

Limitations of the Kerberos Authentication System

*S.M. Bellovin**
smb@ulysses.att.com
M. Merritt†
mischu@allegra.att.com

AT&T Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

The Kerberos authentication system, a part of MIT's Project Athena, has been adopted by other organizations. Despite Kerberos's many strengths, it has a number of limitations and some weaknesses. Some are due to specifics of the MIT environment; others represent deficiencies in the protocol design. We discuss a number of such problems, and present solutions to some of them. We also demonstrate how special-purpose cryptographic hardware may be needed in some cases.

1. INTRODUCTION

The Kerberos authentication system^[Ste88, Mill87, Brya88] was introduced by MIT to meet the needs of Project Athena. It has since been adopted by a number of other organizations for their own purposes, and is being discussed as a possible standard. In our view, both these decisions may be premature. Kerberos has a number of limitations and weaknesses; a decision to adopt or reject it cannot properly be made without considering these issues. (A *limitation* is a feature that is not as general as one might like, while a *weakness* could be exploited by an attacker to defeat the authentication mechanism.) Some improvements can be made within the current design. Support for optional mechanisms would extend Kerberos's applicability to environments radically different from MIT.

These problems fall into several categories. Some stem from the Project Athena environment. Kerberos was designed for that environment; if the basic assumptions differ, the authentication system may need to be changed as well. Other problems are simply deficiencies in the protocol design. Some of these are corrected in the proposed Version 5 of Kerberos,^[Koh189] but not all. Even the solved problems merit discussion, since the code for Version 4 has been widely disseminated. Finally, some problems with Kerberos are not solvable without employing special-purpose hardware, no matter what the design of the protocol. We will consider each of these areas in turn.

We wish to stress that we are not suggesting that Kerberos is useless. Quite the contrary — an attacker capable of carrying out any of the attacks listed here could penetrate a typical network of UNIX[®] systems far more easily. Adding Kerberos to a network will, under virtually all circumstances, significantly increase its security; our criticisms focus on the extent to which security is improved. Further, we recommend changes to the protocols that substantially increase security.

Beyond its specific utility in production, Kerberos serves a major function by focusing interest on practical solutions to the network authentication problem. The elegant protocol design and wide availability of the code has galvanized a wide audience. Far from a condemnation, our critique is intended to contribute to an understanding of Kerberos's properties and to influence its evolution into a

* Author's address: Steven M. Bellovin, Room 3C-536B, AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, New Jersey 07974.

† Author's address: Michael Merritt, Room 3D-458, AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, New Jersey 07974.

tool of greater power and utility.

Several of the problems we point out are mentioned in the original Kerberos paper or elsewhere.^[Davi89b] For some of these, we present protocol improvements that solve, or at least ameliorate, the problem; for others, we place them squarely in the context of the intended Kerberos environment.

1.1 Focus on Security

Kerberos is a security system; thus, though we address issues of functionality and efficiency, our primary emphasis is on the security of Kerberos in a general environment. This means that security-critical assumptions must be few in number and stated clearly. For the widest utility, the network must be considered as completely open. Specifically, the protocols should be secure even if the network is under the complete control of an adversary.¹ This means that defeating the protocol should require the adversary to invert the encryption algorithm or to subvert a principal specifically assumed to be trustworthy. Only such a strong design goal can justify the expense of encryption. (No “steel doors in paper walls”.) We believe that Kerberos can meet this ambitious goal with only minor modifications, retaining its essential character.

Some of our suggestions bear a performance penalty; others complicate the design of suggested enhancements. As more organizations make use of Kerberos, pressures to enhance or augment its functionality and efficiency will increase. Security has real costs, and the benefits are intangible. There must be a continuing and explicit emphasis on security as the overriding requirement.

1.2 Validation

It is not sufficient to design and implement a security system. Such systems, though apparently adequate when designed, may have serious flaws. Consequently, systems must be subjected to the strongest scrutiny possible. A consequence of this is that they must be designed and implemented in a manner that facilitates such scrutiny. Kerberos has a number of problems in this area as well.

2. WHAT’S A KERBEROS?

Before discussing specific problem areas, it is helpful to review Kerberos Version 4. Kerberos is an *authentication* system; it provides evidence of a *principal*’s identity. A principal is generally either a user or a particular service on some machine. A principal consists of the three-tuple

$\langle \text{primaryname}, \text{instance}, \text{realm} \rangle$.

If the principal is a user — a genuine person — the *primary name* is the login identifier, and the *instance* is either null or represents particular attributes of the user, i.e., `root`. For a service, the service name is used as the primary name and the machine name is used as the instance, i.e., `rlogin.myhost`. The *realm* is used to distinguish among different authentication domains; thus, there need not be one giant — and universally trusted — Kerberos database serving an entire company.

Kerberos principals may obtain *tickets* for services from a special server known as the *ticket-granting server*, or *TGS*. A ticket contains assorted information identifying the principal, encrypted in the private key of the service. (Notation is summarized in Table 1.)

$$T_{c,s} = \{s, c, \text{addr}, \text{timestamp}, \text{lifetime}, K_{c,s}\}K_s$$

Since only Kerberos and the service share the private key K_s , the ticket is known to be authentic. The ticket contains a new private session key, $K_{c,s}$, known to the client as well; this key may be used to

1. The Project Athena Technical Plan^[Mill87, section 2] describes a simpler threat environment, where eavesdropping and host impersonation are of primary concern. While this may be appropriate for MIT, it is by no means generally true. Consider, for example, a situation where general-purpose hosts also function as routers, and packet modification or deletion become significant concerns.

TABLE 1. Notation

c	client principal
s	server principal
tgs	ticket-granting server
K_x	private key of “ x ”
$K_{c,s}$	session key for “ c ” and “ s ”
$\{info\}K_x$	“ $info$ ” encrypted in key K_x
$T_{c,s}$	Ticket for “ c ” to use “ s ”
A_c	Authenticator for “ c ”
tgs	ticket-granting service
$addr$	client’s IP address

encrypt transactions during the session.²

To guard against *replay attacks*, all tickets presented are accompanied by an *authenticator*:

$$A_c = \{c, addr, timestamp\}K_{c,s}$$

This is a brief string encrypted in the session key and containing a timestamp; if the time does not match the current time within the (predetermined) clock skew limits, the request is assumed to be fraudulent.

For services where the client needs bidirectional authentication, the server can reply with

$$\{timestamp + 1\}K_{c,s}$$

This demonstrates that the server was able to read *timestamp* from the authenticator, and hence that it knew $K_{c,s}$; that in turn is only available in the ticket, which is encrypted in the server’s private key.

Tickets are obtained from the TGS by sending a *request*

$$s, \{T_{c,tgs}\}K_{tgs}, \{A_c\}K_{c,tgs}$$

In other words, an ordinary ticket/authenticator pair is used; the ticket is known as the *ticket-granting ticket*. The TGS responds with a ticket for server s and a copy of $K_{c,s}$, all encrypted with a private key shared by the TGS and the principal:

$$\{\{T_{c,s}\}K_s, K_{c,s}\}K_{c,tgs}$$

That key, $K_{c,s}$, and the ticket-granting ticket itself, are obtained at session-start time. The client sends a message to Kerberos with a principal name; Kerberos responds with

$$\{K_{c,tgs}, \{T_{c,tgs}\}K_{tgs}\}K_c$$

The client key K_c is derived from a non-invertible transform of the user’s typed password. Thus, all privileges depend ultimately on this one key.

Note that servers must possess private keys of their own, in order to decrypt tickets. These keys are stored in a secure location on the server’s machine.

3. THE KERBEROS ENVIRONMENT

The Project Athena computing environment consists of a large number of more or less anonymous workstations, and a smaller number of large autonomous server machines. The servers provide volatile

2. Technically speaking, $K_{c,s}$ is a *multi-session key*, since it is used for all contacts with that server during the life of the ticket.

file storage, print spooling, mailboxes, and perhaps some computing power; the workstations are used for most interaction and computing. Generally, they possess local disks, but these disks are effectively read-only; they contain no long-term user data. Furthermore, they are not physically secure; someone so inclined could remove, read, or alter any portion of the disk without hindrance.

Within this environment the primary need is for user-to-server authentication. That is, when a user sits down at a workstation, that person needs access to private files residing on a server. The workstation itself has no such files, and hence has no need to contact the server or even to identify itself.

This is in marked contrast to a typical UNIX system's view of the world. Such systems do have an identity, and they do own files. Assorted network daemons transfer files in the background, clock daemons perform management functions, electronic mail and news arrives, etc. If such a machine relied on servers to store its files, it would have to assert, and possibly prove, an identity when talking to these servers. The Project Athena workstations are neither capable nor in need of such; they in effect function as very smart terminals with substantial local computing power, rather than as full computer systems.³

What does this mean for Kerberos? Simply this: Kerberos is designed to authenticate the end-user — the human being sitting at the keyboard — to some number of servers. It is not a peer-to-peer system; it is not intended to be used by one computer's daemons when contacting another computer. Attempting to use Kerberos in such a mode can cause trouble.⁴

We make this statement for several reasons. First and foremost, typical computer systems do not have a secure key storage area. In Kerberos, a plaintext key must be used in the initial dialog to obtain a ticket-granting ticket. But storing plaintext keys in a machine is generally felt to be a bad idea;^[Morr79] if a Kerberos key that a machine uses for itself is compromised, the intruder can likely impersonate any user on that computer, by impersonating requests vouched for by that machine (i.e., file mounts). Additionally, the session keys returned by the TGS cannot be stored securely; of necessity, they are stored in some area accessible to `root`. Thus, if the intruder can crack the protection mechanism on the local computer — or, perhaps more to the point, work around it for some limited purposes — all current session keys can be stolen. This is less serious than a breach of the primary Kerberos key, of course, since session keys are limited in lifetime and scope; nevertheless, one does not wish these keys exposed.

This points out a second flaw when multi-user computers employ Kerberos, either on their own behalf or for their users: the cached keys are accessible to attackers logged in at the same time. In a workstation environment, only the current user has access to system resources; there is little or no need even to enable remote login to that workstation. There are many reasons for this; a consequence, though, is that the intruder simply cannot approach the safe door to try to pick its lock.⁵ Only when the legitimate user leaves can the attacker attempt to find the keys. But the keys are no longer available; Kerberos attempts to wipe out old keys at logoff time, leaving the attacker to sift through the debris. With a multi-user computer, on the other hand, an attacker has concurrent access to the keys if there are flaws in the host's security.

There are two other minor flaws in Kerberos directly attributable to the environment. First, there is some question about where keys should be cached. Since all of the Project Athena machines have local disks, the original code used `/tmp`. But this is highly insecure on diskless workstations, where `/tmp` exists on a file server; accordingly, a modification was made to store keys in shared memory. However, there is no guarantee that shared memory is not paged; if this entails network traffic, an intruder can capture these keys.

3. We regard this as a feature, not a bug.

4. More precisely, Kerberos is not a *host-to-host* protocol. In Version 5, it has been extended to support user-to-user authentication.^[Davi89b]

5. On Project Athena machines, remote access to most workstations is disabled by the *toehold* service.

Finally, the Kerberos protocol binds tickets to IP addresses. Such usage is problematic on multi-homed hosts (i.e., hosts with more than one IP address). Since workstations rarely have multiple addresses, this feature — intended to enhance security — was not a problem at MIT. Multi-user hosts often do have multiple addresses, however, and cannot live with this limitation. This problem has been fixed in Version 5.

4. PROTOCOL WEAKNESSES

4.1 Replay Attacks

The Kerberos protocol is not as resistant to penetration as it should be. A number of weaknesses are apparent; the most serious is its use of an authenticator to prevent replay attacks.

The authenticator relies on use of a timestamp to guard against reuse. This is problematic for several reasons. The claim is made that no replays are likely within the lifetime of the authenticator (typically five minutes). This is reinforced by the presence of the IP address in both the ticket and the authenticator. We are not persuaded by this logic. An intruder would not start by capturing a ticket and authenticator, and then develop the software to use them; rather, everything would be in place before the ticket-capture was attempted. Let us consider two examples.

Some years ago, Morris described an attack based on the slow increment rate of the initial sequence number counter in some TCP implementations.^[Mor85] He demonstrated that it was possible, under certain circumstances, to spoof one half of a preauthenticated TCP connection without ever seeing any responses from the targeted host. In a Kerberos environment, his attack would still work if accompanied by a stolen live authenticator, but not if a challenge/response protocol was used. Alternatively, an intruder may simply watch for a “mail-checking” session, wherein a user logs in briefly, reads a few messages, and logs out. A number of valuable tickets would be exposed by such a session, notably the one used to mount the user’s home directory. Note that the lifetime of the authenticators — 5 minutes — contributes considerably to this attack.

Further, the proposed Version 5 of Kerberos anticipates alternative communication protocols in which such replays may be trivial to implement. If Kerberos is to be considered as a general-purpose utility, it must make few security-critical assumptions about the underlying network, and those must be explicit.

It has been suggested that the proper defense is for the server to store all live authenticators; thus, an attempt to reuse one can be detected.^[Ste88] In fact, the original design of Kerberos required such caching, though this was never implemented.

For several reasons, we do not think that this solves the problem. First, on UNIX systems it is difficult for TCP-based^[Post81] servers to store authenticators. Servers generally operate by forking a separate process to handle each incoming request. The child processes do not share any memory with the parent process, and thus have no convenient way to inform it — and hence any other child servers — of the value of the authenticator used. There are a number of obvious solutions — pipes, authenticator servers, shared memory segments and the like — but all are awkward, and some even raise authentication questions of their own. To date, we know of no multi-threaded server implementation which caches authenticators.

UDP-based^[Post80] query servers can store the authenticators more easily, as a single process generally handles all incoming requests; however, they might have problems with legitimate retransmissions of the client’s request if the answer was lost. (UDP does not provide guaranteed delivery; thus, all retransmissions happen from application level, and are visible to the application.) Legitimate requests could be rejected, and a security alarm raised inappropriately. One possible solution would be for the application to generate a new authenticator when retransmitting a request; were it not for the other weaknesses of the authenticator scheme, this would be acceptable.

4.2 Secure Time Services

As noted, authenticators rely on machines’ clocks being roughly synchronized. If a host can be misled about the correct time, a stale authenticator can be replayed without any trouble at all. Since some time

synchronization protocols are unauthenticated,^[Post83, Mill88] and hosts are still using these protocols despite the existence of better ones,^[Mill89] such attacks are not difficult.

The design philosophy of building an authentication service on top of a secure time service is itself questionable. That is, it makes little sense to build an authentication system assuming an already-authenticated underlying system. Furthermore, while spoofing an unauthenticated time service may be a difficult programming task, it is not cryptographically difficult.⁶ As Kerberos is proposed for more varied environments, its dependence on a secure time service is more problematic and must be stressed.

As an alternative, we propose the use of a challenge/response authentication mechanism. As is done today, the client would present a ticket, though without an authenticator. The server would respond with a nonce identifier encrypted with the session key $K_{c,s}$; the client would respond with some function of that identifier, thereby proving that it possesses the session key.

Such an implementation is not without its costs, of course. An extra pair of messages must be exchanged each time a ticket is used, which rules out the possibility of authenticated datagrams. More seriously, all servers must then retain state to complete the authentication process. While not a problem for TCP-based servers, this may require substantial modification to UDP-based query servers. (This change is quite comparable in magnitude to the change needed to store live authenticators.)

There is a significant philosophical difference between the two techniques, however. In the current Kerberos implementation, with its assumptions about the network environment, retained state is only necessary to enhance security. The challenge/response scheme, on the other hand, guarantees security in a more general environment, but requires retained state to function at all.

Instead of substituting challenge/response throughout, a possible compromise is to extend the protocol with a challenge/response option. This option could be used, for example, to authenticate the user in the initial ticket-granting ticket exchange and to access a time service.⁷ Subsequent client-server interactions could use the current time-based protocol. But synchronizing the servers remains a problem; not synchronizing them will lead to denial of service, and if they access the time service as a client, they must somehow obtain and store a ticket and key to authenticate it. (See above on storing keys in servers.) Given these complexities and possible weaknesses, it would seem reasonable to allow any service to insist on the challenge/response option.

Summarizing, we emphasize that the security of Kerberos depends critically on synchronized clocks. In essence, the Kerberos protocols involve mutual trust among four parties: the client, server, authentication server and time server.

4.3 Password-Guessing Attacks

A second major class of attack on the Kerberos protocols involves an intruder recording login dialogs in order to mount a password-guessing assault. When a user requests $T_{c,tgs}$ (the ticket-granting ticket), the answer is returned encrypted with K_c , a key derived by a publicly-known algorithm from the user's password. A guess at the user's password can be confirmed by calculating K_c and using it to decrypt the recorded answer. An intruder who has recorded many such login dialogs has good odds of finding several new passwords; empirically, users do not pick good passwords unless forced to.^[Morr79, Gram84, Stol88]

We propose the use of exponential key exchange^[Diff76] to provide an additional layer of encryption. Without describing the algorithm in detail, it involves the two parties exchanging numbers that each can use to compute a secret key. An outsider, not knowing how the numbers were calculated, cannot easily

6. In some environments, programming is not even necessary. Low-powered fake WWV transmitters are not hard to build, and, if properly located, could easily block out the legitimate signal.

7. This was suggested to us by Clifford Neuman.

derive the key.

Such a use of exponential key exchange would prevent a passive wiretapper from accumulating the network equivalent of `/etc/passwd`. While exponential key exchange is normally vulnerable to active wiretaps, such attacks are comparatively rare, especially if dedicated network routers are used.

Using exponential key exchange has its costs. LaMacchia and Odlyzko^[LaMa] have demonstrated that exchanging small numbers is quite insecure, while using large ones is expensive in computation time. Additionally, we have added extra messages to the login dialog, and imposed the requirement for considerable extra state in the server. Given the trend towards hiding even encrypted passwords on UNIX systems, and given estimates that half of all logins at MIT are used within a two-week period, the investment may be justifiable. Perhaps the best solution is to support this feature as a domain-specific option.

Even exponential key exchange will not prevent all password-guessing attacks. Depending on how carefully the Kerberos logs are analyzed, an intruder need not even eavesdrop. Requests for tickets are not themselves encrypted; an attacker could simply request ticket-granting tickets for many different users. An enhancement to the server, to limit the rate of requests from a single source, may be useful.

Alternatively, some portion of the initial ticket request may be encrypted with K_c , providing a minimal authentication of the user to Kerberos, such that true eavesdropping would be required to mount this attack. (As we are preparing this manuscript, just such a suggestion is being hotly debated on the Kerberos mailing list. We originally overlooked an alternative avenue for mounting a password-guessing attack. Clients may be treated as services, and tickets to the client, encrypted by K_c , may be obtained by any user. This capability has been suggested as the basis for user-to-user authentication and enhanced mail services.^[Salt90] But any such scheme would require repeated re-entry of the user's password, an inconvenience we suspect will not be tolerated. We would prefer to provide the same functionality by having clients register separate instances as services, with truly random keys. Keys could be supplied to the client by the *keystore*, described below.)

An alternative approach is a protocol described by Lomas, Gong, Saltzer, and Needham.^[Loma89] They present a dialog with a server that does not expose the user to password-guessing attacks. However, their protocol relies on public-key cryptography, an approach explicitly rejected for Kerberos.

4.4 Spoofing Login

In a workstation environment, it is quite simple for an intruder to replace the `login` command with a version that records users' passwords before employing them in the Kerberos dialog. Such an attack negates one of Kerberos's primary advantages, that passwords are never transmitted in cleartext over a network. While this problem is not restricted to Kerberos environments, the Kerberos protocol makes it difficult to employ the standard countermeasure: one-time passwords.

A typical one-time password scheme employs a secret key shared between a server and some device in the user's possession. The server picks a random number and transmits it to the user. Both the server and the user (with the aid of the device) encrypt this number using the secret key; the result is transmitted back to the server. If the two computed values match, the user is assumed to possess the appropriate key.

Kerberos makes no provision for such a challenge/response dialog at login time. The server's response to the login request is always encrypted with K_c , a key derived from the user's password. Unless a "smart card" is employed that understands the entire Kerberos protocol, this precludes any use of one-time passwords.

An alternative (first suggested to us by T.H. Foregger) requires that the server pick a random number R , and use K_c to encrypt R . This value $\{R\}K_c$, rather than K_c , would be used to encrypt the server's response. R would be transmitted in the clear to the user. If a hand-held authenticator was in use, the user would employ it to calculate $\{R\}K_c$; otherwise, the login program would do it automatically.

Several objections may be raised to this scheme. First, hand-held authenticators are often thought to be inconvenient. This is true; however, they offer a substantial increase in security in high-threat environments. If they are not used, the cost of our scheme is quite low, simply one extra encryption on each end.

A second, more cogent, objection is that if the client's workstation cannot be trusted with a user's password, it cannot be trusted with session keys provided by Kerberos. This is, to some extent, a valid criticism, though we believe that compromise of the login password is much more serious than the capture of a few limited-lifetime session keys. This problem cannot be solved without the use of special-purpose hardware, a subject we shall return to below.

Finally, it has been pointed out that a user can always supply a known-clean boot device, or boot via the network. The former we regard as improbable in practice unless removable media are employed; the latter is insecure because the boot protocols are unauthenticated.

4.5 Inter-Session Chosen Plaintext Attacks

According to the description in the Version 5 draft,^[Koh189] servers using the KRB_PRIV format are susceptible to a *chosen plaintext attack*. (A chosen-plaintext attack is one where an attacker may choose all or part of the plaintext and, typically, use the resulting cipher text to attack the cipher. Here we use the cipher text to attack the protocol. Mail and file servers are examples of servers susceptible to such attacks.) Specifically, the encrypted portion of messages of this type have the form

$$X = (\text{DATA}, \text{timestamp} + \text{direction}, \text{hostaddress}, \text{PAD})$$

Since cipher-block chaining^[FIPS81, Davi89a] has the property that prefixes of encryptions are encryptions of prefixes, if *DATA* has the form

$$(\text{AUTHENTICATOR}, \text{CHECKSUM}, \text{REMAINDER})$$

then a prefix of the encryption of *X* with the session key is the encryption of $(\text{AUTHENTICATOR}, \text{CHECKSUM})$, and can be used to spoof an entire session with the server.

It may be argued that most servers are not susceptible to chosen plaintext attacks. Given that there are easy counters to this attack, it seems foolish to advocate a general format for private servers that does not also protect against it.

It should be noted that the simple attack above does not work against Kerberos Version 4, in which the encrypted portion of the KRB_PRIV message is of the form

$$(\text{length}(\text{DATA}), \text{DATA}, \text{msectime}, \text{hostaddress}, \text{timestamp} + \text{direction}, \text{PAD})$$

as the leading $\text{length}(\text{DATA})$ field disrupts the prefix-based attack. We leave it to the reader to discover a more complicated chosen ciphertext attack against this format, even allowing for the fact that Version 4 uses the nonstandard PCBC mode of encryption. (Hint: assume the initial vector is fixed and public.) However, it is interesting to note that the order of concatenation of message fields can have security-critical implications. We return to this question in the later section on message encoding.

4.6 Exposure of Session Keys

The term "session key" is a misnomer in the Kerberos protocol. This key is contained in the service ticket and is used in the multiple sessions between the client and server that use that ticket. Thus, it is more properly called a "multi-session key". Making this point explicit leads naturally to the suggestion that true session keys be negotiated as part of the Kerberos protocol. This limits the exposure to cryptanalysis^[Kahn67, Beke82, Deav85] of the multi-session key contained in the ticket, and precludes attacks which substitute messages from one session in another. (The chosen-plaintext attack of the previous section is one such example.) The session key could be generated by the server or could be computed as a session-specific function of the multi-session key.

4.7 The Scope of Tickets

Kerberos tickets are limited in both time and space. That is, tickets are usable only within the realm of the ticket-granting server, and only for a limited period of time. The first is necessary to the design of Kerberos; the TGS would not have any keys in common with servers in other realms. The latter is a security measure; the longer a ticket is in use, the greater the risk of it being stolen or compromised.

A further restriction on tickets, in Version 4, is that they cannot be forwarded. A user may obtain tickets at login time, and use these to log in to some other host; however, it is not possible to obtain authenticated network services from that host unless a new ticket-granting ticket is obtained. And that in turn would require transmission of a password across the network, in violation of fundamental principles of Kerberos's design.⁸

Version 5 incorporates provisions for ticket-forwarding; however, this introduces the problem of cascading trust. That is, a host *A* may be willing to trust credentials from host *B*, and *B* may be willing to trust host *C*, but *A* may not be willing to accept tickets originally created on host *C*, which *A* believes to be insecure. Kerberos has a flag bit to indicate that a ticket was forwarded, but does not include the original source.

A second problem with forwarding is that the concept only makes sense if tickets include the network address of the principal. If the address is omitted — as is permitted in Version 5 — a ticket may be used from any host, without any further modifications to the protocol. All that is necessary to employ such a ticket is a secure mechanism for copying the multi-session key to the new host. But that can be accomplished by an encrypted file transfer mechanism layered on top of existing facilities; it does not require flag bits in the Kerberos header.

Is it useful to include the network address in a ticket? We think not. Given our assumption that the network is under full control of the attacker, no extra security is gained by relying on the network address. In fact, the primary benefit of including it appears to be preventing immediate reuse of authenticators from a different host.

Even with the protection provided by network addresses, replay attacks that involve faked addresses are easy; again, see [Morr85]. Furthermore, an attacker can always wait until the connection is set up and authenticated, and then take it over, thus obviating any security provided by the presence of the address. Given these problems, and the cascading trust issue raised earlier, we suggest that ticket-forwarding be deleted.

A new inter-realm authentication mechanism is also introduced in Version 5. Briefly, if a user wishes to access a service in another realm, that user must first obtain a ticket-granting ticket for that realm. This is done by making the ticket-granting server in a realm the client of another realm's TGS. It in turn may be a client of yet another realm's TGS. A user's ticket request is signed by each TGS and passed along; realms will normally be configured in a hierarchical fashion, though "tandem links" are permitted.

Unfortunately, this scheme, while appearing to solve the problem, is deficient in several respects. First, and most serious, there is no discussion of how a TGS can determine which of its neighboring realms should be the next hop. Moving up the tree, towards the root, is an obvious answer for leaf nodes; however, each parent node would need complete knowledge of its entire subtree's realms in order to determine how to pass the request downwards. There are obvious analogies here to network-layer routing issues; note, though, that any "realm routing protocol" must include strong authentication provisions.

8. Actually, a special-purpose ticket-forwarder was built for Version 4. However, the implementation was of necessity awkward, and required participating hosts to run an additional server.

Another answer is to say that static tables should be used. This, too, has its security limitations: should realm administrators rely on electronic mail messages or telephone calls to set up their routing tables? If such calls are not authenticated, the security risks are obvious; if they are, the security of a Kerberos realm is subordinated to the security of a totally different authentication system.

There is also an evident link between inter-realm authentication and the cascading-trust problem. Kerberos Version 5 attempts to solve this by including path information in the ticket request. However, in the absence of a global name space, it is not clear that this is useful. If a realm is not a neighbor, its name may not carry any global significance, whether by malice or coincidence. Furthermore, to assess the validity of a request, a server needs global knowledge of the trustworthiness of all possible transit realms. In a large internet, such knowledge is probably not possible.

5. KERBEROS HARDWARE DESIGN CRITERIA

5.1 A Host Encryption Unit

One of the major reasons we question the suitability of Kerberos for multi-user hosts is the need for plaintext key storage. What if the host were equipped with an attached cryptographic unit? We consider the design parameters for such a box.

The primary goal is to perform cryptographic operations without exposing any keys to compromise. These operations must include validating tickets presented by remote users, creating requests for both ticket-granting tickets and application tickets, and encrypting and decrypting conversations. Consequently, there must be secure storage for an adequate number of keys, and the operating system must be able to select which key should be used for which function.

The next question, of course, is how keys are entered into the secure storage area. If tickets are decrypted by the encryption box but transferred to the host's memory for analysis, the embedded session key is exposed.⁹ Therefore, we conclude that the encryption box itself must understand the Kerberos protocols; nothing less will guarantee the security of the stored keys.

Entry of user keys is more problematic, since they must travel through the host. Unless user terminals are connected directly to the encryption unit, there is little choice. Storing them off the host, though, is a significant help, as the period of exposure is then minimized. Host-owned keys — service keys, or the keys that `root` would use to do NFS mounts — should be loaded via a Kerberos-authenticated service resident in the encryption unit. We shall return to this point below.

We must now ensure that the protocol itself does not provide a mechanism to obtain keys. Looking at the message definitions, we see that only session keys are ever sent, and these are always sent encrypted. Furthermore, user machines never generate any such messages; they merely forward them. Thus, the box need not have the ability to transmit a key, thereby providing us with a very high level of assurance that it will not do so.

If an encryption box is used for the Kerberos server itself, the problem is somewhat more complex. There are two places where keys are transmitted. First, when a ticket is granted, the ticket itself contains a session key, and a copy of that session key is sent back encrypted in the client's ticket-granting session key. Second, during the initial dialog with Kerberos, the ticket-granting session key must be sent out, encrypted in the client's password key. Note, though, that permanent keys are *never* sent; again, this assures us that the encryption box will not give away keys. Furthermore, since these session keys are intended to be random, we can buy ourselves a great deal of security by including a hardware random number generator on-board.

9. This is not a hypothetical concern. A program to do just that (for conventional passwords) was posted to *netnews* as long ago as 1984. It operated by reading `/dev/kmem`. The existence of this program was a principal factor motivating the current restrictive permission settings on `/dev/kmem`.

We are not too concerned about having to load client and server keys onto the board. This operation is done only by the Kerberos master server, for which strong physical security must be assumed in any event. It is possible that such an encryption unit can be made sufficiently tamper-resistant that even workstations can use them; certainly, there are commercial cryptographic devices that claim such strengths.

One major objection to this entire scheme is that ultimately, the encryption box is controlled by the host computer. Thus, if `root` is compromised, the host could instruct the box to create bogus tickets. Such concerns are certainly valid. However, as noted above, we consider such temporary breaches of security to be far less serious than the compromise of a key. Furthermore, using a separate unit allows us to create untamperable logs, etc.

It is also desirable to prevent misuse of keys. For example, we do not want the login key used to decrypt the arbitrary block of text that just happens to be the ticket-granting ticket. Accordingly, keys should be tagged with their purpose. A login key should be used only to decrypt the ticket-granting ticket; the key associated with it should be used only for obtaining service tickets, etc. Since the encryption box is performing all of the key management, this is not a difficult problem.

5.2 The Key Storage Unit

A variety of technologies may be used to implement encryption units, ranging from special boards to dedicated microcomputers connected to server hosts by physically-secure lines. If the latter is used, there is the temptation to use its disk storage to hold the service keys associated with the attached host, but we feel that that is inadvisable. Any media of that sort must be backed up, and the backups must be carefully guarded. Such a high degree of security may be impractical in some environments. Instead, we suggest that keys be kept in volatile memory, and downloaded from a secure *keystore* on request, via an encryption-protected channel. Thus, only one master key need be stored within the box; this key could either be in non-volatile storage, or be supplied by an operator when necessary.

More generally, the keystore is a secure, reliable repository for a limited amount of information. A client of the keystore could package arbitrary data to be retained by the keystore, and retrieved at a later date. This data — the service keys and tags, in the case of an encryption unit, or even a conventional Kerberos host — would be uninterpreted by the keystore. Storage and retrieval requests would be authenticated by Kerberos tickets, of course. Only encrypted transfer (KRB_PRIV) should be employed, as insurance against disclosure of such sensitive material.

As noted, the same keystore protocol could be used to supply additional keys for new instances of the same client. For example, a user *pat* could have a separate instance *pat.email*, for receiving encrypted electronic mail. The key for that instance would be restricted to that user, of course.

Generally, transactions with the keystore are initiated by the client. However, there is some question about how to create the additional user keys, as user workstations are not particularly good sources of random keys. The best alternative is to provide a (secure) random number service on the network. When a new client instance is added, this service would be consulted to generate the key; both Kerberos and the keystore would be told about the key.

6. SECURITY VALIDATION

Is Kerberos correct? By that we are asking if there are bugs (or trapdoors!) in the design or implementation of Kerberos, bugs that could be used to penetrate a system that relies on Kerberos. Some would say that by making the code widely available, the implementors have enabled would-be penetrators to gain a detailed knowledge of the system, thereby simplifying their task considerably. We reject that notion.

In the late nineteenth century, Kerckhoffs formulated the basic principal under which the security of cryptographic systems should be evaluated: all details of the system design should be assumed to be known by the adversary. Only cryptographic keys specifically assumed to be secret should be unavailable to an attacker.^[Kahn67, Kerc83] Given this basic premise, the security of a cryptographic system

is evaluated based on concerted efforts at cryptanalysis.

Kerberos is designed primarily as an authentication system incorporating a traditional cryptosystem (the Data Encryption Standard) as a component. Nevertheless, the philosophy guiding Kerckhoffs' evaluation criterion applies to the evaluation of the security of Kerberos. The details of Kerberos's design and implementation must be assumed known to a prospective attacker, who may also be in league with some subset of servers, clients, and (in the case of hierarchically-configured realms) some authentication servers. Kerberos is secure if and only if it can protect other clients and servers, beginning only with the premise that these client and server keys are secret, and that the encryption system is secure. Moreover, in the absence of a central, trusted "validation authority", each prospective user of Kerberos is responsible for judging its security. Of course, a public discussion of system security and publication of security evaluations will facilitate such judgements.

By describing the Kerberos design in publications and making the source code publically available, the Kerberos designers and implementors at Project Athena have made a commendable effort to encourage just such a public system validation. Obviously, this document is itself part of that process. However, the system design and its implementation have undergone significant modification, in part as a consequence of this public discussion. We stress that each modification to the design and implementation results in a new system whose security properties must be considered anew. (Examples of such modifications are the incorporation of hierarchically-organized servers and forwardable tickets in Version 5.)

Hence, on-going modification of Kerberos makes it a moving target for security validation attempts. A detailed security analysis would thus be premature. However, the proposed changes to Kerberos in the next few sections are intended, not so much to defeat specific attacks, as to facilitate the validation process. In particular, these suggestions are intended to make Kerberos more modular, in design and implementation. Doing so should make the security consequences of modifications more apparent, and facilitate an incremental approach to Kerberos security validation.

6.1 Message Encoding and Cut-and-Paste Attacks

The most simple analysis of the security of the Kerberos protocols should check that there is no possibility of ambiguity between messages sent in different contexts. That is, a ticket should never be interpretable as an authenticator, or vice versa. Such an analysis depends on redundancy in the pre-encryption binary encodings of each of the ticket and authenticator information. Currently, that analysis must be repeated with every modification to the protocol. This repetitive and often intricate analysis would be unnecessary if standard encodings (such as ASN.1^[ASN1, BER]) were used. These encodings should include the overall message type (such as KRB_TGS_REP or KRB_PRIV). Together with reasonable assumptions about the encryption layer (see the next section), such an encoding scheme would greatly simplify the protocol validation process, particularly as the protocol is modified or extended.

The use of ASN.1 encodings has been adopted for other reasons in Version 5. We reinforce here that there are design principles other than standards compatibility that motivate such a change.

6.2 The Encryption Layer

Version 4 of Kerberos uses the nonstandard PCBC mode of encryption, *propagating cipher block chaining*, in which plaintext block $i + 1$ is exclusive-or'ed with both the plaintext and ciphertext of block i before encryption. This mode was observed to have poor propagation properties that permit message-stream modification: specifically, if two blocks of ciphertext are interchanged, only the corresponding blocks are garbled on decryption. Version 5 replaces PCBC mode with the standard CBC mode, *cipher block chaining*, which exclusive-or's just the ciphertext of block i with the plaintext of block $i + 1$ before encryption. A checksum — as of Draft 2, the exact form had not been determined — is used to detect message modification. In order to ensure that duplicate messages have different encryptions, random initial "confounders" are added to some message formats. In addition, Version 5 supports alternative encryption algorithms as options.

Both the confounder and checksum mechanisms are meant to augment the security of CBC encryption. They belong in a separate encryption layer, not at the level of the Kerberos protocols themselves. Further, the confounder mechanism should be replaced by using the standard initial vector mechanism of cipher-block chaining.^[FIPS81, Davi89a]

To prevent message-stream modification during authenticated or private sessions, Version 5 uses a timestamp field to prevent entire encrypted messages from being replayed. This is another concern more properly delegated to the encryption layer, where chaining across the packets of the entire session is the more standard mechanism. (Such chaining avoids both the dependence on a clock and the need to cache recent timestamps.)

Separating the Kerberos protocols from the details of encryption would facilitate both validation of the security of the Kerberos protocols, and implementations and validations involving alternative cryptosystems. Too much focus on mechanism, while endemic to cryptographic protocol design, leads away from the need to state the basic properties required of the encryption layer. We would suggest the following adversarial analysis as the starting point for such a specification: allow an adversary to submit, one after the other, any number of messages for encryption under an unknown key K . The adversary also has the ability to take prefixes and suffixes of known messages, exclusive-or known messages, and encrypt or decrypt with known keys. At the end of this process, the adversary should not be able to produce any encrypted messages other than those specifically submitted for encryption. Such an analysis would preclude encryption schemes susceptible to simple chosen-plaintext attacks, as described in a previous section.

Given the intractability of reasoning about DES, or of proving complexity properties of any cryptosystem with bounded key size, such analyses will be no guarantee of overall security. But they can be used to preclude the existence of trivial cut-and-paste attacks.^[DeMi83, Moor88]

7. RECOMMENDED CHANGES TO THE KERBEROS PROTOCOL

Below, we list our recommended changes to the Kerberos protocol. Our ranking is governed by our estimate of the likelihood and consequences of the attack, balanced against the difficulty of implementing the modification.

- a. A challenge/response protocol should be offered as an optional alternative to time-based authentication.
- b. Use a standard message encoding, such as ASN.1, which includes identification of the message type within the encrypted data.
- c. Alter the basic login protocol to allow for handheld authenticators, in which $\{R\}K_C$, for a random R , is used to encrypt the server's reply to the user, in place of the key K_C obtained from the user password. This allows the login procedure to prompt the user with R , who obtains $\{R\}K_C$ from the handheld device and returns that value instead of the password itself.
- d. Mechanisms such as random initial vectors (in place of confounders), block chaining and message authentication codes should be left to a separate encryption layer, whose information-hiding requirements are clearly explicated. Specific mechanisms based on DES should be validated and implemented.
- e. The client/server protocol should be modified so that the multi-session key is used to negotiate a true session key, which is then used to protect the remainder of the session.
- f. Support for special-purpose hardware should be added, such as the keystore. More importantly, future enhancements to the Kerberos protocol should be designed under the assumption that a host is using encryption and key-storage hardware.
- g. To protect against trivial password-guessing attacks, the protocol should not distribute tickets for users (encrypted with the password-based key), and the initial exchange should authenticate the user to the Kerberos server.

- h. Support for optional extensions should be included. In particular, an option to protect against password-guessing attacks via eavesdropping may be a desirable feature.

8. ACKNOWLEDGEMENTS

We would like to thank T.H. Foregger and C. Neuman for their comments on a draft of this manuscript.

REFERENCES

- [FIPS81] “DES Modes of Operation,” Federal Information Processing Standards Publication 81 (December 1980). National Bureau of Standards, U.S. Department of Commerce
- [ASN1] “Information Processing Systems — Open Systems Interconnection — Specification of Abstract Syntax Notation One (ASN.1),” International Standard 8824 (1987). International Organization for Standardization and International Electrotechnical Committee
- [BER] “Information Processing Systems — Open Systems Interconnection — Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1),” International Standard 8825 (1987). International Organization for Standardization and International Electrotechnical Committee
- [Beke82] H. Beker and F. Piper, *Cipher Systems*, John Wiley & Sons (1982).
- [Brya88] B. Bryant, *Designing an Authentication System: A Dialogue in Four Scenes*, Draft, February 8, 1988.
- [Davi89a] D.W. Davies and W.L. Price, *Security for Computer Networks*, John Wiley & Sons (1989). Second Edition
- [Davi89b] D. Davis and R. Swick, *Workstation Services and Kerberos Authentication at Project Athena*, MIT Project Athena (March 3, 1989).
- [Deav85] C.A. Deavours and L. Kruh, *Machine Cryptography and Modern Cryptanalysis*, Artech House (1985).
- [DeMi83] R. DeMillo and M. Merritt, “Protocols for Data Security,” *Computer* **16**(2), pp. 39-50 (February 1983).
- [Diff76] W. Diffie and M.E. Hellman, “New Directions in Cryptography,” *IEEE Transactions on Information Theory* **6**, pp. 644-654 (November, 1976).
- [Gram84] F.T. Grampp and R.H. Morris, “UNIX Operating System Security,” *AT&T Bell Laboratories Technical Journal* **63**(8, Part 2), pp. 1649-1672, A&T (October, 1984).
- [Kahn67] D. Kahn, *Codebreakers: The Story of Secret Writing*, Macmillan (1967).
- [Kerc83] A. Kerckhoffs, *La Cryptographie Militaire*, Libraire Militaire de L. Baudoin & Cie., Paris (1883).
- [Kohl89] J. Kohl, B. Clifford Neuman, and J. Steiner, *The Kerberos Network Authentication Service*, MIT Project Athena (November 6, 1989). Version 5, Draft 2
- [LaMa] B.A. LaMacchia and A.M. Odlyzko, *Computation of Discrete Logarithms in Prime Fields*, (Manuscript in preparation).
- [Loma89] T.M.A. Lomas, L. Gong, J.H. Saltzer, and R.M. Needham, “Reducing Risks from Poorly Chosen Keys,” *Operating Systems Review* **23**(5), pp. 14-18, ACM (December 1989).
- [Mill87] S.P. Miller, B.C. Neuman, J.I. Schiller, and J.H. Saltzer, “Kerberos Authentication and Authorization System,” in *Project Athena Technical Plan* (December 1987). Section

E.2.1

- [Mill88] D.L. Mills, “Network Time Protocol,” RFC 1059 (July 1988).
- [Mill89] D.L. Mills, “Network Time Protocol,” RFC 1119 (September 1989).
- [Moor88] J.H. Moore, “Protocol Failures in Cryptosystems,” *Proc. IEEE* **76**(5), pp. 594-602 (May 1988).
- [Morr79] R. Morris and K. Thompson., “UNIX Password Security,” *Communications of the ACM* **22**(11), p. 594 (November 1979).
- [Morr85] R.T. Morris, “A Weakness in the 4.2BSD UNIX TCP/IP Software,” Computing Science Technical Report No. 117, AT&T Bell Laboratories, Murray Hill, New Jersey (February 1985).
- [Post80] J.B. Postel, “User Datagram Protocol.,” RFC 768 (August 28, 1980).
- [Post81] J.B. Postel, “Transmission Control Protocol.,” RFC 793 (September 1981).
- [Post83] J.B. Postel and K. Harrenstien, “Time Protocol.,” RFC 868 (May 1983).
- [Salt90] J.H. Saltzer, private communication, June 19, 1990.
- [Steir88] J. Steiner, C. Neuman, and J.I. Schiller, “Kerberos: An Authentication Service for Open Network Systems,” in *Proc. Winter USENIX Conference*, Dallas (1988).
- [Stoll88] C. Stoll, “Stalking the Wiley Hacker,” *Communications of the ACM* **31**(5), p. 484 (May 1988).