

Naval Research Laboratory

Washington, DC 20375-5320



AD-A268 579



NRL/FR/5544--93-9561

An Internetwork Authentication Architecture

RANDALL J. ATKINSON

*Center for High Assurance Computing Systems
Information Technology Division*

August 5, 1993



93 8 25 013

93-19908



Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE August 5, 1993		3. REPORT TYPE AND DATES COVERED
4. TITLE AND SUBTITLE An Internetwork Authentication Architecture			5. FUNDING NUMBERS PE - 03390NN	
6. AUTHOR(S) Randall J. Atkinson				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory Washington, DC 20375-5320			8. PERFORMING ORGANIZATION REPORT NUMBER NRL/FR/5544-93-9561	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Space and Naval Warfare Systems Command Washington, DC 20363-5100			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Current internetworks do not have effective host-to-host authentication mechanisms. The lack of these mechanisms contributes substantially to currently widespread network security problems. Public key authentication is a good approach to providing optional authentication in internetworks. An internetwork authentication architecture using public key authentication technology is proposed as a possible mechanism to substantially improve the security of large internetworks. Limitations of the proposed authentication architecture are also described.				
14. SUBJECT TERMS Network Authentication Security Digital signature Internet			15. NUMBER OF PAGES 15	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

CONTENTS

INTRODUCTION	1
AUTHENTICATION NEEDS	2
PROPOSED ARCHITECTURE	3
BENEFITS	5
LIMITATIONS	6
CONCLUSIONS	8
REFERENCES	8
APPENDIX A — Application to the Internet Domain Name System	11
APPENDIX B — Application to the XTP Protocol	13

DTIC QUALITY INSPECTED 3

DTIC QUALITY INSPECTED 3

Accession For	
DTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

AN INTERNETWORK AUTHENTICATION ARCHITECTURE

INTRODUCTION

Historically, most networking protocols and architectures have not included solid authentication or confidentiality mechanisms. The Massachusetts Institute of Technology (MIT) Athena project has been the exception to this rule with its development of the Kerberos authentication system [Miller88, Dyer88]. This is beginning to be implemented at some sites and some workstation manufacturers are considering implementing Kerberos in their standard operating system (OS) releases, but the overwhelming majority of networked sites have no authentication or confidentiality mechanisms in their network architectures. The International Standards Organization Open Systems Interconnection (ISO OSI) suite provides for confidentiality services in the upper layers but does not require authentication of any of the lower layer protocols. Steven Bellovin has pointed out a number of security problems in protocols commonly used in the Internet and has also pointed out certain limitations in the Kerberos protocols [Bellovin89, Bellovin91]. The security issues in the ISO OSI suite appear to have gotten less attention than in the Internet suite because the Internet suite is more widely implemented at present.

Recently, the Internet Engineering Task Force has begun to incorporate authentication and confidentiality mechanisms in some protocols, notably the Simple Network Management Protocol (SNMP) and Privacy Enhanced Mail (PEM) [Galvin92, Linn93, Kent93, Linn89]. A few other recent protocol specifications, such as for the Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF) routing protocols provide hooks for authentication to be added later but do not define or mandate any real authentication mechanism [Lougheed91, Hedrick88, Moy91]. The BGP Version 3 specification explicitly states that the definition of authentication mechanisms other than the default "no authentication" option are out of the scope of the specification. Similarly, the OSPF Version 2 specification asserts that "OSPF also provides for the authentication of routing updates, ..." when in fact the only authentication mechanisms specified are "no authentication" or "cleartext password." Overall, no fundamental systemic security architecture is apparent in the Internet protocol suite at present.

Similarly, the ISO protocol suite has not paid sufficient attention to building security mechanisms into the network, transport, or routing protocols. Like some of the protocols in the Internet suite, hooks for security have been provided. In the transport layer, TP1 through TP4 provide an option in the variable length part of the header that can contain user-defined security data having a user-defined length. This option appears to be for sensitivity labelling similar to that in the Internet Protocol (IP) Security Option [Kent91], though its use isn't entirely clear because it isn't fully defined in the Connectionless Network Protocol (CLNP) specification [ISO88]. This incomplete definition makes interoperability unlikely. The European Computer Manufacturers Association (ECMA) Connectionless Internetwork Protocol specification permits user-defined options, but does not itself define

mandatory or optional authentication or confidentiality mechanisms [ECMA84]. The ISO OSI protocol suite has historically been more concerned with providing security mechanisms in their applications. X.400 defines security mechanisms for e-mail and X.509 defines security mechanisms for the X.500 directory service.

The goals of the architecture presented in this report are to provide systemic authentication services to all hosts and users in a large internetwork where hosts are trustworthy systems operating at or above the Orange Book's C2 level [DoD85]. User-to-host authentication is not addressed in this architecture. The architecture should work with conventional passwords or other techniques such as biometric authentication being employed for user-to-host authentication and identification. This goal is unlike that of Kerberos which seeks to authenticate a user of an untrusted workstation to various networked servers and services. However, the proposed architecture is designed to authenticate networked trusted systems to each other and deliberately does not address user-to-host authentication. The use of trusted computer systems technology as an enabling base has only recently become practical. Several major workstation vendors have commercial B1 trusted system products currently in evaluation at the National Computer Security Center [DoD85].

AUTHENTICATION NEEDS

The Internet Protocol does not currently provide any mechanism for authentication of an IP frame that is received by a host on the Internet [ISI81a]. As a result, there are a number of possible attacks on the network and hosts using the network. Bellovin [Bellovin89] describes a number of attacks at the network layer. IP source routing can be used by an intruder to masquerade as a trusted host. The Address Resolution Protocol (ARP) and its cousin, the Reverse Address Resolution Protocol (RARP), lack authentication and make it easy for an intruder to masquerade as a trusted host [Plummer82, Finlayson84]. False routing information may be disseminated and cause denial of service or host masquerading attacks. All of these attacks can be prevented by providing authentication in IP frames.

The ISO protocol suite is not really any more secure than the Internet protocol suite. User-defined options may be created to hold authentication information, but currently no authentication mechanisms are defined in ISO TP4 or ISO CLNP [ISO86, ISO88]. Without such authentication, it is not possible to implement dependable policy-based routing. Similarly, attacks similar to those described by Bellovin [Bellovin89] are believed to be possible against the ISO protocol suite. For example, because the Subnetwork Dependent Convergence Function of the ISO CLNP definition lacks any authentication mechanism, attacks on address resolution and on underlying X.25 virtual circuits appear feasible [IEEE87, ISO88].

Aside from concerns about attacks, there is recently much interest in implementing policy-based routing, network usage accounting, and network auditing. None of these can be dependably implemented unless the network protocol headers can be authenticated by routers as well as the end hosts. If there is no intermediate authentication, then it is straightforward to spoof policy-based routing and to cause others to pay for one's network traffic. Without authentication, auditing cannot yield meaningful results. It is clear that network protocol header authentication is essential for both existing and future services. It makes more sense to authenticate the entire network protocol frame than the header data alone. The incremental cost of authenticating the entire frame instead of just the headers is not significant and the increased entropy and size of the authenticated information makes many cryptanalytic attacks on the authentication harder, while also ensuring the authenticity of the data.

Bellovin and Morris [Bellovin89, Morris85] have each described a number of attacks at the transport layer, such as using TCP sequence number prediction to masquerade as another host's connection. Even trustworthy hosts need to isolate user connections from one another and to ensure that no user is capable of masquerading as another user via networking mechanisms. The ability to provide circuit-oriented confidentiality mechanisms is also desirable. Neither TCP nor TP4 currently provides either authentication or confidentiality mechanisms [ISI81b, ISO86].

While it is possible to support transport authentication using entirely different mechanisms than those used to provide network authentication, it is desirable to devise a common approach to authentication so that the overhead of implementation is minimized and so that the different services integrate nicely. Moreover, use of a common authentication mechanism should reduce the size of the trusted code required to implement the authentication services, thereby allowing easier verification of the implementation.

Another critical service that needs authentication is the network name service. If an intruder can masquerade as the legitimate name service provider, he can cause denial-of-service attacks, modify data in transit, and make other attacks on users of the internetwork. If however, the name service were authenticated, these attacks would not be possible. Moreover, as Needham and Schroeder observed, an authenticated name service is a good mechanism to use for distributing host authentication or encryption keys [Needham78, Needham87]. The name server could be used to distribute keys used in various layers and protocols, including network layer and transport layer authentication keys. As mentioned above, the Consultative Committee for International Telephone and Telegraph (CCITT) X.509 recommendation describes security mechanisms for use in an X.500 Directory Service [CCITT88]. An authenticated name service might also be used to distribute user "keys" or a host might distribute keys for its users via a new user key distribution protocol.

PROPOSED ARCHITECTURE

Cryptographic mechanisms provide the greatest assurance of the authenticity of data. Cryptographic systems come in two varieties, symmetric key and asymmetric key. In a symmetric key system, the same key is used for encryption and decryption. When providing confidentiality using an asymmetric system, each party has two keys, one public and one private, and data are encrypted using the sender's private key and the recipient's public key and decrypted using the recipient's private key and the sender's public key. When providing authentication using an asymmetric system, the data and the keys are used to generate a digital signature without encrypting the data itself. That signature can be verified by the recipient using the data received and the appropriate decryption keys. Asymmetric cryptography was chosen as the basis for this architecture because the ability to publish all of the public keys precludes the more severe key distribution problem which arises when symmetric key technology is used. Another advantage is that protocols for asymmetric cryptography tend to have less complexity than do their private key equivalents and, hence, it is easier to provide formal assurance that the protocols used are correct [Needham87].

By authenticating all name service responses, it becomes straightforward to distribute the keys to all hosts and users of the internetwork. Public keys for hosts are included in the name service database and all name service responses are authenticated by the recipient of the response. This means that all the host's public keys are distributed in an authenticated manner. Name service requests need not be authenticated or confidential in the general case. However, if the visibility of some data in the name service database is to be controlled, then authenticated confidential requests would be required to access nonpublished data and authenticated confidential responses to such requests would also be required. The value of the public keys for the root name servers should be made readily available by

telephone and postal mail so that system administrators may have confidence in the authenticity of the root public key and so that an intruder would not be able to easily masquerade as the legitimate name server.

Many networking services can be enhanced with confidentiality and authentication capabilities if the name service is authenticated. For example, there might be more than one public key for each host or network address. For example, the host might use a different key for network management than it uses for network protocol frame authentication. Each user could have his own public keys for use with privacy enhanced mail, virtual terminals, or other applications. User's own public keys should probably not be in the network name service database, but rather should be available only from the user's normal host (or possibly from a specific host out of the user's normal hosts). The name service could indicate which host held published keys for a domain, subdomain, or host. Obtaining a user's authentic public key would thus require one additional transaction. However, the reduction in the network name service database size might be significant and make it worthwhile.

To obtain the public key of any arbitrary desired host in the internetwork, the following sequence is used. First, the host seeking the public key would make a name service request of the root name server or a locally trusted name server containing cached root name server data. That request would be for the public key of the desired host. The response would contain either the name of the name server for the domain containing the desired host or the public key of the desired host. All responses would be authenticated using the public key of the name server and any inauthentic responses would be discarded and ignored. It might be valuable to audit all inauthentic responses. This process would be repeated as necessary until the requesting host received an authentic response containing the public key of the desired other host. If the locally trusted name server used caching of data, response time would be reasonable despite having authentication. Using local name servers and caching is a good implementation strategy for name service regardless of whether authentication is used.

To enable any intermediate network device to authenticate the contents of the network frame header, the decryption key for each host is published, and the encryption key is kept private by that host. The sending host uses its private encryption key plus the header data to generate a cryptographic checksum which is embedded in the header. Any intermediate node reads the header to determine the sender's identity and attempts to confirm that the claimed sender's published decryption key produces the correct results when applied to the embedded cryptographic checksum. If it does, then the sender and the other header data are authentic. Otherwise, some part of the header data isn't authentic. This permits policy-based routing and usage-based accounting to be dependably implemented. This same technology can be used to provide authentication of message contents by simply computing the cryptographic checksum over the entire frame rather than just the frame header. This procedure does not depend on any particular cryptographic checksum or message digest. The MD5 Message Digest Algorithm is an example of a cryptographic checksum that is in the public domain and might be suitable [Rivest92].

Personal or application level keys are distributed by the user's host and are authenticated with that host's authentication key. Transport keys and session keys are also distributed in this manner. Using this approach, confidentiality and authentication may be built into applications above the transport layer and also into the transport layer itself. In some cases, it might be desirable instead to use mechanisms built into the upper layer protocol that are independent of these mechanisms. For example, the Secure SNMP proposals build authentication and optional confidentiality mechanisms into the SNMP protocols [Galvin92]. This approach has the advantage that a security breach at a lower layer does not necessarily compromise the security at the upper layer. However, host masquerading attacks still appear feasible even in the presence of Secure SNMP. These mechanisms could be used either to distribute keys or key certificates.

BENEFITS

Implementation of this architecture would have a number of benefits. The obvious benefit is that a large number of networking security problems, including those identified above, are eliminated. This is a significant step also towards enhancing network reliability since all of the attacks above also impair reliability. The security and reliability of a computer system are inseparably related. The ability to implement network auditing, network traffic counting, and policy-based routing are made feasible.

Additionally, this authentication architecture could be used to implement the Clark-Wilson commercial security policy over a network or internetwork [Clark87]. To support Clark-Wilson, authentication of users' real identities is essential. In the approach suggested here, the hosts would be authenticated to each other and could provide user authentication keys or such keys could be placed in a central directory service with its responses being authenticated. Full protection from host masquerading and network traffic control policies could be easily enforced. Since the Clark-Wilson policy is more concerned with integrity than confidentiality, this might be sufficient for a commercial firm or educational institution. Confidentiality could easily be added at the transport layer or above if it were needed and need not degrade performance for applications or users that don't need it.

With a few extensions, the approach outlined here could also support a multilevel security policy using either a pink architecture or a red/black architecture [Cole89]. For example, there might be encryption of user data immediately above the transport layer or the transport layer itself might be encrypted. Either asymmetric or symmetric keys could be used, though use of the latter would complicate key management. Because the network layer is fully authenticated, the receiving host can be confident of where the transmission originated. Also, vulnerability to certain kinds of denial of service attacks can be significantly reduced by precluding the attacks described earlier. Use of link encryption below the network layer to minimize the effectiveness of traffic analysis remains feasible and is unaffected by network layer or higher mechanisms such as these.

Additionally, support for confidentiality between hosts and between users is now practical at a slight incremental cost paid only by users of such confidentiality services. Authenticated public encryption keys for use at the network, transport, or higher layers can be easily distributed with authentication. The encryption of data for commercial and educational users might best be placed above the transport layer because traffic analysis of the headers is not a significant concern in those environments. In many computer systems, the protocols above the transport layer are not in privileged parts of the operating system, while the transport layer and below frequently are. This means that an implementation of the encryption scheme above the transport layer would help separate the functions, thus making verification easier. This also means that the encryption implementation could be changed by local authorities without having to modify the operating system software. Moreover, it would avoid complicating the lower level protocols with services easily provided above.

It appears feasible to implement the required protocol changes in a way that would retain interoperability with older versions. Moreover, this architecture scales nicely to large internetworks such as the current Internet. There are a number of hardware implementations of Data Encryption Standard (DES) available already and it is feasible to implement digital signature algorithms and asymmetric key cryptography in hardware as well [NBS77]. If these were integrated into a chipset, the cost of authentication would be minimized. Depending on the algorithm chosen, the cost of authentication might still be too high to permit authentication of all traffic at all intermediate systems. However, hosts that do not wish to use authentication don't have to. Only the root name servers and hosts wishing to use authentication services need pay for its implementation costs and overhead.

LIMITATIONS

The use of confidentiality services in a multicast network layer protocol might not work very well because of key distribution difficulties. In particular, there are no known protocols to initially set up a group or multicast key using either symmetric key technology or public key technology. This makes dynamic changes in group membership difficult. It is feasible to use symmetric keys to authenticate data sent to a group as coming from some indeterminate member of that group. Similarly, public key technology could be used effectively once some pair of group keys were distributed to each member of the group.

There is some cost to implementing either authentication or confidentiality services. An early implementor of the Secure SNMP specifications has reported that Secure SNMP with a tuned MD4 implementation is about 10% slower than ordinary SNMP (without authentication) and that Secure SNMP with an untuned MD5 implementation is 15-20% slower than ordinary SNMP. Public key authentication algorithms exist in the public literature. The costs of authentication depend greatly upon the algorithm chosen and the implementation method. Software implementations of some asymmetric algorithms are believed to be too slow to make their use practical in this architecture. Hardware implementations are more expensive in initial cost but have much better performance. More detailed information about performance costs of various algorithms and implementations is needed. Such information would directly influence the choice of digital signature algorithm. The selection of any particular algorithm is beyond the scope of this work.

There is another cost to implementing network layer authentication. One must have the entire original network frame intact to attempt to authenticate it. Network frames are frequently segmented into smaller frames that will fit within the size limitations of the protocols in and underneath the link layer. This means that at each point where a node wishes to attempt to authenticate the IP frame, it must reassemble all of the components of the original IP frame first. It also means that if any intermediate node does not reassemble the original frame before resending or resends different segments of a given IP frame over different routes, intermediate nodes downstream from that node will be unable to authenticate the segmented IP frames.

In most cases currently, reassembly only happens at the destination node. Intermediate nodes, such as routers, need not pay this cost. Reassembly and potential subsequent refragmentation can cause significant overheads when the link and physical protocols carry very small amounts of data in each lower level frame. Recently there has been much interest in using the CCITT's Asynchronous Transfer Mode (ATM) technology to build very-high-speed computer networks [CCITT91]. ATM is probably the most expensive case today since it has an extremely small frame size. In the cases of Ethernet or FDDI, this cost is not nearly as high as it would be with ATM. However, in the case of ATM, it appears that reassembly of the original IP frame will be necessary if the IP frame is to be transported over some other link medium because an ATM cell cannot contain a full set of IP headers. Also, the small size of the ATM cell means that the efficiency of other media such as FDDI would be greatly decreased if the original IP frame were not reassembled before encapsulation in an FDDI frame. Simulation and experimentation needs to be done to quantify the cost of the proposed changes to segmentation, reassembly, and routing. Routing in the proposed scheme, while still dynamic, is less flexible than in existing IP networks.

The delayed authentication scheme proposed by Tsudik has problems when being used as part of policy based routing because all but the last fragment of a network frame that isn't authentic can be sent over routes that the traffic is not allowed to use [Tsudik89]. This means that the deployment of the delayed authentication scheme could constitute an additional mechanism with which to launch a denial of service attack on a network or link. The other datagram authentication approach proposed

by Tsudik would restrict the size of the originating datagram to that allowed by the smallest of the link protocols used between source and destination and would require source routing the packets. This would be overly restrictive and seriously impair performance if, for example, an ATM network were in the circuit. In fact, it would not work at all for ATM because most network-layer protocols have minimum frame sizes larger than a single ATM cell. Moreover, this approach would require using source routing instead of the highly successful dynamic routing currently used in the Internet. The authentication mechanism proposed here is preferable to either of Tsudik's proposals.

Because the user and application level keys are distributed using mechanisms implemented in the local host, those keys may be changed easily by the user without much concern for the key change being delayed in propagation to all of the directory or network name service providers. Host keys are less easily changed, but such changes should be regularly scheduled to limit damage from compromised keys. If the application has its own built-in key exchange mechanism, it is more expensive to implement but provides additional protection against damage from a compromised host key. The desirability of such additional protection is determined largely by the perceived threats and value of the user data.

Because knowledge of the root private key by an intruder permits the intruder to control name service data and, hence, make host masquerading possible, the root keys should be changed periodically. To control damage in the case where the root keys are believed have been compromised, the change of the root key in systems in the internetwork should require human intervention and the use of reliable telephone or postal mail channels to pass the key. Because of the small number of root name servers, this should not present a terribly difficult key distribution problem for the root private key. It is not safe to use the name service to distribute changes to the root public key. The root private key is a critical single point of failure.

Similarly, knowledge of a particular name server's private key would permit an intruder to masquerade as any host within that domain and facilitate broad denial of service attacks against hosts within that domain or against hosts attempting to communicate with the compromised domain. Consequently, private keys for all name servers are also highly sensitive and should be treated as such.

Knowledge of a particular host's private key permits one to masquerade as that host and facilitates masquerading as any user of that host. So all host private keys are also sensitive. The knowledge of a particular user's private key without access to the private key of that user's host doesn't permit masquerading as that user from his usual host because the intruder cannot masquerade as the host without knowledge of the host's private key. However, if other users on a host know of a particular user's private key, they could successfully masquerade as him if there were no other system controls.

Only the host should know its own private key; no human should have knowledge of the key. The immediate key change approach used in the Secure SNMPS is one way to ensure this. In that scheme, the administrator will initially set the key to an appropriate secret value, but the host will immediately execute the key change procedure to change its public and private keys so that the value of the private key is not known or accessible to any human. The administrator would need to be able to force the host to change any of its keys in the event the administrator believed that that particular key had been compromised. It is also important to ensure that no two hosts use the same keys. Key selection methods are beyond the scope of this report, however, and won't be discussed further. The risk of attack is reduced by the use of well-partitioned trusted software in storing and handling keys internally and in implementing the networking, authentication, and confidentiality services.

CONCLUSIONS

This report outlines some known problems in network security for large internetworks and proposes an end-to-end authentication architecture that resolves the network security problems while also facilitating implementation of confidentiality services in upper layer protocols. The described architecture should work well for large internetworks as well as small networks. Its implementation in a reasonably backward-compatible way appears to be feasible by using the existing networking technology and infrastructure, although the performance of such an implementation is not immediately clear.

REFERENCES

- Bellovin, S.M., "Security Problems in the TCP/IP Protocol Suite," *ACM Computer Communicat. Rev.* 19(2), 32-48, 1989.
- Bellovin, S.M. and M. Merritt, "Limitations of the Kerberos Authentication System," *Proceedings of the Winter 1991 USENIX Conference*, USENIX Association, Berkeley, CA, 1991.
- CCITT, *The Directory—Authentication Framework*, Recommendation X.509, CCITT, Geneva, Switzerland, 1988.
- CCITT, *B-ISDN ATM Layer Specification*, Recommendation I.361, CCITT, Geneva, Switzerland, 1991.
- Chesson, G., *XTP Protocol Definition*, Revision 3.6, Protocol Engine, Inc., Santa Barbara, CA, 1992.
- Clark, D.D. and D.R. Wilson, "A Comparison of Commercial and Military Computer Security Policies," *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, IEEE Computer Society, Oakland, CA, 1987.
- Cole, R. Jr., D. Kallgren, R. Hale, and J.R. Davis, "Multilevel Secure Mixed-Media Communication Networks," *Proceedings of the 1989 IEEE Conference on Military Communications (MIL-COM '89)*, IEEE, NY, NY.
- Department of Defense, *Department of Defense Trusted Computer System Evaluation Criteria*, DoD 5200.28-STD, Ft. Meade, MD, 1985.
- Dyer, S.P., "The Hesiod Name Server," *Proceedings of the 1988 Winter USENIX Conference*, USENIX Association, Berkeley, CA, 1988.
- European Computer Manufacturer's Association, *Standard ECMA-92, Connectionless Internetwork Protocol*, ECMA, Geneva, Switzerland, 1984.
- Finlayson, R., T. Mann, J. Mogul, and M. Theimer, *A Reverse Address Resolution Protocol*, RFC-903, DDN Network Information Center, 1984.
- Galvin, J.M., K. McCloghrie, and J.R. Davin, *SNMP Security Protocols*, Internet Draft, NSFnet Network Service Center, 1992.
- Hedrick, C., *Routing Information Protocol*, RFC-1058, DDN Network Information Center, 1988.

- IEEE, *Logical Link Control (LLC) Protocol*, IEEE Standard 802.2, IEEE Computer Society, 1987.
- Information Sciences Institute, University of Southern California, *Internet Protocol Specification*, RFC-791, DDN Network Information Center, 1981.
- Information Sciences Institute, University of Southern California, *Transport Control Protocol*, RFC-793, DDN Network Information Center, 1981.
- International Standards Organization, *Open Systems Interconnection, Transport Protocol Specification*, IS-8073, ISO, 1986.
- ISO, *Protocol for Providing the Connection-Less Mode Network Service and Provision of Underlying Service (CLNP)*, IS-8473, ISO, 1988.
- Kent, S.T., *Privacy Enhancement for Internet Electronic Mail: Part II - Certificate-based Key Management*, RFC-1422, DDN Network Information Center, 1993.
- Kent, S.T., *US DoD Security Options for the Internet Protocol*, RFC-1108, DDN Network Information Center, 1991.
- Linn, J., *Privacy Enhancement for Internet Electronic Mail: Part I - Message Encryption and Authentication Procedures*, RFC-1421, DDN Network Information Center, 1993.
- Linn, J., *Privacy Enhancement for Internet Electronic Mail: Part III—Algorithms, Modes, and Identifiers*, RFC-1115, DDN Network Information Center, 1989.
- Lougheed, K., *Border Gateway Protocol, Version 3*, RFC-1267, DDN Network Information Center, October 1991.
- Morris, R.T., *A Weakness in the 4.2 BSD Unix TCP/IP Software*, CS Technical Report 117, AT&T Bell Laboratories, Murray Hill, New Jersey, 1985.
- Miller, S.P., B.C. Neuman, J.I. Schiller, and J.H. Saltzer, *Kerberos Authentication and Authorization System*, Project Athena Technical Plan, Section E.2.1, MIT, Cambridge, MA, 1988.
- Moy, J., *OSPF Routing Protocol, Version 2*, RFC-1247, DDN Network Information Center, 1991.
- National Bureau of Standards, Federal Information Processing Standards Publication 46: Data Encryption Standard, U.S. Department of Commerce, Washington, DC, 1977.
- Needham, R.M. and M.D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers," *Communicat. ACM* 21(12), 1978.
- Needham, R.M. and M.D. Schroeder, "Authentication Revisited," *ACM Operating Systems Review* 21(1), 1987.
- Plummer, D.C., *Ethernet Address Resolution Protocol*, RFC-826, DDN Network Information Center, 1982.
- Rivest, R. and S. Dusse, *The MD5 Message-Digest Algorithm*, RFC-1321, DDN Network Information Center, 1992.

Rivest, R.L., A. Shamir, and L. Adleman, "A Method for Obtaining Digital Structures and Public-Key Cryptosystems," *Communicat. ACM* 21(2), 120-127, 1978.

Tsudik, G., "Datagram Authentication in Internet Gateways: Implications of Fragmentation and Dynamic Routing," *IEEE Journal on Selected Areas in Communications* 7(4), 1989.

Appendix A

APPLICATION TO THE INTERNET DOMAIN NAME SYSTEM

This Appendix describes additions and changes to the Internet Protocol suite to enable it to distribute asymmetric keys and to enable its responses to be authenticated.

A new TYPE field is added to the resource records. This new field contains the asymmetric host authentication key to be used by hosts attempting to authenticate IP frames. Each host which transmits any authenticated frames must have this record in the Domain Name System (DNS) and the value of the record must be correctly advertised. The proposed name of this new DNS record type is HAK. The value of the HAK record is represented as hexadecimal numbers using the digits 0 through 9 and letters A through F. The HAK record's value is the value of the authentication key used for that host that the HAK record is associated with. No HAK records may exist that are not associated with a specific host. No host may have more than one HAK record. HAK records could contain key certificates rather than keys, if desired.

All DNS responses from name servers must provide authentication. All DNS requests should provide authentication. Hosts receiving an unauthenticated response should take note of the lack of authentication and may ignore unauthenticated responses if required by the security policy applicable to the domain of the receiving host. Hosts receiving a response containing incorrect authentication data should discard the response without processing it further. It is not suggested that an ICMP "IP Packet not authentic" message may be sent to the alleged sender of the frame because this would be easily subverted to a denial of service attack. Such an attack would only require that the attacker send out frames to arbitrary recipients that have random data in the authentication data field and have the address of the host to be attacked in the source address field.

To provide user asymmetric keys for encryption or authentication, it is suggested that a new service, the Key Information Protocol (KIP), be provided on a to-be-assigned TCP port. This service would accept requests for user public keys and would respond only if such information were available. The "no key exists for that user" and "that user not valid here" cases would both cause an "invalid request" to be sent back to the requestor. All responses would use IP authentication. The KIP would also use the host's public authentication key in the KIP response to enable the recipient to authenticate the response. KIP should provide for separate authentication and confidentiality keys. Depending on perceived need, KIP could even be extended to use a mechanism such as that described by Needham and Schroeder* to set up and use symmetric keys for some session with the two KIPs handling the key setup securely (each on behalf of its local user). The author feels that the use of the Needham and Schroeder symmetric key mechanism is less desirable than using asymmetric key technology because of the increased complexity.

If the KIP concept were to be implemented, it is suggested that a new DNS record be added that would point to the name of the host providing KIP service for a host or domain (similar to the way MX records work at present).

*Needham, R.M. and M.D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers," *Communicat. ACM* 21(12), 1978.

Appendix B

APPLICATION TO THE XTP PROTOCOL

This section describes proposed changes to the recently developed Xpress Transfer Protocol (XTP) [Chesson92]. XTP, like the other existing protocols discussed in the body of this report, does not attempt to provide for authentication or sensitivity labelling of XTP frames. In XTP, however, it would be much more straightforward to add such information because XTP uses a modern header-trailer frame format. An extension could be added in a way that would permit those interested in sensitivity labels or authentication to implement that without significant cost to those not interested in such features.

The trailer of XTP would need to be modified to add an additional 32-bit quadword at the end of the currently defined trailer. The first octet of this quadword would be used for the sensitivity (classification level) data and the second octet would be used for the protection authority data. If the data sensitivity were not being labelled, then both of these octets would have values of 0. The third octet would be used to identify the authentication data format being used for this frame. No authentication in use would have a value of 0. The final octet would be the length in quadwords (32 bits) of the authentication data that followed. If no authentication data followed, this length would, of course, be 0. If the authentication data did not fall on a 32-bit boundary, it would be padded with 0s appended to the right of the authentication information and the authentication data format code specification would specify what the actual data length was for each particular authentication data format. This preserves the neat 32-bit alignment of the frame.

Users not interested in authentication or labelling could ignore the values in the first 24 bits of the proposed extension and would use the value of the length field only to determine how many quadwords to skip, thus avoiding significant additional cost. Users of the labelling option that receive an unlabelled frame may treat it as unclassified or discard it, depending on the local security policy.

We propose that the values for the sensitivity label fields be as defined by the extant Request for Comments (RFC) describing the "IP Security Option" (sic) [Kent91].

There are only three initially suggested authentication modes. One is the degenerate case of no authentication and two actually provide some authentication. The existence of the no authentication case permits hosts or networks not interested in the offered security properties to go without them and not have to pay for what they do not seek to use. The first real authentication mode suggested would use the MD5 digital signature algorithm applied across the header of the XTP frame and then encoded using previously agreed upon DES encryption using the chained block mode of DES. The second real authentication mode would use the MD5 digital signature algorithm having been applied across the header of the XTP frame and then encoded using Rivest-Shamir-Adleman (RSA) encryption [Rivest78]. The third real authentication mode would use the MD5 digest algorithm having been applied across the entire XTP frame (exclusive of the authentication information field) and then have that encoded using RSA encryption.