Security Problems in the TCP/IP Protocol Suite

*S.M. Bellovin**
smb@ulysses.att.com

AT&T Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

The TCP/IP protocol suite, which is very widely used today, was developed under the sponsorship of the Department of Defense. Despite that, there are a number of serious security flaws inherent in the protocols, regardless of the correctness of any implementations. We describe a variety of attacks based on these flaws, including sequence number spoofing, routing attacks, source address spoofing, and authentication attacks. We also present defenses against these attacks, and conclude with a discussion of broad-spectrum defenses such as encryption.

1. INTRODUCTION

The TCP/IP protocol suite[1][2], which is very widely used today, was developed under the sponsorship of the Department of Defense. Despite that, there are a number of serious security flaws inherent in the protocols. Some of these flaws exist because hosts rely on IP source address for authentication; the Berkeley "*r*-utilities"[3] are a notable example. Others exist because network control mechanisms, and in particular routing protocols, have minimal or non-existent authentication.

When describing such attacks, our basic assumption is that the attacker has more or less complete control over some machine connected to the Internet. This may be due to flaws in that machine's own protection mechanisms, or it may be because that machine is a microcomputer, and inherently unprotected. Indeed, the attacker may even be a rogue system administrator.

1.1 Exclusions

We are not concerned with flaws in particular implementations of the protocols, such as those used by the Internet "worm"[4][5][6]. Rather, we discuss generic problems with the protocols themselves. As will be seen, careful implementation techniques can alleviate or prevent some of these problems. Some of the protocols we discuss are derived from Berkeley's version of the UNIX® system; others are generic Internet protocols.

We are also not concerned with classic network attacks, such as physical eavesdropping, or altered or injected messages. We discuss such problems only in so far as they are facilitated or possible because of protocol problems.

For the most part, there is no discussion here of vendor-specific protocols. We do discuss some problems with Berkeley's protocols, since these have become de facto standards for many vendors, and not just for UNIX systems.

2. TCP SEQUENCE NUMBER PREDICTION

One of the more fascinating security holes was first described by Morris[7]. Briefly, he used TCP sequence number prediction to construct a TCP packet sequence without ever receiving any responses from the server. This allowed him to spoof a trusted host on a local network.

---

* Author's address: Room 3D-558, AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, New Jersey 07974.

The normal TCP connection establishment sequence involves a 3-way handshake. The client selects and transmits an initial sequence number ISN_C, the server acknowledges it and sends its own sequence number ISN_S, and the client acknowledges that. Following those three messages, data transmission may take place. The exchange may be shown schematically as follows:

$C \rightarrow S: SYN\,(ISN_C)$
$S \rightarrow C: SYN\,(ISN_S), ACK\,(ISN_C)$
$C \rightarrow S: ACK\,(ISN_S)$
$C \rightarrow S: data$
and/or
$S \rightarrow C: data$

That is, for a conversation to take place, C must first hear ISN_S, a more or less random number.

Suppose, though, that there was a way for an intruder X to *predict* ISN_S. In that case, it could send the following sequence to impersonate trusted host T:

$X \rightarrow S: SYN\,(ISN_X), SRC=T$
$S \rightarrow T: SYN\,(ISN_S), ACK\,(ISN_X)$
$X \rightarrow S: ACK\,(ISN_S), SRC=T$
$X \rightarrow S: ACK\,(ISN_S), SRC=T, nasty\text{-}data$

Even though the message $S \rightarrow T$ does not go to X, X was able to know its contents, and hence could send data. If X were to perform this attack on a connection that allows command execution (i.e., the Berkeley *rsh* server), malicious commands could be executed.

How, then, to predict the random *ISN*? In Berkeley systems, the initial sequence number variable is incremented by a constant amount once per second, and by half that amount each time a connection is initiated. Thus, if one initiates a legitimate connection and observes the ISN_S used, one can calculate, with a high degree of confidence, ISN_S' used on the next connection attempt.

Morris points out that the reply message

$S \rightarrow T: SYN\,(ISN_S), ACK\,(ISN_X)$

does not in fact vanish down a black hole; rather, the real host T will receive it and attempt to reset the connection. This is not a serious obstacle. Morris found that by impersonating a server port on T, and by flooding that port with apparent connection requests, he could generate queue overflows that would make it likely that the $S \rightarrow T$ message would be lost. Alternatively, one could wait until T was down for routine maintenance or a reboot.

A variant on this TCP sequence number attack, not described by Morris, exploits the *netstat*[8] service. In this attack, the intruder impersonates a host that is down. If *netstat* is available on the target host, it may supply the necessary sequence number information on another port; this eliminates all need to guess[1].

Defenses

Obviously, the key to this attack is the relatively coarse rate of change of the initial sequence number variable on Berkeley systems. The TCP specification requires that this variable be incremented approximately 250,000 times per second; Berkeley is using a much slower rate. However, the critical factor is the granularity, not the average rate. The change from an increment of 128 per second in 4.2BSD to 125,000 per second in 4.3BSD is meaningless, even though the latter is within a factor of two of the specified rate.

1. The *netstat* protocol is obsolete, but is still present on some Internet hosts. Security concerns were not behind its elimination.

Let us consider whether a counter that operated at a true 250,000 hz rate would help. For simplicity's sake, we will ignore the problem of other connections occurring, and only consider the fixed rate of change of this counter.

To learn a current sequence number, one must send a SYN packet, and receive a response, as follows:

$$X \rightarrow S: SYN(ISN_X)$$
$$S \rightarrow X: SYN(ISN_S), ACK(ISN_X) \quad (1)$$

The first spoof packet, which triggers generation of the next sequence number, can immediately follow the server's response to the probe packet:

$$X \rightarrow S: SYN(ISN_X), SRC=T \quad (2)$$

The sequence number ISN_S used in the response

$$S \rightarrow T: SYN(ISN_S), ACK(ISN_X)$$

is uniquely determined by the time between the origination of message (1) and the receipt at the server of message (2). But this number is precisely the round-trip time between X and S. Thus, if the spoofer can accurately measure (and predict) that time, even a 4 μ-second clock will not defeat this attack.

How accurately can the trip time be measured? If we assume that stability is good, we can probably bound it within 10 milliseconds or so. Clearly, the Internet does not exhibit such stability over the long-term[9], but it is often good enough over the short term.[2] There is thus an uncertainty of 2500 in the possible value for ISN_S. If each trial takes 5 seconds, to allow time to re-measure the round-trip time, an intruder would have a reasonable likelihood of succeeding in 7500 seconds, and a near-certainty within a day. More predictable (i.e., higher quality) networks, or more accurate measurements, would improve the odds even further in the intruder's favor. Clearly, simply following the letter of the TCP specification is not good enough.

We have thus far tacitly assumed that no processing takes places on the target host. In fact, some processing does take place when a new request comes in; the amount of variability in this processing is critical. On a 6 MIPS machine, one tick — 4 μ-seconds — is about 25 instructions. There is thus considerable sensitivity to the exact instruction path followed. High-priority interrupts, or a slightly different TCB allocation sequence, will have a comparatively large effect on the actual value of the next sequence number. This randomizing effect is of considerable advantage to the target. It should be noted, though, that faster machines are *more* vulnerable to this attack, since the variability of the instruction path will take less real time, and hence affect the increment less. And of course, CPU speeds are increasing rapidly.

This suggests another solution to sequence number attacks: randomizing the increment. Care must be taken to use sufficient bits; if, say, only the low-order 8 bits were picked randomly, and the granularity of the increment was coarse, the intruder's work factor is only multiplied by 256. A combination of a fine-granularity increment and a small random number generator, or just a 32-bit generator, is better. Note, though, that many pseudo-random number generators are easily invertible[10]. In fact, given that most such generators work via feedback of their output, the enemy could simply compute the next "random" number to be picked. Some hybrid techniques have promise — using a 32-bit generator, for example, but only emitting 16 bits of it — but brute-force attacks could succeed at determining the seed. One would need at least 16 bits of random data in each increment, and perhaps more, to defeat probes from the network, but that might leave too few bits to guard against a search for the seed. More research or simulations are needed to determine the proper parameters.

---

2. At the moment, the Internet may not have such stability even over the short-term, especially on long-haul connections. It is not comforting to know that the security of a network relies on its low quality of service.

Rather than go to such lengths, it is simpler to use a cryptographic algorithm (or device) for ISN_S generation. The Data Encryption Standard[11] (DES) in *electronic codebook mode*[12] is an attractive choice as the ISN_S source, with a simple counter as input. Alternatively, DES could be used in *output feedback mode* without an additional counter. Either way, great care must be taken to select the key used. The time-of-day at boot time is not adequate; sufficiently good information about reboot times is often available to an intruder, thereby permitting a brute-force attack. If, however, the reboot time is encrypted with a per-host secret key, the generator cannot be cracked with any reasonable effort.

Performance of the initial sequence number generator is not a problem. New sequence numbers are needed only once per connection, and even a software implementation of DES will suffice. Encryption times of 2.3 milliseconds on a 1 MIPS processor have been reported[13].

An additional defense involves good logging and alerting mechanisms. Measurements of the round-trip time — essential for attacking RFC-compliant hosts — would most likely be carried out using ICMP *Ping* messages; a "transponder" function could log excessive ping requests. Other, perhaps more applicable, timing measurement techniques would involve attempted TCP connections; these connections are conspicuously short-lived, and may not even complete *SYN* processing. Similarly, spoofing an active host will eventually generate unusual types of *RST* packets; these should not occur often, and should be logged.

3. THE JOY OF ROUTING

Abuse of the routing mechanisms and protocols is probably the simplest protocol-based attack available. There are a variety of ways to do this, depending on the exact routing protocols used. Some of these attacks succeed only if the remote host does source address-based authentication; others can be used for more powerful attacks.

A number of the attacks described below can also be used to accomplish denial of service by confusing the routing tables on a host or gateway. The details are straight-forward corollaries of the penetration mechanisms, and will not be described further.

3.1 Source Routing

If available, the easiest mechanism to abuse is IP source routing. Assume that the target host uses the reverse of the source route provided in a TCP open request for return traffic. Such behavior is utterly reasonable; if the originator of the connection wishes to specify a particular path for some reason — say, because the automatic route is dead — replies may not reach the originator if a different path is followed.

The attacker can then pick any IP source address desired, including that of a trusted machine on the target's local network. Any facilities available to such machines become available to the attacker.

Defenses

It is rather hard to defend against this sort of attack. The best idea would be for the gateways into the local net to reject external packets that claim to be from the local net. This is less practical than it might seem since some Ethernet[3] network adapters receive their own transmissions, and this feature is relied upon by some higher-level protocols. Furthermore, this solution fails completely if an organization has two trusted networks connected via a multi-organization backbone. Other users on the backbone may not be trustable to the same extent that local users are presumed to be, or perhaps their vulnerability to outside attack is higher. Arguably, such topologies should be avoided in any event.

A simpler method might be to reject pre-authorized connections if source routing information was present. This presumes that there are few legitimate reasons for using this IP option, especially for

3. Ethernet is a registered trademark of Xerox Corporation.

relatively normal operations. A variation on this defense would be to analyze the source route and accept it if only trusted gateways were listed; that way, the final gateway could be counted on to deliver the packet only to the true destination host. The complexity of this idea is probably not worthwhile.

Some protocols (i.e., Berkeley's *rlogin* and *rsh*) permit ordinary users to extend trust to remote host/user combinations. In that case, individual users, rather than an entire system, may be targeted by source routing attacks.[4] Suspicious gateways[14] will not help here, as the host being spoofed may not be within the security domain protected by the gateways.

3.2 Routing Information Protocol Attacks

The *Routing Information Protocol*[15] (RIP) is used to propagate routing information on local networks, especially broadcast media. Typically, the information received is unchecked. This allows an intruder to send bogus routing information to a target host, and to each of the gateways along the way, to impersonate a particular host. The most likely attack of this sort would be to claim a route to a particular unused host, rather than to a network; this would cause all packets destined for that host to be sent to the intruder's machine. (Diverting packets for an entire network might be too noticeable; impersonating an idle work-station is comparatively risk-free.) Once this is done, protocols that rely on address-based authentication are effectively compromised.

This attack can yield more subtle, and more serious, benefits to the attacker as well. Assume that the attacker claims a route to an active host or workstation instead. All packets for that host will be routed to the intruder's machine for inspection and possible alteration. They are then resent, using IP source address routing, to the intended destination. An outsider may thus capture passwords and other sensitive data. This mode of attack is unique in that it affects outbound calls as well; thus, a user calling out from the targeted host can be tricked into divulging a password. Most of the earlier attacks discussed are used to forge a source address; this one is focused on the destination address.

Defenses

A RIP attack is somewhat easier to defend against than the source-routing attacks, though some defenses are similar. A paranoid gateway — one that filters packets based on source or destination address — will block any form of host-spoofing (including TCP sequence number attacks), since the offending packets can never make it through. But there are other ways to deal with RIP problems.

One defense is for RIP to be more skeptical about the routes it accepts. In most environments, there is no good reason to accept new routes to your own local networks. A router that makes this check can easily detect intrusion attempts. Unfortunately, some implementations rely on hearing their own broadcasts to retain their knowledge of directly-attached networks. The idea, presumably, is that they can use other networks to route around local outages. While fault-tolerance is in general a good idea, the actual utility of this technique is low in many environments compared with the risks.

It would be useful to be able to authenticate RIP packets; in the absence of inexpensive public-key signature schemes, this is difficult for a broadcast protocol. Even if it were done, its utility is limited; a receiver can only authenticate the immediate sender, which in turn may have been deceived by gateways further upstream.

Even if the local routers don't implement defense mechanisms, RIP attacks carry another risk: the bogus routing entries are visible over a wide area. Any router (as opposed to host) that receives such data will rebroadcast it; a suspicious administrator almost anywhere on the local collection of networks could notice the anomaly. Good log generation would help, but it is hard to distinguish a genuine intrusion from the routing instability that can accompany a gateway crash.

---

4. Permitting ordinary users to extend trust is probably wrong in any event, regardless of abuse of the protocols. But such concerns are beyond the scope of this paper.

3.3 Exterior Gateway Protocol

The *Exterior Gateway Protocol* (EGP)[16] is intended for communications between the core gateways and so-called *exterior gateways*. An exterior gateway, after going through a *neighbor acquisition* protocol, is periodically polled by the core; it responds with information about the networks it serves. These networks must all be part of its *autonomous system*. Similarly, the gateway periodically requests routing information from the core gateway. Data is not normally sent except in response to a poll; furthermore, since each poll carries a sequence number that must be echoed by the response, it is rather difficult for an intruder to inject a false route update. Exterior gateways are allowed to send exactly one spontaneous update between any two polls; this, too, must carry the sequence number of the last poll received. It is thus comparatively difficult to interfere in an on-going EGP conversation.

One possible attack would be to impersonate a second exterior gateway for the same autonomous system. This may not succeed, as the core gateways could be equipped with a list of legitimate gateways to each autonomous system. Such checks are not currently done, however. Even if they were, they could be authenticated only by source IP address.

A more powerful attack would be to claim reachability for some network where the real gateway is down. That is, if gateway G normally handles traffic for network N, and G is down, gateway G' could advertise a route to that network. This would allow password capture by assorted mechanisms. The main defense against this attack is topological (and quite restrictive): exterior gateways must be on the same network as the core; thus, the intruder would need to subvert not just any host, but an existing gateway or host that is directly on the main net.

A sequence number attack, similar to those used against TCP, might be attempted; the difficulty here is in predicting what numbers the core gateway is using. In TCP, one can establish arbitrary connections to probe for information; in EGP, only a few hosts may speak to the core. (More accurately, the core could only speak to a few particular hosts, though as noted such checks are not currently implemented.) It may thus be hard to get the raw data needed for such an attack.

3.4 The Internet Control Message Protocol

The *Internet Control Message Protocol* (ICMP)[17] is the basic network management tool of the TCP/IP protocol suite. It would seem to carry a rich potential for abuse. Surprisingly, ICMP attacks are rather difficult; still, there are often holes that may be exploited.

The first, and most obvious target, is the ICMP *Redirect* message; it is used by gateways to advise hosts of better routes. As such it can often be abused in the same way that RIP can be. The complication is that a Redirect message must be tied to a particular, existing connection; it cannot be used to make an unsolicited change to the host's routing tables. Furthermore, Redirects are only applicable within a limited topology; they may be sent only from the first gateway along the path to the originating host. A later gateway may not advise that host, nor may it use ICMP Redirect to control other gateways.

Suppose, though, that an intruder has penetrated a secondary gateway available to a target host, but not the primary one. (It may suffice to penetrate an ordinary host on the target's local network, and have it claim to be a gateway.) Assume further that the intruder wishes to set up a false route to trusted host T through that compromised secondary gateway. The following sequence may then be followed. Send a false TCP open packet to the target host, claiming to be from T. The target will respond with its own open packet, routing it through the secure primary gateway. While this is in transit, a false Redirect may be sent, claiming to be from the primary gateway, and referring to the bogus connection. This packet will appear to be a legitimate control message; hence the routing change it contains will be accepted. If the target host makes this change to its global routing tables, rather than just to the per-connection cached route, the intruder may proceed with spoofing host T.

Some hosts do not perform enough validity checks on ICMP Redirect messages; in such cases, the impact of this attack becomes similar to RIP-based attacks.

ICMP may also be used for targeted denial of service attacks. Several of its messages, such as *Destination Unreachable* and *Time to Live Exceeded*, may be used to reset existing connections. If the

intruder knows the local and remote port numbers of a TCP connection, an ICMP packet aimed at that connection may be forged[5]. Such information is sometimes available through the *netstat* service.

A more global denial of service attack can be launched by sending a fraudulent *Subnet Mask Reply* message. Some hosts will accept any such message, whether they have sent a query or not; a false one could effectively block all communications with the target host.

Defenses

Most ICMP attacks are easy to defend against with just a modicum of paranoia. If a host is careful about checking that a message really does refer to a particular connection, most such attacks will not succeed. In the case of TCP, this includes verifying that the ICMP packet contains a plausible sequence number in the returned-packet portion. These checks are less applicable to UDP, though.

A defense against Redirect attacks merits additional attention, since such attacks can be more serious. Probably, the best option is to restrict route changes to the specified connection; the global routing table should not be modified in response to ICMP Redirect messages[6].

Finally, it is worth considering whether ICMP Redirects are even useful in today's environment. They are only usable on local networks with more than one gateway to the outside world. But it is comparatively easy to maintain complete and correct local routing information. Redirect messages would be most useful from the core gateways to local exterior gateways, as that would allow such local gateways to have less than complete knowledge of the Internet; this use is disallowed, however.

Subnet Mask attacks can be blocked if the Reply packet is honored only at the appropriate time. In general, a host wants to see such a message only at boot time, and only if it had issued a query; a stale reply, or an unsolicited reply, should be rejected out of hand. There is little defense against a forged reply to a genuine Subnet Mask query, as a host that has sent such a query typically has few resources with which to validate the response. If the genuine response is not blocked by the intruder, though, the target will receive multiple replies; a check to ensure that all replies agree would guard against administrative errors as well.

4. THE "AUTHENTICATION" SERVER

As an alternative to address-based authentication, some implementations use the *Authentication Server*[18]. A server that wishes to know the identity of its client may contact the client host's Authentication Server[7], and ask it for information about the user owning a particular connection. This method is inherently more secure than simple address-based authentication, as it uses a second TCP connection not under control of the attacker. It thus can defeat sequence number attacks and source routing attacks. There are certain risks, however.

The first, and most obvious, is that not all hosts are competent to run authentication servers. If the client host is not secure, it does not matter who the user is claimed to be; the answer cannot be trusted. Second, the authentication message itself can be compromised by routing table attacks. If RIP has been used to alter the target's idea of how to reach some host, the authentication query will rely on the same altered routing data. Finally, if the target host is down, a variant on the TCP sequence number attack may be used; after the server sends out a TCP open request to the presumed authentication server, the attacker can complete the open sequence and send a false reply. If the target runs a *netstat* server, this is even easier; as noted, *netstat* will often supply the necessary sequence numbers with no need to guess.

---

5. In fact, such programs are available today; they are used as administrative tools to reset hung TCP connections.

6. This has other benefits as well, especially in environments where ICMP-initiated route changes are not timed out. The author has seen situations where RIP instability following a gateway crash has led to erroneous ICMP Redirect messages. These had the effect of permanently corrupting the routing tables on other hosts.

7. The Internet Activities Board does not currently recommend the Authentication Server for implementation[19]. However, the decision was not made because of security problems[5].

A less-obvious risk is that a fake authentication server can always reply "no". This constitutes a denial of service attack.

Defenses

A server that wishes to rely on another host's idea of a user should use a more secure means of validation, such as the Needham-Schroeder algorithm[20][21][22]. TCP by itself is inadequate.

5. HERE BE DRAGONS

Some protocols, while not inherently flawed, are nevertheless susceptible to abuse. A wise implementor would do well to take these problems into account when providing the service.

5.1 The "Finger" Service

Many systems implement a *finger* service [23]. This server will display useful information about users, such as their full names, phone numbers, office numbers, etc. Unfortunately, such data provides useful grist for the mill of a password cracker.[24] By running such a service, a system administrator is giving away this data.

5.2 Electronic Mail

Electronic mail is probably the most valuable service on the Internet. Nevertheless, it is quite vulnerable to misuse. As normally implemented[25][26], the mail server provides no authentication mechanisms. This leaves the door wide open to faked messages. RFC 822 does support an *Encrypted* header line, but this is not widely used. (However, see RFC 1040[27] for a discussion of a proposed new encryption standard for electronic mail.)

5.2.1 The Post Office Protocol

The *The Post Office Protocol* (POP)[28] allows a remote user to retrieve mail stored on a central server machine. Authentication is by means of a single command containing both the user name and the password. However, combining the two on a single command mandates the use of conventional passwords. And such passwords are becoming less popular; they are too vulnerable to wire-tappers, intentional or accidental disclosure, etc.

As an alternative, many sites are adopting "one-time passwords"[8]. With one-time passwords, the host and some device available to the user share a cryptographic key. The host issues a random challenge; both sides encrypt this number, and the user transmits it back to the host. Since the challenge is random, the reply is unique to that session, thereby defeating eavesdroppers. And since the user does not know the key — it is irretrievably stored in the device — the password cannot be given away without depriving the user of the ability to log in.

The newest version of POP[30] has split the user name and password into two commands, which is useful. However, it also defines an optional mechanism for preauthenticated connections, typically using Berkeley's mechanisms. Commendably, the security risks of this variant are mentioned explicitly in the document.

5.2.2 PCMAIL

The *PCMAIL* protocol[31] uses authentication mechanisms similar to those in POP2. In one major respect, PCMAIL is more dangerous: it supports a password-change command. This request requires that both the old and new passwords be transmitted unencrypted.

---

8. One-time passwords were apparently first used for military IFF (Identification Friend or Foe) systems[29].

5.3 The Domain Name System

The *Domain Name System* (DNS)[32][33] provides for a distributed database mapping host names to IP addresses. An intruder who interferes with the proper operation of the DNS can mount a variety of attacks, including denial of service and password collection. There are a number of vulnerabilities.

In some resolver implementations, it is possible to mount a sequence number attack against a particular user. When the target user attempts to connect to a remote machine, an attacker can generate a domain server response to the target's query. This requires knowing both the UDP port used by the client's resolver and the DNS sequence number used for the query. The latter is often quite easy to obtain, though, since some resolvers always start their sequence numbers with 0. And the former may be obtainable via *netstat* or some analogous host command.

A combined attack on the domain system and the routing mechanisms can be catastrophic. The intruder can intercept virtually all requests to translate names to IP addresses, and supply the address of a subverted machine instead; this would allow the intruder to spy on all traffic, and build a nice collection of passwords if desired.

For this reason, domain servers are high-value targets; a sufficiently determined attacker might find it useful to take over a server by other means, including subverting the machine one is on, or even physically interfering with its link to the Internet. There is no network defense against the former, which suggests that domain servers should only run on highly secure machines; the latter issue may be addressed by using authentication techniques on domain server responses.

The DNS, even when functioning correctly, can be used for some types of spying. The normal mode of operation of the DNS is to make specific queries, and receive specific responses. However, a *zone transfer* (AXFR) request exists that can be used to download an entire section of the database; by applying this recursively, a complete map of the name space can be produced. Such a database represents a potential security risk; if, for example, an intruder knows that a particular brand of host or operating system has a particular vulnerability, that database can be consulted to find all such targets. Other uses for such a database include espionage; the number and type of machines in a particular organization, for example, can give away valuable data about the size of the organization, and hence the resources committed to a particular project.

Fortunately, the domain system includes an error code for "refused"; an administrative prohibition against such zone transfers is explicitly recognized as a legitimate reason for refusal. This code should be employed for zone transfer requests from any host not known to be a legitimate secondary server. Unfortunately, there is no authentication mechanism provided in the AXFR request; source address authentication is the best that can be done.

Recently, a compatible authentication extension to the DNS has been devised at M.I.T. The Hesiod name server[34] uses Kerberos[35] tickets to authenticate queries and responses. The *additional information* section of the query carries an encrypted ticket, which includes a session key; this key, known only to Hesiod and the client, is used to compute a cryptographic checksum of the both the query and the response. These checksums are also sent in the additional information field.

5.4 The File Transfer Protocol

The *File Transfer Protocol* (FTP)[36] itself is not flawed. However, a few aspects of the implementation merit some care.

5.4.1 FTP Authentication

FTP relies on a login and password combination for authentication. As noted, simple passwords are increasingly seen as inadequate; more and more sites are adopting one-time passwords. Nothing in the FTP specification precludes such an authentication method. It is vital, however, that the "331" response to a *USER* subcommand be displayed to the user; this message would presumably contain the challenge. An FTP implementation that concealed this response could not be used in this mode; if such implementations are (or become) common, it may be necessary to use a new reply code to indicate that

the user must see the content of the challenge.

5.4.2 Anonymous FTP

A second problem area is "anonymous FTP". While not required by the FTP specification, anonymous FTP is a treasured part of the oral tradition of the Internet. Nevertheless, it should be implemented with care.

One part of the problem is the implementation technique chosen. Some implementations of FTP require creation of a partial replica of the directory tree; care must be taken to ensure that these files are not subject to compromise. Nor should they contain any sensitive information, such as encrypted passwords.

The second problem is that anonymous FTP is truly anonymous; there is no record of who has requested what information. Mail-based servers will provide that data; they also provide useful techniques for load-limiting[9], background transfers, etc.

5.5 Simple Network Management Protocol

The *Simple Network Management Protocol* (SNMP)[37] has recently been defined to aid in network management. Clearly, access to such a resource must be heavily protected. The RFC states this, but also allows for a null authentication service; this is a bad idea. Even a "read-only" mode is dangerous; it may expose the target host to *netstat*-type attacks if the particular Management Information Base (MIB)[38] used includes sequence numbers. (The current standardized version does not; however, the MIB is explicitly declared to be extensible.)

5.6 Remote Booting

Two sets of protocols are used today to boot diskless workstations and gateways, *Reverse ARP* (RARP)[39] with the *Trivial File Transfer Protocol* (TFTP)[40] and BOOTP[41] with TFTP. A system being booted is a tempting target; if one can subvert the boot process, a new kernel with altered protection mechanisms can be substituted. RARP-based booting is riskier because it relies on Ethernet-like networks, with all the vulnerabilities adhering thereto. One can achieve a modest improvement in security by ensuring that the booting machine uses a random number for its UDP source port; otherwise, an attacker can impersonate the server and send false DATA packets.

BOOTP adds an additional layer of security by including a 4-byte random *transaction id*. This prevents an attacker from generating false replies to a workstation known to be rebooting. It is vital that these numbers indeed be random; this can be difficult in a system that is freshly powered up, and hence with little or no unpredictable state. Care should be taken when booting through gateways; the more networks traversed, the greater the opportunity for impersonation.

The greatest measure of protection is that normally, the attacker has only a single chance; a system being booted does not stay in that state. If, however, communications between the client and the standard server may be interrupted, larger-scale attacks may be mounted.

6. TRIVIAL ATTACKS

A few attacks are almost too trivial to mention; nevertheless, completeness demands that they at least be noted.

---

9. Recently, a host was temporarily rendered unusable by massive numbers of FTP requests for a popular technical report. If this were deliberate, it would be considered a successful denial of service attack.

6.1 Vulnerability of the Local Network

Some local-area networks, notably the Ethernet networks, are extremely vulnerable to eavesdropping and host-spoofing. If such networks are used, physical access must be strictly controlled. It is also unwise to trust any hosts on such networks if any machine on the network is accessible to untrusted personnel, unless authentication servers are used.

If the local network uses the Address Resolution Protocol (ARP)[42] more subtle forms of host-spoofing are possible. In particular, it becomes trivial to intercept, modify, and forward packets, rather than just taking over the host's role or simply spying on all traffic.

It is possible to launch denial of service attacks by triggering *broadcast storms*. There are a variety of ways to do this; it is quite easy if most or all of the hosts on the network are acting as gateways. The attacker can broadcast a packet destined for a non-existent IP address. Each host, upon receiving it, will attempt to forward it to the proper destination. This alone will represent a significant amount of traffic, as each host will generate a broadcast ARP query for the destination. The attacker can follow up by broadcasting an ARP reply claiming that the broadcast Ethernet address is the proper way to reach that destination. Each suspectible host will then not only resend the bogus packet, it will also receive many more copies of it from the other suspectible hosts on the network.

6.2 The Trivial File Transfer Protocol

TFTP[40] permits file transfers without any attempt at authentication. Thus, any publicly-readable file in the entire universe is accessible. It is the responsibility of the implementor and/or the system administrator to make that universe as small as possible.

6.3 Reserved Ports

Berkeley-derived TCPs and UDPs have the notion of a "privileged port". That is, port numbers lower than 1024 may only be allocated to privileged processes. This restriction is used as part of the authentication mechanism. However, neither the TCP nor the UDP specifications contain any such concept, nor is such a concept even meaningful on a single-user computer. Administrators should never rely on the Berkeley authentication schemes when talking to such machines.

7. COMPREHENSIVE DEFENSES

Thus far, we have described defenses against a variety of individual attacks. Several techniques are broad-spectrum defenses; they may be employed to guard against not only these attacks, but many others as well.

7.1 Authentication

Many of the intrusions described above succeed only because the target host uses the IP source address for authentication, and assumes it to be genuine. Unfortunately, there are sufficiently many ways to spoof this address that such techniques are all but worthless. Put another way, source address authentication is the equivalent of a file cabinet secured with an $100 lock; it may reduce the temptation level for more-or-less honest passers-by, but will do little or nothing to deter anyone even slightly serious about gaining entry.

Some form of cryptographic authentication is needed. There are several possible approaches. Perhaps the best-known is the Needham-Schroeder algorithm[20][21][22]. It relies on each host sharing a key with an authentication server; a host wishing to establish a connection obtains a session key from the authentication server and passes a sealed version along to the destination. At the conclusion of the dialog, each side is convinced of the identity of the other. Versions of the algorithm exist for both private-key and public-key[43] cryptosystems.

How do these schemes fit together with TCP/IP? One answer is obvious: with them, preauthenticated connections can be implemented safely; without them, they are quite risky. A second answer is that the DNS provides an ideal base for authentication systems, as it already incorporates the necessary name structure, redundancy, etc. To be sure, key distribution responses must be authenticated and/or

encrypted; as noted, the former seems to be necessary in any event.

In some environments, care must be taken to use the session key to encrypt the entire conversation; if this is not done, an attacker can take over a connection via the mechanisms described earlier.

7.2 Encryption

Suitable encryption can defend against most of the attacks outlined above. But encryption devices are expensive, often slow, hard to administer, and uncommon in the civilian sector. There are different ways to apply encryption; each has its strengths and weaknesses. A comprehensive treatment of encryption is beyond the scope of this paper; interested readers should consult Voydock and Kent[44] or Davies and Price[45].

Link-level encryption — encrypting each packet as it leaves the host computer — is an excellent method of guarding against disclosure of information. It also works well against physical intrusions; an attacker who tapped in to an Ethernet cable, for example, would not be able to inject spurious packets. Similarly, an intruder who cut the line to a name server would not be able to impersonate it. The number of entities that share a given key determines the security of the network; typically, a key distribution center will allocate keys to each pair of communicating hosts.

Link-level encryption has some weaknesses, however. Broadcast packets are difficult to secure; in the absence of fast public-key cryptosystems, the ability to decode an encrypted broadcast implies the ability to send such a broadcast, impersonating any host on the network. Furthermore, link-level encryption, by definition, is not end-to-end; security of a conversation across gateways implies trust in the gateways and assurance that the full concatenated internet is similarly protected. (This latter constraint may be enforced administratively, as is done in the military sector.) If such constraints are not met, tactics such as source-routing attacks or RIP-spoofing may be employed. Paranoid gateways can be deployed at the entrance to security domains; these might, for example, block incoming RIP packets or source-routed packets.

Many portions of the DARPA Internet employ forms of link encryption. All Defense Data Network (DDN) IMP-to-IMP trunks use DES encryption, even for non-classified traffic; classified lines use more secure cryptosystems[46]. These, however, are point-to-point lines, which are comparatively easy to protect.

A multi-point link encryption device for TCP/IP is the *Blacker Front End* (BFE)[47]. The BFE looks to the host like an X.25 DDN interface, and sits between the host and the actual DDN line. When it receives a call request packet specifying a new destination, it contacts an Access Control Center (ACC) for permission, and a Key Distribution Center (KDC) for cryptographic keys. If the local host is denied permission to talk to the remote host, an appropriate diagnostic code is returned. A special "Emergency Mode" is available for communications to a restricted set of destinations at times when the link to the KDC or ACC is not working.

The permission-checking can, to some extent, protect against the DNS attacks described earlier. Even if a host has been mislead about the proper IP address for a particular destination, the BFE will ensure that a totally unauthorized host does not receive sensitive data. That is, assume that a host wishes to send Top Secret data to some host *foo*. A DNS attack might mislead the host into connecting to penetrated host 4.0.0.4, rather than 1.0.0.1. If 4.0.0.4 is not cleared for Top Secret material, or is not allowed communications with the local host, the connection attempt will fail. To be sure, a denial of service attack has taken place; this, in the military world, is far less serious than information loss.

The BFE also translates the original ("Red") IP address to an encrypted ("Black") address, using a translation table supplied by the ACC. This is done to foil traffic analysis techniques, the bane of all multi-point link encryption schemes.

End-to-end encryption, above the TCP level, may be used to secure any conversation, regardless of the number of hops or the quality of the links. This is probably appropriate for centralized network management applications, or other point-to-point transfers. Key distribution and management is a greater problem, since there are more pairs of correspondents involved. Furthermore, since encryption

and decryption are done before initiation or after termination of the TCP processing, host-level software must arrange for the translation; this implies extra overhead for each such conversation[10].

End-to-end encryption is vulnerable to denial of service attacks, since fraudulently-injected packets can pass the TCP checksum tests and make it to the application. A combination of end-to-end encryption and link-level encryption can be employed to guard against this. An intriguing alternative would be to encrypt the data portion of the TCP segment, but not the header; the TCP checksum would be calculated on the cleartext, and hence would detect spurious packets. Unfortunately, such a change would be incompatible with other implementations of TCP, and could not be done transparently at application level.

Regardless of the method used, a major benefit of encrypted communications is the implied authentication they provide. If one assumes that the key distribution center is secure, and the key distribution protocols are adequate, the very ability to communicate carries with it a strong assurance that one can trust the source host's IP address for identification.

This implied authentication can be especially important in high-threat situations. A routing attack can be used to ''take over'' an existing connection; the intruder can effectively cut the connection at the subverted machine, send dangerous commands to the far end, and all the while translate sequence numbers on packets passed through so as to disguise the intrusion.

It should be noted, of course, that any of these encryption schemes provide privacy. Often that is the primary goal of such systems.

7.3 Trusted Systems

Given that TCP/IP is a Defense Department protocol suite, it is worth asking to what extent the Orange Book[48] and Red Book[49] criteria would protect a host from the attacks described above. That is, suppose that a target host (and the gateways!) were rated B1 or higher. Could these attacks succeed? The answer is a complex one, and depends on the assumptions we are willing to make. In general, hosts and routers rated at B2 or higher are immune to the attacks described here, while C2-level systems are susceptible. B1-level systems are vulnerable to some of these attacks, but not all.

In order to understand how TCP/IP is used in secure environments, a brief tutorial on the military security model is necessary. All *objects* in the computer system, such as files or network channels, and all data exported from them, must have a *label* indicating the sensitivity of the information in them. This label includes hierarchical components (i.e., Confidential, Secret, and Top Secret) and non-hierarchical components. *Subjects* — i.e., processes within the computer system — are similarly labeled. A subject may *read* an object if its label has a higher or equal hierarchical level and if all of the object's non-hierarchical components are included in the subject's label. In other words, the process must have sufficient clearance for the information in a file. Similarly, a subject may write to an object if the object has a *higher* or equal level and the object's non-hierarchical components include all of those in the subject's level. That is, the sensitivity level of the file must be at least as high as that of the process. If it were not, a program with a high clearance could write classified data to a file that is readable by a process with a low security clearance.

A corollary to this is that for read/write access to any file, its security label must exactly match that of the process. The same applies to any form of bidirectional interprocess communication (i.e., a TCP virtual circuit): both ends must have identical labels.

We can now see how to apply this model to the TCP/IP protocol suite. When a process creates a TCP connection, that connection is given the process's label. This label is encoded in the IP security option. The remote TCP must ensure that the label on received packets matches that of the receiving process.

---

10. We are assuming that TCP is handled by the host, and not by a front-end processor.

Servers awaiting connections may be eligible to run at multiple levels; when the connection is instantiated, however, the process must be forced to the level of the connection request packet.

IP also makes use of the security option[50]. A packet may not be sent over a link with a lower clearance level. If a link is rated for Secret traffic, it may carry Unclassified or Confidential traffic, but it may not carry Top Secret data. Thus, the security option constrains routing decisions. The security level of a link depends on its inherent characteristics, the strength of any encryption algorithms used, the security levels of the hosts on that network, and even the location of the facility. For example, an Ethernet cable located in a submarine is much more secure than if the same cable were running through a dormitory room in a university.

Several points follow from these constraints. First, TCP-level attacks can only achieve penetration at the level of the attacker. That is, an attacker at the Unclassified level could only achieve Unclassified privileges on the target system, regardless of which network attack was used[11]. Incoming packets with an invalid security marking would be rejected by the gateways.

Attacks based on any form of source-address authentication should be rejected as well. The Orange Book requires that systems provide secure means of identification and authentication; as we have shown, simple reliance on the IP address is not adequate. As of the B1 level, authentication information must be protected by cryptographic checksums when transmitted from machine to machine[12].

The *authentication* server is still problematic; it can be spoofed by a sequence number attack, especially if *netstat* is available. This sort of attack could easily be combined with source routing for full interactive access. Again, cryptographic checksums would add significant strength.

B1-level systems are not automatically immune from routing attacks; RIP-spoofing could corrupt their routing tables just as easily. As seen, that would allow an intruder to capture passwords, perhaps even some used on other trusted systems. To be sure, the initial penetration is still restricted by the security labelling, but that may not block future logins captured by these means.

Routing attacks can also be used for denial of service. Specifically, if the route to a secure destination is changed to require use of an insecure link, the two hosts will not be able to communicate. This change would probably be detected rather quickly, though, since the gateway that noticed the misrouted packet would flag it as a security problem.

At the B2 level, secure transmission of routing control information is required. Similar requirements apply to other network control information, such as ICMP packets.

Several attacks we have described rely on data derived from "information servers", such as *netstat* and *finger*. While these, if carefully done, may not represent a direct penetration threat in the civilian sense, they are often seen to represent a *covert channel* that may be used to leak information. Thus, many B-division systems do not implement such servers.

In a practical sense, some of the technical features we have described may not apply in the military world. Administrative rules[51] tend to prohibit risky sorts of interconnections; uncleared personnel are not likely to have even indirect access to systems containing Top Secret data. Such rules are, most likely, an accurate commentary on anyone's ability to validate any computer system of non-trivial size.

8. CONCLUSIONS

Several points are immediately obvious from this analysis. The first, surely, is that in general, relying on the IP source address for authentication is extremely dangerous[13]. Fortunately, the Internet

---

11. We are assuming, of course, that the penetrated system does not have bugs of its own that would allow further access.

12. More precisely, user identification information must be protected to an equal extent with data sensitivity labels. Under certain circumstances, described in the Red Book, cryptographic checks may be omitted. In general, though, they are required.

13. There are some exceptions to this rule. If the entire network, and all of its components (hosts, gateways, cables, etc.) are physically protected, and if all of the operating systems are sufficiently secure, there would seem to be little risk.

community is starting to accept this on more than an intellectual level. The Berkeley manuals[3] have always stated that the authentication protocol was very weak, but it is only recently that serious attempts (i.e., Kerberos[35] and SunOS 4.0's DES authentication mode[52]) have been made to correct the problem. Kerberos and SunOS 4.0 have their weaknesses, but both are far better than their predecessor. More recently, an extension to the *Network Time Protocol* (NTP)[53] has been proposed that includes a cryptographic checksum[54].

A second broad class of problems is sequence number attacks. If a protocol depends on sequence numbers — and most do — it is vital that they be chosen unpredictably. It is worth considerable effort to ensure that these numbers are not knowable even to other users on the same system.

We may generalize this by by stating that hosts should not give away knowledge gratuitously. A *finger* server, for example, would be much safer if it only supplied information about a known user, rather than supplying information about everyone logged on. Even then, some censorship might be appropriate; a refusal to supply the last login date and other sensitive information would be appropriate if the account was not used recently. (Never-used accounts often have simple default passwords. Infrequently-used accounts are often set up less carefully by the owner.) We have also seen how *netstat* may be abused; indeed, the combination of *netstat* with the *authentication* server is the single strongest attack using the standardized Internet protocols.

Finally, network control mechanisms are dangerous, and must be carefully guarded. Static routes are not feasible in a large-scale network, but intelligent use of default routes and verifiable point-to-point routing protocols (i.e., EGP) are far less vulnerable than broadcast-based routing.

9. ACKNOWLEDGEMENTS

Dave Presotto, Bob Gilligan, Gene Tsudik, and especially Deborah Estrin made a number of useful suggestions and corrections to a draft of this paper.

REFERENCES

1. E.J. Feinler, O.J. Jacobsen, M.K. Stahl, C.A. Ward, eds. *DDN Protocol Handbook*. DDN Network Information Center, SRI International, 1985.

2. Comer, D. *Internetworking with TCP/IP : Principles, Protocols, and Architecture*. Prentice Hall, 1988

3. Computer Systems Research Group. *UNIX User's Reference Manual (URM). 4.3 Berkeley Software Distribution Virtual Vax-11 Version.* Computer Science Division, Department of Electrical Engineering and Computer Science, University of California, Berkeley. 1986.

4. Spafford, E.H. *The Internet Worm Program: An Analysis*. Purdue Technical Report CSD-TR-823, Department of Computer Sciences Purdue University, West Lafayette, IN. 1988

5. Seeley, D. *A Tour of the Worm*. Department of Computer Science, University of Utah. 1988.

6. Eichin, M. and Rochlis, J. *With Microscope and Tweezers: An Analysis of the Internet Virus of November 1988*. Massachussetts Institute of Technology, 1988.

7. Morris, R.T. 1985. *A Weakness in the 4.2BSD UNIX TCP/IP Software*. Computing Science Technical Report No. 117, AT&T Bell Laboratories, Murray Hill, New Jersey.

8. Reynolds, J.K., and J. Postel. *Assigned Numbers*. RFC 990, 1986

9. Mills, D.L. *Internet Delay Experiments*, RFC 889, 1983.

10. Blum, M. and Micali, S. "How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits". *SIAM J. Computing*, vol. 13, no. 4, pp. 850-864, Nov. 1984.

11. US Federal Information Processing Standards Publication (FIPS PUB) 46, *Data Encryption Standard*, 15 January 1977.

12. US Federal Information Processing Standards Publication (FIPS PUB) 81. *DES Modes of Operation*, 2 December 1980.

13. Bishop, M. *An Application of a Fast Data Encryption Standard Implementation*. Technical Report PCS-TR88-138, Department of Mathematics and Computer Science, Dartmouth College, Hanover, NH. 1988.

14. Mogul, J. "Simple and Flexible Datagram Access Controls for UNIX-based Gateways", *Proceedings, Summer USENIX*, 1989, Baltimore, Maryland (to appear).

15. Hedrick, C. *Routing Information Protocol*. RFC 1058, 1988.

16. Mills, D.L. *Exterior Gateway Protocol Formal Specification*. RFC 904, 1984.

17. Postel, J. *Internet Control Message Protocol*. RFC 792, 1981.

18. St. Johns, M. *Authentication Server*. RFC 931, 1985.

19. Defense Advanced Research Projects Agency, Internet Activities Board. *IAB Official Protocol Standards*. RFC 1083, 1988

19. Postel, J. Private communication. 1989.

20. Needham, R.M. and Schroeder, M.D. "Using Encryption for Authentication in Large Networks of Computers". *Communications of the ACM*, vol. 21, no. 12, pp. 993-999, December 1978.

21. Denning, D.E. and Sacco, G.M. "Timestamps in Key Distribution Protocols", *Communications of the ACM*, vol. 24, no. 8, pp. 533-536, August 1981.

22. Needham, R.M. and Schroeder, M.D. "Authentication Revisited", *Operating Systems Review*, vol. 21, no. 1, p. 7, January 1987.

23. Harrenstien, K. *NAME/FINGER Protocol*, RFC 742, 1977.

24. Grampp, F.T. and Morris, R.H. "UNIX Operating System Security", *AT&T Bell Laboratories Technical Journal*, vol. 63, no. 8, part 2, October, 1984.

25. Crocker, D. *Standard for the Format of ARPA-Internet Text Messages*. RFC 822, 1982.

26. Postel, J. *Simple Mail Transfer Protocol*. RFC 821, 1982.

27. Linn, J. *Privacy Enhancement for Internet Electronic Mail: Part I: Message Encipherment and Authentication Procedures*. RFC 1040, 1988.

28. Butler, M.; Postel, J.B.; Chase, D.; Goldberger, J.; Reynolds, J.K. *Post Office Protocol - Version 2*. RFC 937, 1985.

29. Diffie, W. "The First Ten Years of Public Key Cryptography". *Proc. IEEE*, vol. 76, no. 5, pp. 560-577, May 1988.

30. Rose, M. *Post Office Protocol - Version 3*. RFC 1081, 1988

31. Lambert, M.L. *PCMAIL: A Distributed Mail System for Personal Computers*. RFC 1056, 1988

32. Mockapetris, P. *Domain Names - Concepts and Facilities*. RFC 1034, 1987.

33. Mockapetris, P. *Domain Names - Implementations and Specifications*. RFC 1035, 1987.

34. Dyer, S.P. "Hesiod", *Proceedings, Winter USENIX*, 1988, Dallas, Texas.

35. Steiner, J.G, Neuman, C., Schiller, J.I. "Kerberos: An Authentication Service for Open Network Systems", *Proceedings, Winter USENIX*, 1988, Dallas, Texas.

36. Postel, J. *File Transfer Protocol*. RFC 959, 1985.

37. Case, J., Fedor, M., Schoffstall, J., and Davin, J. *A Simple Network Management Protocol*. RFC 1067, 1988.

38. McCloghrie, K. and Rose, M. *Management Information Base for Network Management of TCP/IP-based Internets*. RFC 1066. 1988.

39. Finlayson, R.; Mann, T.; Mogul, J.; Theimer, M. *Reverse Address Resolution Protocol*. RFC 903, 1984.

40. Sollins, K.R. *The TFTP Protocol (Revision 2)*. RFC 783, 1981.

41. Croft, W.J.; Gilmore, J. *Bootstrap Protocol*. RFC 951, 1985.

42. Plummer, D.C. *An Ethernet Address Resolution Protocol*. RFC 826, 1982.

43. Diffie, W. and Hellman, M.E. "New Directions in Cryptography." *IEEE Transactions on Information Theory*, vol. IT-22, no. 6, pp. 644-654.

44. Voydock, V.L. and Kent, S.T. "Security Mechanisms in High-Level Network Protocols". *ACM Computer Surveys*, vol. 15, no. 2, pp. 135-171, June 1983.

45. Davies, D.W. and Price, W.L. *Security for Computer Networks: An Introduction to Data Security in Teleprocessing and Electronic Funds Transfer*. Wiley. 1984.

46. Defense Communications Agency. *Defense Data Network Subscriber Security Guide*. 1983.

47. "Blacker Front End Interface Control Document", in *DDN Protocol Handbook*. DDN Network Information Center, SRI International, vol. 1, 1985.

48. DoD Computer Security Center. *DoD Trusted Computer System Evaluation Criteria*, 1983, CSC-STD-001-83.

49. National Computer Security Center. *Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria*. NCSC-TG-005, Version 1, July 31, 1987.

50. St. Johns, M. *Draft Revised IP Security Option*. RFC 1038, 1988.

51. DoD Computer Security Center. *Technical Rationale Behind CSC-STD-003-85: Computer Security Requirements*, CSC-STD-004-83, 1983.

52. Taylor, B. and Goldberg, D. "Secure Networking in the Sun Environment". *Proceedings, Summer USENIX*, 1986, Atlanta, Georgia.

53. Mills, D.L. *Network Time Protocol (Version 1) Specification and Implementation*. RFC 1059, 1988.

54. Mills, D.L. Mailing list message <8901192354.aa03743@Huey.UDEL.EDU>, January 19, 1989.