988, 909-920.