

1 EXECUTIVE SUMMARY	1
1.1 Major Results/Conclusions.....	1
1.1.1 Diversity	1
1.1.2 Varying Requirements.....	2
1.1.3 Remove Obstacles to Deployment	2
1.1.4 Leverage experience	2
1.1.5 Reduce Barriers to Inter-agency Collaboration	2
1.1.6 Eliminate Clear-Text Passwords ted.....	2
1.2 Recommendations	3
2 INTRODUCTION	5
2.1 Background.....	5
2.2 Project Goals and Scope	6
3 TESTBED FINDINGS	7
3.1 Kerberos Testbed	7
3.1.1 Background.....	7
3.1.2 Purpose	9
3.1.3 Relationship to Other Testbeds.....	9
3.1.4 Activities of the Kerberos Testbed	9
3.1.5 What's Missing?.....	17
3.1.6 Next Steps.....	17
3.2 FORTEZZA Testbed	17
3.2.1 Background.....	18
3.2.2 Purpose	18
3.2.3 Relationship to Other Testbeds.....	18
3.2.4 Activities of the FORTEZZA Testbed	19
3.2.5 What's Missing?	22
3.2.6 Next Steps.....	23
3.2.7 References and Notes	23
3.3 Advanced Authentication Testbed.....	26
3.3.1 Background.....	26
3.3.2 Purpose	27
3.3.3 Authentication Information (Auth-Info).....	27
3.3.4 Authentication Protocols (Algorithms)	30
3.3.5 Relation to Other Testbeds	30
3.3.6 Activities of the Advanced Authentication Testbed.....	30
3.3.7 What's Missing.....	33
3.3.8 Next Steps.....	34

3.4	Public Key Infrastructure Testbed	34
3.4.1	Purpose	34
3.4.2	Background.....	34
3.4.3	Activities.....	35
3.4.4	Conclusions	36
3.4.5	What's Missing?.....	36
3.4.6	Next Steps.....	36
3.5	INCIDENT RESPONSE TESTBED	37
3.5.1	Background.....	37
3.5.2	Purpose	38
3.5.3	Relation to Other Testbeds	38
3.5.4	Activities of the Incident Response Testbed	38
3.5.7	What's Missing?.....	43
3.5.8	Next Steps.....	43
3.6	Security Testing and Certification Testbed	43
3.6.1	Background.....	44
3.6.2	Purpose	44
3.6.3	Relationship to Other Testbeds.....	45
3.6.4	Activities of the Testing and Certification Testbed.....	45
3.6.5	What's Missing?.....	46
3.6.6	Next Steps.....	46
3.7	Privacy, Digital Signature & Secure Messaging Testbeds.....	47
3.7.1	Background.....	47
3.7.2	Purpose	48
3.7.3	Relationship to Other Testbeds.....	48
3.7.4	Activities of the Privacy, Digital Signature & Secure Messaging Testbed.....	49
3.7.5	What's Missing?.....	50
3.7.6	Next Steps.....	50
3.7.7	References and Notes	51
3.8	Secure Web Testbed.....	51
3.8.1	Background.....	51
3.8.2	Purpose	52
3.8.3	Relationship to Other Testbeds.....	52
3.8.4	Activities of the Secure Web Testbed	52
3.8.5	What's Missing.....	55
3.8.6	Next Steps.....	56
4	CONCLUSIONS	58
4.1	Security Solutions.....	58
4.1.1	Feasible Alternatives to Passwords	58
4.2	Obstacles	58
4.2.1	User reluctance	58

4.2.2	Cost.....	59
4.2.3	Cross-Platform Support	59
4.2.4	Lack of Ideal Solutions.....	60
4.3	Deployment Strategies	60
4.4	Issues.....	61
4.4.1	Encryption Export Controls.....	61
4.4.2	Collaboration with Industry leaders MS and Netscape	62
4.4.3	Secure Web Technologies are Cost Effective and Provide a High Payoff.....	63
4.4.4	Leverage Other's Experience	63
4.4.5	Government Procedures Create Barriers to Effective Interagency Collaboration and Cooperation	63
4.4.6	Security Policies and Practices for Levels of Trust Between Government Agencies Need to be Developed	64
4.4.7	There are no "one size fits all" solutions	64
4.4.8	Personnel and security information resources need to be kept up to date..	65
4.5	Security Management Infrastructure (SMI)	65
5	RECOMMENDATIONS	67
5.1	Recommendations for Internet Managers	67
5.1.1	Don't Forget the Basics	67
5.1.2	Focus on Requirements before Technology	67
5.1.3	One Size Does Not Fit All.....	67
5.1.4	Look for Small Victories	68
5.1.5	Focus First on High Payoff Applications	68
5.1.6	Make Use of Available Expertise and Resources.....	68
5.2	Recommendations for the Federal Internet Community	68
5.2.1	Encourage Working Level Collaborative Projects	68
5.2.2	Remove Barriers to Collaborations	69
5.2.3	Encourage Multiple Solutions and Interoperability.....	69
5.2.4	Institutionalizing the CIS Concept	70
6	FOLLOW ON ACTIVITIES	71
6.1	Kerberos Testbed	71
6.2	FORTEZZA Testbed	71
6.3	Advanced Authentication Testbed.....	71
6.7	Privacy, Digital Signature and Secure Messaging Testbeds	72
6.8	Certificate Management Testbed.....	73

6.9	Electronic Laboratory Notebook Testbed.....	73
6.10	Intrusion Detection (ID) Testbed.....	75
6.11	DNSSEC Testbed.....	75
6.12	Push Technologies/Software Distributions	76
6.13	Network Scanning	76
APPENDIX A. EXAMPLES OF SECURITY BASICS.....		78
APPENDIX B. REFERENCES/LINKS.....		80
B.1	Participating Agencies	80
B.2	Incident Response & Monitoring:.....	81
B.3	Internet Security.....	81
B.4	Internet Privacy.....	83
B.5	Kerberos.....	84
B.6	FORTEZZA	84
B.7	Other Links	85
APPENDIX C: OTHER TECHNOLOGIES WE DID NOT ADDRESS.....		87
C.1	Firewalls.....	87
C.2	Technologies for Classified Systems	87
C.3	Hot New Technologies	88
C.4	Authorization Certificates.....	88
C.5	IPV6	88
C.6	Custom Applications.....	88

1 Executive Summary

Collaborations on Internet Security (CIS) is an interagency task force, set up under the auspices of the Privacy and Security Working Group (PSWG) of the federal Networking Council (FNC). Although most government agencies recognize that computer security is important, enforcement of security tends to be lax and good security tools have not been widely deployed. CIS members represented a broad cross-section of government agencies with varying security requirements. They answer the following questions:

- What are the differences and similarities in security requirements across government agencies?
- What good, broadly applicable computer security solutions exist today?
- What security solutions are going to be available in the near future?
- What are the obstacles to implementing these solutions in the government?
- In which areas can increased security provide the most cost-effective results?

The long-term goal of these inquiries was to develop a sustainable process for implementing interoperable security solutions across government agencies using a process geared towards openness and inclusiveness.

To address these questions, participants held regular meetings and engaged in lively e-mail discussions. Promising security approaches and tools were discussed and task forces were formed to gain hands-on experience with these techniques. The areas examined were Kerberos; FORTEZZA; Advanced Authentication; Public Key Infrastructure; Incident Response; Security Testing and Certification; Privacy, Digital Signature and Secure Messaging; and Secure Web.

To make the CIS work available for other agencies to use, Web pages were set up to contain CIS documents and provide resources to help others apply the investigated security techniques.

1.1 Major Results/Conclusions

The CIS found that although the participating agencies were quite different, they actually had comparable security needs and similar obstacles to their deployment.

1.1.1 Diversity

All agencies were home to a wide variety of computers ranging from hand-held personal digital assistants to Cray supercomputers, and they ran a wide variety of different

operating systems. This diversity prevented any one solution from being applicable, even within one agency. Because new security solutions are not being made available for old systems, there is a strong incentive for the government to upgrade its computer inventory. It is probably cheaper to buy a new computer and new security tool than it is to develop security patches for older equipment.

1.1.2 Varying Requirements

Even within an agency, the computer needs of a research scientist differ from those of a secretary. Thus, different computer security solutions are probably required. The trick is to be sure that all portions of an organization are covered to prevent attack at the unprotected “weak links.”

1.1.3 Remove Obstacles to Deployment

Security is not free. Resources must be allocated to purchase security tools, and their support, maintenance and upgrading constitute on-going expenses.

Because the computer world changes so fast, there is a real risk in deploying a computer security solution; the changing environment can make solutions obsolete over night. Accordingly, security administrators must look into the future to pick solutions that are likely to be viable for some time.

The culture of good computer security needs to be promulgated from the bottom up and from the top down in each organization. Active and enthusiastic management support is difficult to obtain in a downsizing environment, however it is crucial. Education of computer users, administrators and managers is essential.

1.1.4 Leverage Experience

It is difficult to implement a security solution without advice and some “hand holding.” Experts in all areas exist within the government, and mechanisms to make their expertise available (without undue burden) must be found. One successful method is via a resource Web page similar to the Kerberos page (<http://www.epm.ornl.gov/~jar/HowToKerb.html>) set up by a CIS testbed.

1.1.5 Reduce Barriers to Inter-agency Collaboration

The CIS effort pointed out how difficult it is to do seemingly simple things such as inter-agency money transfer. Synergism and cooperation should be encouraged, and the government infrastructure should be modified to make this possible.

1.1.6 Eliminate Clear-Text Passwords

The testbeds found that there exist good tools for almost all venues that can be used to eliminate clear-text passwords from networks. Kerberos and the Distributed Computing

Environment (DCE) provide large-scale single sign-on services that are especially appropriate in a UNIX environment. Digital certificates provide inexpensive, easy to use, strong authentication for web clients and servers (as well as encryption of traffic). Hardware tokens can be used in environments where higher levels of security are required.

In general, CIS found that the emerging public key infrastructure (PKI) and the use of identity and authorization certificates provides the most secure, flexible, and user-friendly approach to enhanced security. Even today, popular Web browsers support the use of identity certificates for digital signatures and for encrypted e-mail.

1.2 Recommendations

Some CIS recommendations apply to managers of Internet systems and applications:

- **Don't forget the basics.** The basic elements of computer security always apply. Systems must be up-to-date, patched, and well managed.
- **Focus on the requirements before the technology.** Technology is not always the answer to security problems. Start with your organizations needs, develop a security policy, and then determine the best way to implement it.
- **Remember, one size does not fit all.** CIS found that due to diverse environments and needs, there is no single magic computer security solution. Practice "divide and conquer." Split your organization into groups with similar properties and apply the best solutions to each piece.
- **Look for small victories.** For example, if you can provide a good means for remote users to access your system without sending clear-text passwords over the Internet, security will be greatly advanced. It is a big problem, so you probably have to solve it piece by piece.
- **Focus on high payoff applications.** CIS felt that it is easy, cheap and user friendly to secure e-mail and Web access.
- **Make use of available expertise and resources.** Don't reinvent the wheel. There are many sources for help. The government should determine a process for making them available to all agencies.

These recommendations are for the federal Internet community:

- **Encourage working level collaborative projects.** CIS brought together participants from many agencies at a working level. This cross-fertilization resulted in a valuable exchange of ideas, techniques, and actual code.
- **Remove institutional barriers to interagency collaborations..**

- **Encourage multiple solutions and interoperability.** Large-scale, government-wide solutions have not worked. The use of commercial-off-the-shelf software (COTS) products adhering to international standards results in cheaper, more up-to-date, and interoperable solutions that are tailored for each environment.
- **Institutionalize the CIS effort.** Open and active collaboration among agencies, as exemplified by the CIS effort, has been successful. It provides sufficient process that there should be some examination of how this collaboration model could be institutionalized in and among Government agencies.

2 Introduction

The federal Internet has reached a critical juncture in its development. What was once considered an "experimental" networking technology by many agencies has become the primary means for computer communication within agency enclaves as well as with outside organizations and the general public. There is a trend toward "Internet-enabling" critical business applications and making valuable or sensitive information available to the public. With the federal government's increasing dependency on computers and distributed systems, there is a greater urgency to develop and deploy information security technologies and practices.

In recognition of this need, the FNC PSWG was awarded a National Performance Review (NPR) Innovation Fund Grant. The project that came out of this grant was Collaborations in Internet Security (CIS) -- a cross-agency effort that compared and validated agency approaches to security and tested the strength of various technologies beyond the boundaries of closed agency networking environments. The ultimate goal of this project is to lay the foundations for a more trusted environment for communication among agencies and with the general public.

This report is an end-of-year summary of the work of CIS. It is written to help technical managers and those implementing information security technologies at a workgroup or enterprise level. This document describes the objectives, activities, conclusions, and recommendations of the CIS project through the end of Fiscal Year 1997. As additional work is recommended, this work has been labeled Phase I.

2.1 Background

The FNC and its PSWG were created in 1990 at the time of the High Performance Computing (HPC) Act to coordinate networking research and operational aspects of the networking infrastructure that supported other computer research. As the Internet expanded in reach and scope, the PSWG expanded its reach and focus from strictly research activities and infrastructure to the needs of less specialized federal agency security requirements.

In 1993, under the auspices of the Office of the Vice President, the National Performance Review charged the FNC with the task of formulating a federal Internet security plan. The resulting effort represented a first attempt to synthesize Internet security research into a coherent set of actions aimed at improving the security of operation of the federally funded portion of the Internet and to present a strategy for raising the state of practice in Internet security.

Since the development of the federal Internet Security Plan, the PSWG has continued to work with various organizations to augment the plan with details for implementing the plan's recommendations. To this end, the PSWG prepared a proposal in response to the Government Information Technology Services (GITS) Working Group's call for proposals under the Information Technology Innovation Fund Pilot Program. In February

of 1996, the PSWG was awarded a grant to CIS that brought together experts from FNC agencies and national laboratories to focus on key aspects of information security technology.

2.2 Project Goals and Scope

Three primary, over-arching goals guide the CIS effort:

- Develop a new and sustainable process for designing, integrating, and deploying information security technology that is interoperable at all levels of the federal Government.
- Set a precedent for future federal and commercial sector collaboration in the security networking sector, by employing a process geared toward openness and inclusiveness.
- Improve awareness among FNC agencies, other federal organizations, and commercial and academic sectors regarding the critical importance of network security and the range of options/technologies available to meet specific needs.

At a more tangible level, this effort will result in the identification of the security requirements, concerns, and objectives of all participating networks, as well as the corresponding existing or emerging technologies and applications that address those requirements. Testbeds will be established to implement these technologies, with the results being fully analyzed and documented. From these reports, will come security technology recommendations for implementations within participating agency networks and for future Internet security collaborations.

The effort will also include development of a laboratory accreditation program for testing and certifying Internet security software and systems. This process will be modeled on the National Institute of Standards and Technology's (NIST's) National Voluntary Laboratory Accreditation Program that tests and accredits systems and products for private sector and federal users.

3 Testbed Findings

A number of technology areas were investigated under the CIS project. They are Kerberos; FORTEZZA; Advanced Authentication; Public Key Infrastructure; Incident Response Testbed; Security Testing and Certification; Privacy, Digital Signatures, and Secure Messaging; and Secure Web. Clearly, several of these technology areas intersect and have had impact on the direction and conclusions of other testbeds.

The testbed efforts were guided by simple and direct principles:

- **Openness** — Invite commercial, academic, and government communities;
- **Multivendor** — Encompass technologies from a variety of federal and commercial developers;
- **Commercially-Used Technologies** — Test technologies currently in use in the commercial environment as well as new federal research and development (R&D) products;
- **Proactive** — Provide solutions based on current and anticipated Internet technologies/uses;
- **Privacy Considerations** — Use the appropriate privacy mechanisms at all levels.

3.1 Kerberos Testbed

Kerberos is a suite of client and server applications that was developed and is maintained by the Massachusetts Institute of Technology (MIT). It is a network security service that provides for strong authentication over unsecured networks. Kerberos can provide for mutual authentication between two parties, called principals, usually a client and a server using a third party trust model. Kerberos normally uses a software-only method of authentication, so it relies upon "something you know," a password, to provide this authentication. With modifications, it can provide additional security by using hardware devices, such as smart cards. The power of Kerberos is that it can identify a user to a remote computer without sending that user's password over the network in clear text.

3.1.1 Background

To use Kerberos securely, a Kerberos client must be installed on the user's local computer. Kerberos uses the concept of a ticket, that can be employed only by the user and only for a specific server. This local client requests tickets from the Kerberos server, known as a key distribution computer (KDC). The KDC normally sends back a ticket encrypted such that only the user can decrypt and use the ticket. The local client accepts the user's password from the keyboard, and uses it to decrypt the ticket. Since only the client using the user's key can decrypt the ticket, the ticket is useless to anyone else. Since symmetric encryption keys are usually used, and exposing one exposes the other

(with DES they are identical), it is extremely important to locate the KDC in a secure environment and to protect it carefully. To avoid requiring the user to enter in the password for every ticket request, the concept of a ticket that can be used to obtain other tickets, a "ticket granting ticket" (TGT), is normally used. Tickets have many properties, including an expiration time.

A detailed explanation of Kerberos can be found in the book *Network Security, PRIVATE Communication in a PUBLIC World*, by Charlie Kaufman, Radia Perlman, and Mike Spencer (Prentice Hall, 1995).

The computer user must not share his password or else security will be compromised. To counteract this, some implementations of Kerberos replace (or augment) the password with a token such as a smart card.

Kerberos has been deployed in large organizations, especially university communities such as MIT. Although security flaws have been found in Kerberos over the years, no significant break-ins have been reported. The bottom line is that if a user is content to use the services provided by Kerberos, then it provides a user-friendly method for once-a-day sign-on for networked resources while offering enhanced security.

The Kerberos client software consists of many parts. One component handles tickets:

- *kinit* asks the user for a password and securely gets the Kerberos ticket.
- *klist* displays the user's current tickets and their properties.
- *kdestroy* lets the user destroy unneeded tickets.
- *kpasswd* lets the user change a password on the KDC without sending it in clear-text over the Internet.

The more useful group of tools replaces the ordinary, UNIX-like remote communication services:

- *ktelnet* provides encrypted *telnet* sessions. However, any X sessions started by this connection will not be encrypted.
- *kftp* provides file transfer services without clear-text passwords. The file contents can be encrypted.
- *krlogin* is analogous to the usual *rlogin* function and works like *ktelnet*.
- *ksu* lets a user become superuser securely.
- *krcp* provides a secure way to copy files from one computer to another.
- *krsh* provides the secure analog to the UNIX *rsh* command.

There are also other utilities for remotely managing the KDC and for interoperability with Kerberos Version 4.

3.1.2 Purpose

The Kerberos testbed was created to answer the following questions:

- What sort of security solution does Kerberos provide?
- How does one obtain and install Kerberos?
- How do you deploy Kerberos across a large organization?
- How do the different versions of Kerberos interoperate?
- How does Kerberos interoperate with other security solutions?
- Why isn't Kerberos the whole security solution?

3.1.3 Relationship to Other Testbeds

Kerberos is probably the best protocol for securely transferring an identity across the Internet without the use of additional hardware. Kerberos has been integrated with hardware tokens to provide additional security, or, the Kerberized login can be used to replace other protocols such as in the secure shell (SSH). Microsoft announced that Kerberos 5 will serve as the basis for security in the next version of the NT operating system.

Because Kerberos can be used as a basis for authentication and the means for encrypting network services, the results from this testbed interact with those of the FORTEZZA and Advanced Authentication Testbeds.

3.1.4 Activities of the Kerberos Testbed

3.1.4.1 Installing Kerberos

The first step was to deploy Kerberos at testbed sites to gain experience in this process. Unfortunately, Kerberos is a large and complicated program and building it successfully on various platforms using different compile-time options is not simple. To codify this knowledge, participants created a Kerberos Web Site (<http://www.epm.ornl.gov/~jar/HowToKerb.html>) that leads the user through this process. This web page has had over 3000 visitors and is now linked into the official MIT Kerberos page (<http://web.mit.edu/kerberos/www/index.html>). Once Kerberos is built, the KDC and Kerberized resources must be configured, and the testbed web page provides step-by-step instructions for these procedures.

3.1.4.2 Diversity of Kerberos

Kerberos has been around for years. It was originally defined and developed at MIT, but the development continues with many other organizations making changes and adding new features to the MIT reference implementation. Over the years, many vendors have taken code and used it as a basis for security in their specific application(s), or with enhancements to provide a total network security solution. Others have "Kerberized" (added Kerberos security to) specific applications.

The Kerberos request for comments (RFC) define the Kerberos network protocol but, they do not define a programming Applications Programming Interface, nor do they define how each system will store needed information on the local machine. Also, there are two versions of Kerberos in current use, Version 4 and Version 5, which have partial interoperability.

Export control issues have also been a factor in Kerberos development. Many vendors have been reluctant to use Kerberos because of these issues, or they will only use its authentication features, and not provide the full range of Kerberos features that provide encryption. Interoperability between the diverse set of Kerberos implementations remains an important issue that is being tracked by testbed participants.

Kerberos Version 5 fixed some of the problems with Version 4, by adding pre-authentication, better cross-realm authentication and by implementing the forwarding of credentials. Renewable, postdated, and proxyable credentials were also added. Some implementations of Kerberos Version 5 will also support many of the Version 4 applications.

One of the largest and oldest products to use Kerberos is the Andrew File System (AFS), developed at Carnegie Mellon University (CMU). AFS is now available as a commercial product from Transarc. It was developed at about the same time as Kerberos Version 4, but it uses Kerberos mostly for internal purposes. It does provide some Kerberized applications, but only to allow the use of AFS with them. An example is a file transfer protocol (FTP) daemon that will allow you to access AFS via FTP. Unfortunately, this daemon requires that you send your clear-text password across the Internet, so the principal advantage of Kerberos is not used.

The DCE from the Open Group is one of the newer products that uses Kerberos. The DCE security service is based on Kerberos Version 5, and all the other DCE services use it for their security. DCE was designed to provide secure, distributed applications using secure remote procedure calls (RPC). Its chief feature, Distributed File System (DFS) is similar to AFS, but uses the DCE-based services. The latest version of DCE comes with some Kerberized applications, *rlogin* in particular, but like AFS does not have the full suite of applications. The DCE security is compliant with the Kerberos RFC and can work with other Kerberos implementations using only the Kerberos portion of DCE, but it does not support Version 4 applications or protocols. The Open Group is also working on a Java version of Kerberos that includes Netscape and Microsoft Internet Explorer

(IE) plug-ins. DCE is now a component of, or is available on almost every major UNIX computer system. IBM, HP and DEC have integrated it into their products. Silicon Graphics, Incorporated (SGI) and Cray also offer it as a product that can be purchased. Third parties have implemented DCE for Solaris, as well as for Windows 95, Windows NT and the Macintosh. Unfortunately, there is about a one year lag between the time the DCE reference version is made available to the manufacturers and the time it is available to the user. This results in less interoperability.

Deleted: ¶

Solaris comes with some of the Kerberos Version 4 clients, and a Kerberized Network File System (NFS). To use them, the other services must be provided using some other package.

Cygnus produces the *KerbNet* product and recently made the source code freely available. This is very close to the MIT Version 5, and it interoperates well with other Kerberos implementations. The *KerbNet* product includes PC and Macintosh *telnet* clients.

Cyber Safe and Open Vision are two other companies that sell total network solutions based on Kerberos. Testbed participants have not had much experience with either.

Platinum Technologies has a PC and Macintosh product called PC-Enterprise that uses Kerberos Version 4 and Version 5. It allows access to AFS and DFS, and has a Kerberized Post Office Protocol (POP) server that will work with Eudora Pro, but it too does not have the full suite of Kerberized applications.

Microsoft has stated that NT 5.0 will use Kerberos Version 5 for its primary authentication mechanism. They have also said they would be compliant with the RFCs and will interoperate with other compliant Kerberos implementations. If this NT implementation is done correctly, it could greatly enhance the image of Kerberos, by bringing it to the mass market, and by giving applications developers a foundation on which to build Kerberized applications on NT.

Kerberos has been added to web browsers as well, but not in either of the two major commercial browsers, Netscape or Microsoft's Internet Explorer. When Microsoft introduces NT 5.0, it is likely that Kerberos will be integrated into Internet Explorer for use in intranet environments.

Gradient Technologies has a number of products that are add-ons to the current browsers, or that use the security of the current browsers to add DCE authentication and authorization.

3.1.4.3 Cross-Realm Kerberos Authentication

A Kerberos realm (or a "cell" in DCE terms) is comprised of the user and service principals, and the KDC, which serves as the third-party in the third-party trust model. If one is a member of a Kerberos realm, and needs access to resources in another realm, one can either become a member of the second realm and obtain a ticket for that realm, or one

can create an inter-realm trust and perform cross-realm authentication. It is very awkward to maintain an identity in two separate realms at once, so in applications where a single application must be distributed across realms, cross-realm authentication looks attractive.

A principal in one realm can authenticate to a server principal in another realm if the two realms are registered in each other's database. Additional realms can also be involved. This is one of the attractive features of Kerberos Version 5.

The cross-realm authentication can be used between organizations, a model pursued by the Department of Energy (DoE) Energy Sciences Network (ESNet). There are more policy and procedural issues to this than technical issues. As compared to the traditional password authentication, one organization is vouching for the authentication of the user, and for the policies used in providing that authentication to the other organization. This mechanism works best where the realms are each large. First, it cuts down the administrative burden and the number of tickets needed for a transaction. Second, it is easier to place trust in a large realm because it probably has stronger controls on its KDCs and on the identification of its users. In such situations, it is necessary to have clearly stated administrative policies for each of the realms that are involved. One such policy for secure servers was created by an ESNet task force led by John Volmer of Argonne National Laboratory (ANL).

There are technical issues as well. The Kerberos RFC 1510 is vague on how the authentication path is determined by the client and server, that are used to determine if it is acceptable to have intermediate realms involved in the authentication. The DCE and MIT reference implementations do this differently, and thus have an interoperability problem when a Kerberos realm and DCE cell participate in cross-realm authentication.

The DoE ESNet Authentication Task Force (ATF) has a solution for this, and it has been incorporated into the MIT reference port. The ATF also had discussion with the Open Group and HP on solving this problem, as well as input to the Internet Engineering Task Force (IETF) for the next version of the Kerberos RFC in an effort to improve the situation.

The Oak Ridge National Laboratory (ORNL) and ANL Kerberos realms were set up to perform cross-realm authentication. In addition, the ANL realm was a member of a large ESNet DCE cell. Participants were able to transparently log in to each other's realms using these mechanisms, and were even able to use a local Kerberos ticket to obtain access to resources on a remote DCE cell.

3.1.4.4 Misuse of Kerberos

At the Spring 1997 IETF meeting, the Internet Architecture Board (IAB) session emphasized security. The board stated that the one thing that should not be done was to send clear-text, reusable passwords across the network. This is an obvious point, but many are still doing it.

Kerberos has long been viewed as an excellent network security system that uses a password to authenticate on the local system, but does not require the password to be sent over the network. Its primary function is to protect the password from exposure over the network. But in many situations when Kerberos was not on the local system, this authentication has been done over the network, and the password sent across the network in clear text. This is clearly a violation of the Kerberos principles, but continues to be overlooked. One example is a popular terminal server vendor that has implemented Kerberos, but requires sending a password over the phone lines to the terminal server. Another example is Transarc's AFS, which uses an older version of Kerberos to get tickets, but does not protect the password when logging in over the network.

Until recently, DCE had the same problem. If a DCE administrator had to log in over the network to do system administration functions on DCE servers, the administrator had to send the root password or cell-admin password over the network in clear-text.

3.1.4.5 Kerberos Version 5 and Multitiered Environments

For Kerberos to be an effective security solution, it needs to be implemented correctly, and it needs to work in multitiered environments. Only Kerberos Version 5 Kerberos Version 5 has the ability to forward credentials rather than passwords; Kerberos Version 4 did not have this capability, and was thus limited to client-server environments. Some developers have tried to extend this, but many of these extensions have introduced additional security risks and are not standardized.

The Open Group's DCE and DFS uses Kerberos Version 5 internally, in much the same way AFS does. DCE internally did allow for multitiered applications and delegation of authority. But only in their latest release do they provide some of the Kerberos applications that allow for sending forwarded credentials instead of passwords to allow for remote authentication.

The DoE ESNet community has been very active in this area. Using the MIT implementation of Kerberos Version 5 with DCE, the ESNet community has the ability to use a Kerberos Version 5, forwarded credential to get a DCE context and/or an AFS token. This community continues to help influence the Open Group and the vendors to toward interoperability. ESNet has also integrated AFS into their environment with a goal of single sign-on.

3.1.4.6 Kerberos and DCE as a Foundation for Security

Kerberos can be combined with other security technologies, and can form the basis for authentication. There are two IETF drafts, *PK-INIT* and *PK-CROSS*, which use public-key technology with Kerberos. *PK-INIT* uses public key technology for the original user authentication, rather than having a password. DCE 1.2.2 has implemented *PK-INIT*, and has included a smart-card API that can be used with *PK-INIT*. *PK-CROSS* deals with using public key technology with cross-realm authentication.

There is also a draft document on how to use FORTEZZA authentication with DCE, and thus Kerberos. The merging of these two technologies will benefit all. DCE's widely installed base (it is bundled with IBM's AIX and HP's HP-UX systems, for example) will benefit from improved security, and the FORTEZZA technology will have a much wider application base helping to offset its cost.

Both PKI and smart-card technologies can be used with Kerberos. Entrust has been talking to the Open Group about integration with DCE. SecureID can already be used with Kerberos due to the efforts of Ken Hornstein of the Naval Research Laboratory (NRL). A lack of application support and reluctance of software vendors to eliminate the password problem has caused delays in these implementations. DCE can use these technologies as well as FORTEZZA for authentication.

3.1.4.7 Kerberos Deployment for Large Installations

The true benefit of Kerberos comes from a fully deployed environment in which software is installed and used for all network transactions. This nirvana of computer security is seldom realized. Deployment of a "Kerberized" environment throughout an organization is usually more challenging than porting and installing the software itself. There are several technical and political hurdles involved such as education, large-scale installation, and support.

The Army Research Laboratory's (ARL's) Aberdeen Proving Ground (APG) site has deployed Kerberos Version 5 to about 2000 users and 800 UNIX hosts over a 2.5 year period (about 4500 Kerberos principals). Although there are several academic sites with over 20,000 principals, deploying Kerberos across a government installation is different. The majority of the UNIX hosts were Sun (SunOS and Solaris) and SGI-Irix machines with a small number of various other architectures such as HP-UX, IBM AIX, Cray Unicos, Pyramid OSX, and DEC Ultrix.

A large portion of the deployment effort went into porting Kerberos to these UNIX platforms. At the time of initial deployment, the MIT Kerberos distribution (Version 5-beta 3) was not as clean and portable as is it today. Changes were made to existing applications to make the transition for users easier such as:

- Changing *kinit* to always request forwardable tickets;
- Making ticket forwarding the default for Berkley Standard Distribution (BSD) commands;
- Building *kinit* into an OS version of */bin/login* and *ftpd*.

Software installation and distribution at ARL was made easy by use of an existing remote distribution (*rdist*) infrastructure. Kerberos was made available to all hosts via a nightly distribution of standard software packages. Standard *inetd.conf* and */etc/services* files were made available as well as text files that included the entries to append to these files.

A sufficient number of system administrators must be given "add" privileges for running *kadmin* to get *krb5.keytab* files for installation of application servers.

Most Kerberized applications run on different port numbers than their un-Kerberized counterparts. This situation makes the transition to a Kerberized environment more flexible because you can run both Kerberized and un-Kerberized services in parallel for a period. The strategy employed at the ARL was to install Kerberos clients and application servers on all UNIX machines and gradually phase out un-Kerberized services. This maintains the benefit of Kerberos application functionality (such as encrypted *rlogin/telnet* connections) while being flexible in the transition of user's day-to-day operations. Obviously, the full benefit of Kerberos is not realized until all un-Kerberized services are discontinued. This transition should be made as quickly as possible.

A short course including an overview of the Kerberos protocol and the administrator and user command suites was given to site system administrators. A site policy for system administrators was established to require remote access of privileged accounts (*/bin/su*) to be done via an encrypted *rlogin* or *telnet* session. This had an immediate benefit of keeping plain-text root passwords off the network. Creation of Kerberos accounts became the standard procedure for getting an account on most large servers. In some cases, passwords were set only in the Kerberos database, not in the host's password file. This makes password management easier for users since there is only one centrally managed password required for access to multiple hosts.

A large obstacle in the initial deployment and early training was resistance by administrators and users to accept the additional workload and inconvenience of using Kerberos software. Many users and administrators claimed that would preclude Kerberos from being a viable security solution on their hosts. Most of these claims were based on a lack of resources to implement the additional security on their hosts, the main complaint being a lack of time. Other claims were based on misunderstandings of Kerberos, UNIX, and/or security in general. Recent media attention has raised management awareness of the need for security. Convincing management to expend the necessary resources is a more difficult issue.

A successful strategy employed at ARL for deployment of Kerberos was to bring all new machines online in Kerberos-only mode. In this manner, access that users once had is not taken away, but the additional access procedures are masked behind a new computing resource. Users will resist, but will eventually realize the need and ease of additional security measures.

User education is an essential element in a successful deployment of Kerberos. Effective training sessions have included a background knowledge of basic networking principles, client-server concepts, examples of security threats (sniffers, TCP hijacking, etc.), real examples and statistics of site security incidents, demonstrations of software usage, and handouts of basic Kerberos commands and troubleshooting. Users are more accepting in a learning environment where they can understand the need for security rather than being forced into usage without adequate preparation.

KDC workload for up to 4000 principals is still trivial. A Sun ELC or Sparc Classic has more than enough horsepower to serve a site of this size. Daily maintenance is minimal on the KDC itself (backups and careful systems administration), and failures are extremely rare. At least one backup KDC is highly recommended for larger deployments, mostly for network failures. Network bandwidth requirements are low and connectivity to high-bandwidth networking (FDDI, ATM, HIPPI, etc.) serves mainly to increase routes to KDC. Database propagation takes less than 30 seconds over Ethernet.

Kerberos is dependent on a few core services such as time, accurate name service, host OS, network connectivity, and power. Most Kerberos errors are the result of failures in these core services. About 95 percent of all reported Kerberos problems are not Kerberos problems. For example, the failure of a Kerberized *rlogin* will generally be reported as a "Kerberos problem" even if the remote host or network is down. A list of common errors, causes, and solutions is helpful for administrators and users. Kerberos error messages are usually verbose and can be confusing to the uninitiated.

Maintenance of a Kerberos realm consists mostly of user support and accounts management. Password management can be accomplished by use of the "policy" mechanism within *kadmin*. This feature sets basic requirements for acceptable passwords. Attributes of a password policy are: maximum lifetime, minimum lifetime, minimum length, minimum classes, and history. "Classes" are sets of characters such as lowercase, uppercase, numbers, punctuation, and all other characters. History is the number of previous passwords from which the user cannot select. The *kpasswd* command conveniently explains password requirements to the user. Many policies can be created and applied to different principals. A policy named 'default' will be applied to all new principals. Another useful password management feature is the definition of a "dict_file" in the *kdc.conf* file. The *kadmin* daemon will check all new passwords against the defined dictionary and not allow matches.

A successful deployment of Kerberos will continue to provide user-friendly access for authorized users while limiting access enough to manage the risk associated with a site's network connectivity. By requiring Kerberos authentication for primary access mechanisms (*telnet*, *rlogin*, *rsh*, *ftp*, etc), you can reduce a large amount of risk. Complete risk avoidance is nearly unobtainable and should never be used a reason not to take the first step.

3.1.4.8 Non-UNIX Kerberos Implementations

In September 1997, Cygnus Solutions released a free implementation Kerberos Version 5 including clients for Windows 95, Windows NT, and Macintosh. This implementation, *Kerbnet*, comes with source code, precompiled binaries and excellent documentation. For the first time, it is possible to set up a KDC on an NT platform that offers some security advantages over UNIX.

Kerberos Testbed participants tested the Kerbnet code on Windows NT (both client and KDC), Macintosh (client only), and AIX 4.2 (client and host daemons). It all worked very well and was much easier to set up than it is using the MIT source code. So, unless one

requires extensive customization, we suggest trying Kerbnet first to see if it meets your organization's requirements. CIS member Doug Engert (ANL) created an alternate Windows 32-bit implementation of *krlogin* and ticket tools.

3.1.5 What's Missing?

Kerberos seems to offer a highly secure, one logon solution for access to networked resources. Some of the institutional resistance to Kerberos conversion has been outlined already. However, another factor is that there are no full suites of Kerberos client tools for Macintoshes and PCs. Kerberos is targeted at UNIX systems to avoid having to integrate the code into the graphical user interface (GUI) environments. For example, the Windows 32 and Macintosh client versions of Kerberos have only *telnet* and ticket management tools.

So, in an environment where PCs and Macintoshes prevail, it is difficult to implement Kerberos without removing important features from the users such as *ftp*. If a user must send his clear-text password over the network to transfer files, is it worth the effort to provide a *telnet* that does not send passwords over the network? Security policies and frameworks must be developed to eliminate such inconsistencies.

3.1.6 Next Steps

Kerberos is gradually spreading to new platforms, and new tools are appearing for existing platforms. If Kerberos Version 5 really appears in the next version of Windows NT, it will constitute a major vote of confidence for the Kerberos security mechanisms.

Kerberos Testbed participants believe that the most effective way to increase the deployment of Kerberos is to provide the information necessary so that site administrators can get things working without wasting too much time. To achieve this, we have created a "How to Kerberize your Site" (<http://www.epm.ornl.gov/~jar/HowToKerb.html>) web page. This page has proven to be useful to thousands of users since its inception, and participants are committed to keeping information on these pages maintained and updated.

3.2 FORTEZZA Testbed

The FORTEZZA Card, a Personal Computer Memory Card International Association (PCMCIA) circuit card that uses NIST/National Security Agency (NSA)-approved algorithms to provide network related security services, was developed by the NSA. When used with the proper applications and network security infrastructure, the FORTEZZA card provides several functions. These include: (1) authentication provided through the hash and digital signature capabilities, (2) confidentiality provided through the encryption algorithm, (3) integrity provided by the hash and signature algorithms, and (4) non-repudiation provided by the signature algorithm. NSA is also considering adding FORTEZZA implemented as software and in a smart card to its product line. The CIS focused on one Multilevel Information Systems Security Initiative (MISSI) product in particular, the FORTEZZA crypto card. The great power of the FORTEZZA approach is

its ability to service a wide variety of different applications (e.g., e-mail, authentication), the strength of its security design (it is a hardware token), and the supporting infrastructure being built by the Department of Defense (DoD) and the General Services Agency (GSA).

3.2.1 Background

NSA is currently developing MISSI products in support of the Defense Message System (DMS) program under the umbrella of the Defense Information Infrastructure (DII). The DMS was conceived as a replacement for Autodin and as a DoD-wide electronic-mail system designed to serve up to 2 million users. The DII is slated to share information with allies and coalition partners, and through the National Information Infrastructure (NII) with U.S. government agencies, universities, and industry. Consequently, the DoD will increasingly use the Internet to access external sources and to exchange electronic mail with external organizations. MISSI tools will be used to protect the DII. They are also tools that can be used in the NII.

The objective of the MISSI Program is to make available an evolving set of solutions that provide secure interoperability among a wide variety of missions that comprise the DII. Providing information security for the "network" is a system problem and a system approach must be taken to provide meaningful solutions. The MISSI program takes a system approach to this problem and it is characterized as an evolving framework for security in an open environment, driven by secure interoperability and the need to ensure adequate security for customer needs. In addition to the FORTEZZA card, MISSI products include the Secure Network Server (SNS), Certification Authority Workstation (CAW), Directory System Agent (DSA), Mail List Agent (MLA), and In-Line Network Encryptor (INE).

3.2.2 Purpose

A FORTEZZA testbed was developed to provide CIS participants with the capability to evaluate Multilevel Information Systems Security Initiative (MISSI) products as tools in their interoperability testbeds.

3.2.3 Relationship to Other Testbeds

MISSI products are applicable to almost all CIS testbeds. For example, FORTEZZA enabled *LockOut*, an identification and authentication software application package, has a direct relationship to the AA testbed. Also, FORTEZZA enabled Netscape Navigator, Enterprise Server, and Proxy Server have a direct relationship to the Secure Web Testbed.

3.2.4 Activities of the FORTEZZA Testbed

3.2.4.1 NSA Testing Overview

NSA has an extensive testing program focused on validating and verifying MISSI product requirements by functionality, security, conformance, performance and interoperability testing. Both the military services and MISSI product vendors also have test programs for operational and interoperability testing. MISSI testing is divided into four interrelated phases discussed below.

The first phase, Security Product Development Testing, focuses on the contractor/vendor product testing, along with creating all documentation needed for formal acceptance at a MISSI test site. This documentation includes the products test plans and procedures for functional, verification, performance, and security testing. Development testing demonstrates attainment of all identified requirements including those for security functionality, standards conformance, performance, and acceptable software engineering practices.

The second phase, MISSI Product Compliance Testing, determines if a product satisfies an identifiable set of MISSI functional, security, and interoperability requirements and is considered MISSI Compliant. The compliance testing consists of two major test areas: the Product Qualification Test (PQT) and the Product Interoperability Test (PIT). The PQT focuses on how the product meets its requirements, including security testing and product requirement testing. The product is pairwise tested with associated products and a MISSI testing infrastructure in the PIT.

The third phase, Generic Security Solution Testing (SST), includes integration and test (I&T) and system testing. Products are integrated with each other to form a generic functional solution. System testing is conducted at I&T completion to ensure that the documented requirements are fulfilled and that the solution operates functionally as a "system."

The fourth phase, User Security Solution Testing, includes integration of MISSI capabilities into the user environment. This testing is generally conducted at the user's site.

Beta testing, may be conducted in addition to the phases above to verify user functionality and to provide feedback on a specific product or solution's viability, field ability, and effectiveness. NSA has established a Beta Test Working Group (BTWG), the purpose of which is to beta test MISSI-based security products as they become available, and to share findings with other working group members. The current focus is FORTEZZA-based products: specifically, FORTEZZA-enabled Netscape clients and servers on a wide variety of operating systems. These "new" products are beta tested before release to the public at large. Testers are responsible for identifying (and solving, whenever possible) problems as they arise, and reporting those problems to the appropriate developers to ensure a smooth development process and public transition.

Approximately 185 Netscape clients and 22 Netscape servers are involved in MISSI beta testing. Testers are military service and civil agency personnel.

3.2.4.2 MISSI Installation Team

To promote rapid integration of MISSI products into operational environments, NSA has created a MISSI installation team. The MISSI installation team has installed Group Technologies, Spyrus, Mytronics FORTEZZA cards, and over 18 FORTEZZA enabled products. They have also installed various Litronic, Spyrus, and Argus PCMCIA card readers in addition to Graystone and other internal card readers. This was done primarily on IBM compatible computer platforms although some work was done on the Macintosh. Installations were performed on platforms that used a variety of operating systems including DOS, Windows, and Macintosh.

A sample of installation issues include:

- The proper Small Computer System Interface (SCSI) drivers are not always available for external SCSI PC card readers.
- Software applications are continually evolving. Sometimes the new releases do not have a version number.
- Some notebooks come with the PC card ports disabled. A special utility from the manufacturer is required to turn them back on.
- Memory and IRQ conflicts are very common when installing internal card readers.

3.2.4.3 CIS Testing Overview

Initially, 14 CIS participants indicated a desire to use the FORTEZZA card and FORTEZZA enabled software applications. A survey was conducted to determine the computer system configurations and the needed application software of the 14 FORTEZZA testbed members. The objective of the survey was to obtain system configurations and then provide all 14 participants with appropriate FORTEZZA "Test" cards, PCMCIA card drivers, and application software. This was needed to develop the tools for testing the FORTEZZA product and applications. Survey results indicated the following:

- Participant computer platforms include IBM compatible PC, Suns, Macintoshes and HPs.
- Operating systems include Windows 3.1, Windows 95, Windows NT, HP UX 10, Solaris, and MAC OS.
- The members currently use a variety of e-mail packages such as Eudora, Lotus cc:Mail, and MS Mail. Almost all members were interested in evaluating the

FORTEZZA enabled versions of Secret Agent, a file encryptor, and the Netscape web browser, in addition to the appropriate FORTEZZA enabled e-mail application.

As a CIS participant, NSA required no funding to support CIS-related operations. Funding was required to provide PCMCIA card readers, and required software to other CIS participants, since the FORTEZZA card requires modified application software to provide security services for data integrity, authentication, non-repudiation, and confidentiality. The FORTEZZA Testbed leaders estimated that those testbeds that need equipment would require at least two FORTEZZA cards and an external PCMCIA card reader. In addition, application software (the type depends on the individual user) would be needed at each testbed. The estimated cost to CIS to support six testbeds was approximately \$10,000. Unfortunately, some products needed to establish the testbed were not procured.

NSA provided 11 FORTEZZA cards, programmed with "Test" certificates, to CIS participants. Some testbed members, such as the ARL, obtained additional FORTEZZA cards via other means. Other members, such as ORNL, obtained additional MISSI products for use in their own programs. Unfortunately, application software and the PCMCIA card readers were not purchased therefore only those members that already had readers could participate further. A limited amount of application software was obtained and provided to interested participants.

3.2.4.4 ORNL FORTEZZA Testbed

ORNL participants have FORTEZZA enabled Lotus cc:Mail working and have installed a CAW. They have used the CAW to program FORTEZZA cards with the appropriate keys, privileges, and certificates. ORNL has also installed a Directory Server Agent (DSA) and a Directory User Agent (DUA). The DSA is a X.500 based MISSI component supporting e-mail transactions between workstations. The DUA is an interface making the DSA available to the user. ORNL has also successfully tested the FORTEZZA-enabled web browser (Netscape Navigator 3.0 for FORTEZZA) using Windows 3.1. They will install and test the Windows 95 version later this year.

3.2.4.5 ARL FORTEZZA Testbed

There is very little commercial software or vendor application support for FORTEZZA on UNIX platforms. This is obviously not the largest market, but is extremely important for a well integrated, secured computing environment. Part of the ARL testbed focused on developing software that uses FORTEZZA cards for various security services.

The ARL used the following UNIX platforms:

- Sun Ultra 1, Solaris 2.5.1, Sun SBus PCMCIA card reader
- SGI Indy, Irix 6.2, external Adtron SCSI PCMCIA reader

The Sun platform required a special alpha-release driver from Litronic for the PCMCIA reader to function. This inexpensive, readily available device from Sun should be better supported. Specifically, this driver only worked with the Litronic Multi Access Crypto Interface (MACI) library version 1.52a (that was not available in source code). The second platform used the Adtron PCMCIA SCSI card reader sold through SGI with FORTEZZA drivers and libraries for Irix 6.2 and later.

Documentation on the Crypto Interface (CI) libraries is good and readily available with the SGI distribution and at the FORTEZZA web site (<http://www.armadillo.hunstville.AL.US>). Several sample applications can be found at this web site as well as documentation on related topics such as Certificate Revocation List (CRL) file formats and Compromised Key List (CKL) file formats. A draft Cryptoki (PKCS #11) implementation for FORTEZZA and an ASN.1 parsing library (for X.509 certificates) are also included.

ARL wrote simple applications (e.g. file encryptors) and a remote file copy to test crypto operations throughput with various configurations of cards, drivers, and devices. The Solaris environment required linking with the *libsocket.a* library and Irix required the *libds.a* library.

The highest encryption rate, obtained with the Group Technologies Card, MACI libraries, alpha driver, Sun SBus PCMCIA reader, on the Sun Ultra running Solaris 2.5.1, was approximately 4Mbps (encrypting 4.6 Mbytes of text). The Spyrus cards were about 12.5 percent slower for most functions and had a mysterious delay (approximately 5 seconds) on the first call to the card. Random number generation was tested and achieved at a rate of approximately 28 kbps.

3.2.5 What's Missing?

- FORTEZZA is relatively expensive because it is implemented as a hardware token and requires a card reader. Less expensive FORTEZZA implementations, such as smart cards and software are needed for applications requiring less security assurance.
- Additional commercially available software applications that have a FORTEZZA capability are needed for the different computer platforms and operating systems.
- An infrastructure is needed that supports deployment of security products and services on a large scale. MISSI is developing such an infrastructure. Other security approaches should also develop and implement such an infrastructure.
- A better understanding of certificate revocation is required.
- A wider exposure of FORTEZZA is needed throughout federal departments and agencies.

3.2.6 Next Steps

The following actions should be conducted immediately for testbed members:

- Purchase and install PCMCIA card readers and appropriate FORTEZZA enabled software on all participants workstations.
- Install FORTEZZA-enabled Netscape server(s).
- Participate in the Beta Test Working Group's Conference. See <http://beta.missilab.com/register> for more information.
- Obtain input from FORTEZZA participants so that a test plan can be developed.

3.2.7 References and Notes

3.2.7.1 The Network Security Initiative and FORTEZZA Directions

Within the past year or so, NSA, Defense Information Systems Agency (DISA), and the services realized that the security and operational architecture towards which they had been working was not feasible. The original architecture envisioned providing writer-to-reader security (only) from trusted desktops for all classified and unclassified messaging over the unclassified Defense Information Systems Network (DISN). Each level of classification was to be provided appropriate protection using NSA products from the MISSI program. However, the fact that obtaining such products was technologically impossible became obvious. That trusted workstations and servers were never going to be affordable, user-friendly, work with the growing variety of applications being used, and provide any protection for the transport backbone. Another approach was necessary.

It was also recognized that not only were the operational modes of our customers changing, the technology they were using was changing faster than feasible and affordable government-off-the-shelf (GOTS) security solutions could be developed. A broad range of functions were being implemented such as information push and pull, web browsing, and increasing volumes of individual messaging, all being COTS products operating over three classification levels of IP router broadband networks. These operational modes also brought out the need to consider security services beyond confidentiality, i.e., user authentication, authenticated access control, data and process integrity, proof of participation (origin and receipt). The complexities of security solutions in these environments mandated that a risk management approach be adopted that provides appropriate security for the information and services being protected. Emerging technologies just add to the problems of achieving security in this widely-connected networked environment. It became clear that the Network Security Initiative had to change. In the past the focus was on GOTS and Multi-Level Security (MLS) but now the focus has to be on COTS and secure sharing.

The Network Security Initiative has defined a network security framework that is based on:

- common security services;
- common security protocols;
- authenticated and controlled access to information and services;
- interoperable security-enabled COTS products, a global security management infrastructure; and,
- development of a new capability of attack sensing and warning.

Given the establishment of many local enterprise networks interconnected via network backbones, security services can be implemented at either the application layer or the network layer.

Applications layer security can:

- protect documents and files for integrity and confidentiality if a user-applied classification label is bound to the data;
- provide strong authentication of the sender;
- provide strong authentication of the recipient.

This approach relies on the operating system for invocation.

Network layer security can:

- protect packets or sessions;
- authenticate transfer of information;
- provide strong access control;
- be implemented workstation-to-workstation or enclave-to-enclave.

The emerging DII architecture is based on messaging and many other applications on system high systems for the Secret and Below Command and Control (C2) community (on SIPRNET), Top Secret/SCI Intelligence Community (IC) (on the IC Nets), and the Unclassified Community (on NIPRNET).

Although NIPRNET is viewed as the DoD unclassified network, it is so richly interconnected with the Internet that the two should be considered as one network. The C2 community operates with data that is Secret and below across SIPRNET and has established protected connections for unclassified data from and to NIPRNET. The IC Community operates with Top Secret and below data across the IC Nets and is in the process of establishing protected connections from and to the C2 community. Both of the classified networks have the transport systems protected by Type 1 cryptography that

provides both confidentiality and availability for those systems. Within each system writer-to-reader security (at the application layer) provides confidentiality or privacy independent of all intermediate systems and from insiders and outsiders, integrity of messages and attachments, authenticity of the sender, and proof of delivery to the intended recipient. All of these services are essential to achieving the ability to securely move messages and attachments between systems and/or classification levels.

Looking back, NSA provides the essential security services at the most efficient layer in the architecture using the technology available today and coming in the near future. The agency has provided the foundations for Type 1 cryptography for the strong confidentiality and availability at the network/transport layers and Type 2 cryptography (FORTEZZA today) at the application layer for the other security services. The FORTEZZA suite of algorithms provides appropriate security services for the foreseeable future. The original plan to use the FORTEZZA Plus card (Type 1 algorithms) for application layer security services was based on the availability of trusted platforms, which is not happening. Therefore, to transition to FORTEZZA Plus, which is also peripheral crypto, on today's untrusted platforms does not make sense. The level of security improvement does not warrant the cost and schedule delays to change all of the applications and other components necessary to support providing security solutions in the operational environments of today.

In response to numerous requests from DoD customers, NSA is expanding the FORTEZZA family of cryptographic services to provide users with varied levels of security services. NSA will support compatible Smart Card and software versions of FORTEZZA in addition to the existing PCMCIA and T2CSS Server Board versions. The current FORTEZZA PC Card is the size of a credit card and is the cryptographic token used by DoD for digital signature and encryption services at the desktop. Expanding the FORTEZZA family recognizes the varied needs of the national security community. Low assurance needs can be satisfied using the software or the Smart Card implementations, while higher levels of security assurance can be satisfied using the PC Card version.

No single policy statement can define where and when each security service can be used. Each organization must conduct an analysis of where the security robustness of hardware is essential and where the risks associated with the software versions are acceptable in light of security policies, desired security services, and the effect of local decisions on the overall wide area network security posture. Broadening the solutions NSA will support for its customers is neither a general NSA endorsement of software-based cryptography, nor the appropriateness of its use in any particular environment. In some specific applications, currently available commercial algorithms are being used where secure interoperability with industry is required. In the longer range, if the next generation of NIST FIPS algorithms are suitable, these could be adopted for Government use at all classification levels in a variety of configurations as available from the various commercial security providers.

3.2.7.2 References

For further information concerning the FORTEZZA card visit:

- <http://www.armadillo.huntsville.AL.US/>
- <http://beta.missilab.com/>

3.3 Advanced Authentication Testbed

Authentication must take place for every transaction, service request, or data access in a secure system, as it is the basis for further authorization decisions. Each type of request or access may require different levels or strengths of authentication. Authentication systems such as Kerberos provide "automatic" authentication between a user and service based on an initial transaction between the user and the Kerberos system. This initial authentication requires user interaction and must be at least as strong or stronger than is required by any further authentication requests.

The initial authentication must prove the identity of the user to a system, usually by testing a combination of the following:

- Something you know;
- Something you have;
- Something about you.

Traditionally, this is accomplished using only a password (i.e., something you know). However, this approach does not work well for remote access. Initial investigations show that network sniffers are being deployed on the Internet at an alarming rate, and many of them go undetected for weeks or months. Password guessing also poses a threat to traditional authentication measures due to ineffective password selection and management policies, and thus turns "something you know" into "something many may know."

3.3.1 Background

An understanding of one's computing and network environment is crucial to determine a appropriate authentication mechanisms. An authentication system must match the types of computers and networking equipment as well as their normal usage patterns.

A closed environment is one without connections to outside networks that could introduce traffic to the system. There is no access for authorized or unauthorized external users. A closed environment is well suited for small organizations that use computers and networks for internal purposes exclusively.

An open environment is one with outside network connectivity for which any user on any connected network can introduce traffic to the local network. This environment allows the greatest flexibility at a cost of higher security requirements. Authorized users can access protected resources from any location on any connected network and unauthorized users from any location can attempt to gain access to protected resources.

Organizations that must communicate with other sites and networks must open their networks to some degree. Filtering of offsite/onsite traffic, firewalls, and virtual private networks (VPNs) can offer a compromise between fully open and fully closed systems. Organizations that cannot afford to compromise the security posture of their entire network, but must still communicate to other sites may choose to build two separate networks. These may be connected by a firewall with a more strict rulebase, or be completely unconnected.

Patterns and requirements for usage can contribute to determining appropriate mechanisms. A security implementation that accommodates particular requirements generally involves additional software and/or hardware. A system such as Kerberos relies on trusted Kerberos client software installed on all user desktop machines. Smartcard systems usually require attached hardware readers and software. This may preclude access from non-participating hosts and sites such as conference terminal rooms, universities, and unsupported architectures or operating systems. Where access from such sites is required, one should assume that only limited software will be available and may not be trusted. Non-reusable authentication data should be used for this environment with the understanding that such a system is susceptible to attacks such as TCP hijacking.

Ease of use and security are usually mutually exclusive properties. Most advanced authentication systems require users to possess a device, that users must not lose or forget to carry and must know how to use properly. The implementation of such systems often comes at the expense of increased user education and a decrease in the speed of information access. A careful requirements analysis is needed to balance all factors for an organization.

3.3.2 Purpose

The purpose of the Advanced Authentication (AA) Testbed is to investigate technologies and methods for improving the strength of authentication in various environments. By analyzing the properties of various types of authentication, we can understand their strengths and shortcomings in different environments. The methods and algorithms used to accomplish the authentication exchange are also paramount in the overall strength of an authentication system. By selecting the appropriate algorithm and information type for a particular environment, the security of a system can be greatly improved.

3.3.3 Authentication Information (Auth-Info)

The purpose of auth-info is to prove to a service providing computer or process that a client computer, process, or user's claimed identity is authentic. Each type of auth-info has derivation and presentation methods. Traditional password system derivation is accomplished by user selection of a static, alphanumeric string. The password is presented to the system when the user types the remembered string into a keyboard. The derivation process has vulnerabilities with respect to poor password selection by users. The presentation process can be monitored and re-used due to the static nature of the auth-info.

Stronger auth-info types use different derivation and presentation techniques that employ token cards, card readers, biometrics devices, etc. The functions of each are different but the goal is the same. The actual implementations of such devices can be categorized as follows:

3.3.3.1 Challenge-Response/Time-Based Token Cards

Challenge-response tokens are hardware devices that a user must possess to derive auth-info (passcodes). These devices use challenge data from the process requesting authentication along with an internal algorithm and/or seed to generate the auth-info. The auth info is based on "something you have" (and possibly "something you know" such as the proper use of the device or the device password). The dynamic nature of the challenge and the difficulty in predicting the associated response without the token makes passcode sniffing and brute force attacks futile.

Time-based token cards derive pseudo-random passcodes based on time synchronization with the authentication verification process and an internal seed (preset, shared secret). Some forms of time-based tokens use a Personal Identification Number (PIN) to further seed the token generator providing proof of "something you know" (the PIN) in addition to "something you have" (the token card). Since the passcode changes periodically (once a minute is typical), it can be invalidated after its initial use by the verification process. Subsequent uses of the passcode will be treated as a replay attack and denied. The random nature of the internal seed plus PIN, combined with passcode invalidation, protects against unauthorized intrusion.

The presentation of auth-info is usually manual. A user must read the auth-info from the token and enter it as they would a traditional password on the keyboard. The tokens do nothing but generate passcodes. They have no data storage capabilities and computational functionality is limited to passcode computations only. Both systems become vulnerable when the token is lost or stolen. The verification system must be able to detect brute force attacks against the "something you know" (PIN and/or process) when the "something you have" is compromised.

These devices are well suited for remote access in an open environment. There is usually no requirement for special software on remote machines. Utilities such as *telnet* will suffice for accessing the host, which can easily be modified on the server side to present a challenge and/or verify auth-info in various ways. Since auth-info is valid only once per challenge/timestep, sniffers pose little or no threat to the authentication system. Cost, ease of use, token failure, and carrying the token around are the downsides of challenge-response/time-based tokens.

3.3.3.2 Smart Cards

These are similar to Challenge-Response/Time-based tokens but have additional functionality. These devices generally have interfaces that allow them to communicate directly to computer hardware, thus providing higher communication bandwidth and less user interaction. As with the token cards described above, a PIN is usually required to

activate or access certain functions adding the "something you know" (PIN) and "something you have" (smart card) to the authentication function of the system.

Smart cards generally contain microprocessors and memory for uses other than auth-info generation. Memory (volatile and nonvolatile) can store encryption keys and credentials for use in various authentication and encryption schemes. Processing power may be used for cryptographic functions such as passcode generation, data confidentiality, and data integrity. Tamper-proof construction can provide secrecy of algorithms and stored keys. Key recovery mechanisms are sometimes associated with smart cards to allow for data recovery in the case of user error and lost or malfunctioning cards.

Auth-info derivation can be done in several ways, including emulation of time-base and challenge-response tokens. A common method is sending "signed" challenge data or key exchange data along with a certificate list, using a common, trusted Certificate Authority (CA). This takes full advantage of the cryptographic processing capabilities of the smart card to create a stronger level of authentication.

Smart cards provide a stronger degree of authentication than the tokens described above at the cost of more limitations. Readers must be installed on all hosts making the per-user and per-host costs of deployment higher than token cards. Remote access is severely restricted due to hardware deployment requirements. Usage is generally easier than token cards. FORTEZZA is an example of a smart card implementation.

3.3.3.3 Biometric Devices

Biometrics is used to test "something about you." Auth-info is derived by using a data input source such as video camera, retinal scanner, or microphone to collect information. The data collected will be relatively static (biometrics can be changed by accidents, disease, etc.) and not necessarily secret (fingerprints can be copied from coffee mugs). A digitized representation of a biometric (fingerprint, voice characteristics, etc.) is presented to the authentication system for verification. The strength of this authentication is based on the ability to provide the correct data to the input mechanism or measuring device.

The digital representation of a biometric can be copied and inserted to an authentication system, but only the authentic user can generate the data via the measuring device. The unique physical characteristics of an individual must be the only way to present the auth-info to the system. The input mechanism must be trusted to use data properly and to disallow insertion of data other than via the measuring device.

Depending on the type of biometric, additional tests must be done to assure that the proper physical characteristics truly represent an authorized user. For example, fingerprint readers should check for temperature and/or pulse.

Biometrics, if implemented properly, provide excellent security, but can be extremely expensive. Measuring devices must be securely installed in physically controlled environments. Due to the nature of the data (public information, but not reproducible to

measuring devices by others), the devices must take an active part in the authentication algorithm to assure that the auth-info came directly from the measuring device. This may include physically secured wiring.

3.3.4 Authentication Protocols (Algorithms)

The types and meaning of the various authentication data produced by authentication devices must be used in a way that takes advantage of the data's properties to accomplish the goal of authentication. A protocol may be needed to provide means to transfer information securely across and "untrusted" network (one that assumes the existence of sniffers and spoofing capabilities). Replay detection, credentials storage, and key management may also be functions of the protocol. The protocol must be robust, tamper-proof, and well defined with user-friendly client interfaces. Encryption, hashes, and signatures are the building blocks on which authentication protocols are constructed.

Public key-based algorithms can securely transit sensitive authentication information as well as prove identity by mere possession or use of a private key. Key management functions, such as those found in a PKI, must be in place and suitable key storage must exist (e.g., in a smart card). Private key-based algorithms are usually easier to implement, but have less functionality in handling sensitive data.

The protocol brings together the users auth-info with the authentication system. Environmental considerations set requirements for how authentication exchanges must handle data and messaging.

3.3.5 Relation to Other Testbeds

Because authentication is the first step in accessing system and network services, activities of the AA Testbed have interest to most of the other CIS testbeds. Advanced authentication is of interest to the Secure Web Testbed as a means for strengthening authentication for resources available via HTTP. Both the Kerberos and the FORTEZZA Testbeds have strong ties to the AA Testbed as both technologies can be used for strong authentication. Finally, as the infrastructure needed by many advanced authentication methods are provided by a PKI, the activities of the PKI Testbed are related to the activities of this Testbed.

3.3.6 Activities of the Advanced Authentication Testbed

The AA Testbed addresses the issues of improving traditional authentication methods by augmenting or replacing them with mechanisms that can provide a stronger proof of authenticity. The project is based on currently available commercial hardware and software as well as public domain software. The goal is to incorporate the use of AA mechanisms into popular security packages and make them easy to use and administer. The results of this project are available to the software and hardware maintainers for possible incorporation into future distributions.

The computing environment of the DoD High Performance Computing Modernization Program (HPCMP) was chosen for this testbed given its non-unique requirements for openness and security in a widely-distributed network. The HPCMP consists of approximately 18 sites around the United States supporting multiple hardware architectures and operating systems that are accessed by scientists and engineers from various military sites as well as universities and private industries. The HPCMP infrastructure relies on the Defense Research and Engineering Network (DREN) for inter-site communication and the Internet for communication to participating non-DoD sites. The HPCMP required that One-Time-Passwords (OTPs) be used for access to program resources and the program would not support the use of specialized software that would contain encryption code, on non-DoD hosts.

It was decided that a combination of OTP and Kerberos would meet the authentication requirements while maintaining a user-friendly, automated environment. The Kerberos integration started with the work done by Cliff Neuman, Glenn Zorn, and Jonathan Trostle as outlined in their Internet Draft *Integrating Single-use Authentication Mechanisms (SAM) with Kerberos*. This draft details a framework for adding various external authentication mechanisms to Kerberos, which adds basic functionality to the clients (*kinit*) and moves the complexity of the specific authentication mechanisms onto the KDC. Policy decisions (such as if and when a principal must use additional authentication) are made by the KDC based on available data such as username, time client IP address, etc. A draft implementation of this framework is distributed in MIT Kerberos Version 5.

The one-time-password mechanism chosen for this testbed was the SecurID time-based token card with pinpad from Security Dynamics. Cards generate pseudo-random passcodes every minute and users enter a PIN onto the card to further seed the passcode generation. An algorithm was chosen based on the environment and available functionality of the client in the SAM framework. An algorithm developed Sandia National Laboratories (SNL) was selected due to its flexibility, strength, and ease of implementation within the SAM framework.

This algorithm adds a preauthentication (preauth) field to the traditional Kerberos AS_REQ message. The preauth field contains a SecurID code encrypted with the user's key that is derived from the password. The KDC will decrypt the preauth message with its copy of the user's key, and verifies the SecurID code. At this point, the KDC knows the user possesses the card, knows the PIN, and knows the Kerberos password. A loss or compromise of any one of these will prevent the compromise of authorized credentials.

The SAM sequence goes as follows:

1. Client (*kinit*) sends an authentication request (AS_REQ).
2. KDC receives the request and checks for preauth requirements for principal. If the `hw_auth_required` flag is set for this principal in the Kerberos database, the KDC determines what type of SAM mechanism will be used and creates an error message that describes the type of preauth required and how to obtain it. A

SAM_CHALLENGE message is generated and sent within as an AS_REQ error message that describes the specific challenge/SAM type required with the following generic parameters:

type	Defined list of SAM types
Flags	Describe how to use Single-Use-Authentication Data (SAD) <ul style="list-style-type: none"> < Use SAD as encryption Key [send AS_REQ encrypted with SAD] < Send encrypted SAD [send SAD to KDC in preauth field encrypted with client key] < Public Key encrypt SAD [not completely defined or implemented yet]
type-name	Text name for human purposes only.
track-id	Used to track multiple SAM messages.
challenge-label	Text to be displayed to user for challenge. [For SecurID there is no challenge label or challenge value.]
Challenge	Actual challenge value. [NULL]
response-prompt	Text prompt to request input from user.
pk-for-sad	Public Key to encrypt SAD. [undefined and not yet implemented]
Nonce	Nonce for challenge message.
Checksum	Checksum of challenge message.

- The client receives and decodes the SAM_CHALLENGE message, displays the challenge, and prompts for the SAD as instructed by the KDC. The client then processes the SAD based on the SAM flags and sends a response to the KDC. The client then processes the SAD based on the SAM flags and sends a response to the KDC. For SecurID integration, this involves getting the SecurID code from the user and encrypting it with the previously obtained password. This data is sent to the KDC along with a copy of several SAM_CHALLENGE fields in a preauth message of type SAM_RESPONSE.

4. The KDC decodes the message and checks the SAM_RESPONSE. The KDC must then send the passcode off to a verification routine specific to the type of OTP used (in our case Security Dynamics sd_check() routine). This routine may return success, failure, or insufficient data status. If successful, the KDC continues and generates a traditional AS_REP message using the key of the client for encryption. If the SAD is insufficient, the process can be repeated by sending another SAM_CHALLENGE with updated parameters. If a SecurID card is suspected to be out of time synchronization with the SecurID server, a subsequent passcode is required from the user. In this case, a second SAM_CHALLENGE message is issued with the prompt set to request the next passcode as in step 2.
5. The client will receive either a successful AS_REP message or another error message describing the reason for authentication failure.

Implementation work was minimal due to the generic functionality already built in by the SAM extensions. A simple change to the client libraries [src/lib/krb5/krb/preauth.c] was made to include a timestamp in the SAM_RESPONSE message for all SAM-flag settings. This fix was submitted and accepted by MIT for inclusion in their V5-1.0 patch level 1 release. This maintains the notion of a generic client interface that will work without modification for any OTP system designed into a Kerberos KDC.

Modifications were made to the get_sam_edata() routine [src/kdc/kdc_preauth.c] to call a new routine get_securid_edata() to build an appropriate SAM_CHALLENGE message for the request. The new routine uses static values for the SecurID mechanism or pre-computed values from a previous call to the verify_securid_data() routine that resulted in an insufficient data error. Similarly, modifications were made to the verify_securid_data() routine that resulted in an insufficient data error. Similarly, modifications were made to the verify_sam_response() routine to call a new function, verify_securid_data(), which simply acts as an interface to the sd_check() function provided by the Security Dynamics client API.

This same code base can be used with other OTP mechanisms such as S/Key (in precomputed mode) by replacing the SAM_CHALLENGE parameters (prompts) and verification routine call. This configuration can provide a vehicle for implementation of OTP on hardware architectures and operating systems that are unsupported by token vendors. Kerberos is the entry point for clients using OTP and only the KDC must communicate with the verification routines. Kerberos support indirectly implies OTP support on any architecture.

3.3.7 *What's Missing*

The implementation for this testbed was used in the FNC.GOV Kerberos realm and a larger test realm within the HPCMP. A problem arose due to the large amount of time it took to verify a SecurID passcode. The AceServer (SecurID verification daemon) will delay all responses for 1–15 seconds to try to detect replays. If the same SecurID passcode is used within the delay period, the AceServer will send failure messages back to both. This delays the KDC, which blocks, and causes the client to timeout the primary

KDC and send a request to the secondary KDC. The secondary KDC makes the call to verify the OTP; it will return with a failure since the passcode has already been used by the primary KDC. The AceServer response to the primary KDC may also be a failure if the delay period has not been completed. If the primary KDC receives successful verification it must still do some post-processing before it sends the AS_REP to the client. Within this time, the secondary KDC may have sent and received its passcode verification failure and sent a KRB_ERR message to the client.

Even if the delay period can be eliminated from the passcode verification process, allowance should be made for potentially slower verification processes and slower machines. A facility exists in the SAM specification to redirect a SAM_RESPONSE message to a set of KDCs that can handle the request. Implementation of this part of the specification requires a restructuring of client library code to restrict the KDC search and does not exist in the current MIT code base. The SAM_REDIRECT message must be authenticated so that a third party cannot insert it into an authentication negotiation.

3.3.8 Next Steps

Efforts to implement the SAM_REDIRECT facility would be useful to this and many other SAM configurations. Reorganization of the client code to pass KDC lists from preauthentication routines to the AS_REQ generation routines is necessary. A more generic interface for selecting a SAM type would be useful as well as modular challenge generation/response verification process similar to the pluggable authentication module specification.

3.4 Public Key Infrastructure Testbed

3.4.1 Purpose

The objective of this effort was to identify the Public Key Infrastructure (PKI) components needed to support both the CIS collaboration effort itself and general agency interaction and collaboration via the Internet. The plan was to leverage the efforts already underway at NIST under a separate GITS-supported effort and to:

- take advantage of a planned root CA for use by CIS participants, and
- develop basic CA security requirements, which would be needed to assure Internet users that the basic PKI was robust and reliable.

3.4.2 Background

PKI components are needed to support public key based cryptographic services over an open network. The primary purpose of the PKI is – normally through the use of certificates – to provide verification of users, keys, and other information needed to

conduct secure (i.e., signed¹ or encrypted²) information exchange over the network. Typically, a User A will consult a CA to obtain a certificate for User B, with whom User A wishes to communicate. The certificate will contain User B's public key and other information, and the certificate itself will be signed by the CA or other entity whose signature User A can verify. This process may repeat until User A can verify the necessary signatures.

The public key infrastructure is not a single, monolithic system or process. Rather, it is the set of components, processes, policies, and other elements that, together, provide the various PKI services needed for cryptographic protection to work. However, there must be certain common features to enable interoperability among users and PKI components. The focus of the PKI testbed activities is on those interoperability features.

3.4.3 Activities

The following is a brief description of each of the principal activities involved in the PKI testbed effort.

3.4.3.1 Minimum Interoperability Specification for PKI Components (MISPC)

NIST has an ongoing PKI development project designed to advance the development and deployment of a PKI in the federal Government. One of the initial activities of this project was the development of an MISPC. This specification is designed to ensure basic interoperability of PKI component products (e.g. certificate servers, directories, etc.). It is based on the X509 version 3 (ITU-T X509v3) certificate specification and specifies minimum functionality for digital signatures. This certificate format is extensible and can also address encryption needs, but the MISPC does not include encryption functionality in its current version.

The CIS project teams were provided copies of the MISPC document and provided comments and recommendations to NIST as part of the document review process. Although one of the original goals of the CIS PKI effort was to incorporate the MISPC work into the CIS testbeds, there were no off-the-shelf products that incorporated the MISPC sufficiently to enable any useful testing. This should be possible during Phase II of the CIS project.

3.4.3.2 Root Certificate Authority

It was initially planned that a NIST-provided root certification authority (Root CA) would be set up to serve all the CIS testbeds having PKI needs. However, due to

¹ In this discussion, the term signing will refer to the process of generating and attaching a digital signature for a message, record, or other set of data.

² Although encryption can refer to any process in which information is transformed by a cryptographic process, we will use the term encryption to refer to such translations performed for the purpose of providing privacy of information.

procurement delays, NIST was unable to initiate a contract for the development of the planned Root CA reference implementation that the CIS testbeds could use. Therefore, each CIS testbed that employed off-the-shelf products, used the PKI components provided by the respective products; it was not possible to conduct useful interoperability tests based on a common root CA. As of October 1997, the root CA development was underway, and it is hoped that it will be available for use in Phase II of the CIS project.

3.4.3.3 Certificate Policy and Certification Practice Statement Framework

One important, although unglamorous element of establishing a PKI is the development of policies and procedures under which various PKI elements will operate. In particular, since CAs are vital elements in the overall structure, the security policies and procedures of those CAs is critical. The NIST PKI team has developed

3.4.4 Conclusions

Due to our inability to implement a fully-functional CA or other PKI components, there are relatively few solid conclusions to be drawn from the effort. The primary products of this phase of the report include the following:

- **MISPC** – The specification itself has been circulated throughout the community, and has been included in the work of the IETF PKIX effort. This should help increase the overall level of interoperability of cryptographic services and PKI components.
- **Draft CA Policy and Practice Statement Framework** – This framework provides much of the non-technical basis for real-world implementation of a PKI.

3.4.5 What's Missing?

The key missing element – and, indeed, the objective of the MISPC – is the absence of basic interoperability among various crypto services products and their PKI elements. It is currently required that, for the most part, users all employ the same product or product family in order to interoperate. Although the MISPC was developed with the cooperation of several industry partners³ who have committed to implementing the specification in their products, such products are not yet widely available.

3.4.6 Next Steps

The following are some next steps that can be taken under the CIS effort to extend the work of the PKI testbed and to foster wider use of PKI based crypto services.

- **Explore additional applications of the MISPC** – The existing MISPC has yet to be implemented widely in off-the-shelf PKI products. It would be very helpful to explore

³ AT&T, BBN, Certicom, Cylink, DynCorp, IRE, Nortel Secure Networks, Motorola, Spyrus, and Verisign

the extent to which the specification can or will be implemented in existing or upcoming products.

- **Implement the use of a Root CA** – Once the pilot root CA is ready, it should be used among CIS and other participants to demonstrate interoperability and efficiency.
- **Security Specifications for CAs** – Since CAs are more than the hardware and software involved, there is beginning to be activity to develop formal requirements and security specifications for CA operations. This could also be explored as a CIS project.

3.5 INCIDENT RESPONSE TESTBED

Most agencies require incident response assistance now because of their rapid and expanding involvement in the use of the Internet and other networking technologies. OMB has recognized this long-term need by requiring agency incident response capabilities in OMB Circular A-130 (Appendix III).

3.5.1 Background

On June 3, 1996, the GITS Innovation Fund Committee granted \$2,796,000 to the NIST to establish a federal Computer Incident Response Capability (FedCIRC). The capability which is now operational assists federal civilian agencies in their incident handling efforts by providing proactive and reactive computer security related services.

NIST subcontracts the operational incident handling capability to the Defense Advanced Research Project Agency's CERT(SM) Coordination Center (CERT/CC) and to the Department of Energy's Computer Incident Advisory Capability (CIAC). NIST is responsible for operational management and for facilitating the development of incident handling standards and guidelines by utilizing the vulnerability data collected by FedCIRC. The vulnerability information will also be used in the analysis and testing of software and other products.

FedCIRC combines the experience and expertise of NIST's Computer Security Division, CERT/CC, and CIAC to provide agencies with cost reimbursable, direct technical assistance and incident handling support.

The need for an incident handling capability that crosses agency boundaries has never been greater. Almost all federal agencies are now connected to the Internet and exchange information regularly. The number of Internet related incidents that have occurred in the past year along with the increase and complexity of threats, requires agencies to take seriously their incident handling capability. The Office of Management and Budget (OMB) has emphasized this need in OMB Circular A-130, Appendix III, by requiring agencies to be able to respond in a manner that both protects their own information and helps to protect the information of others who might be affected by the incident. The private sector is undergoing the same rapid growth in network dependency as the federal community and are in need of the same incident handling support. Several private sector organizations have foreseen this need and have begun to offer incident handling services.

3.5.2 Purpose

FedCIRC was created to leverage incident response capabilities developed by scientific and research agencies (DoE, DARPA, NASA) to the broader federal networking community

- Provide on-call support for incident handling.
- Educate the federal community in computer security incident handling.
- Assist in providing an ongoing, sustainable process for incident response capability.

3.5.3 Relation to Other Testbeds

The core activity of FedCIRC, providing an incident response capability to federal agencies, has little relation to the work done by other testbeds. Yet, the outreach and education mission of FedCIRC has relevancy to nearly all other testbeds as the technologies employed to reduce the likelihood of an incident are being explored in each of the other CIS testbeds.

3.5.4 Activities of the Incident Response Testbed

There are six primary services that FedCIRC provides to the federal civilian agencies. The amount of service depends on the subscription level to which an agency or major operating unit subscribes. A description of each of the six services, as well as an explanation of the derivation for the cost structure, is given in the sections below.

3.5.4.1 Incident Handling Substructure

The incident handling substructure is the foundation of FedCIRC. The availability of the incident response hotline support and the collection, analysis, and publication of threat, vulnerability, and other security related data can only be accomplished if the underlying infrastructure is in place. The infrastructure consists of the following activities:

- alert creation;
- interaction with other incident handling organizations, law enforcement, and vendors;
- threat and trend analysis;
- hotline availability;
- data tracking;
- vulnerability analysis;
- report generation;

- database maintenance;
- guidance documents (e.g., example practices);
- web site maintenance; and,
- technology watch.

The substructure is, in essence, a shared resource among all subscribers and its cost will be shared proportionately. Each subscribing agency or organizational unit account will be debited monthly for one-twelfth of the total cost listed for this service under the selected subscription level.

The direct products derived from the substructure are described in the following sections.

3.5.4.2 Quarterly Activity Summaries

FedCIRC subscribers will receive by e-mail sanitized data and statistics about types of incidents and trends and information on new tools and guidance on preventing and handling incidents. The Quarterly Summary will contain aggregate data which is available to the public. The complete list of reported statistics can be found at the FedCIRC web site.

3.5.4.3 Web Site Access to a Repository of Tools and Example Practices

The FedCIRC created a public FedCIRC web site to provide access to a repository of tools and to give example practices. Topics available at the web server include: alerts, advisories, security tools, best practices documents, and links to other security servers.

3.5.4.4 Security Alerts, Advisories, and Bulletins

A set of FedCIRC security alerts, modeled after the CIAC Bulletins and the CERT Advisories, will be made available to FedCIRC members as events require.

These alerts will include a description of the vulnerability problem, the platform(s) and operating system(s) affected, the impact of the vulnerability problem, and patches and work-arounds, if available. The alerts will reflect the combined expertise and perspectives of the nation's most experienced incident response teams.

3.5.4.5 Incident Response and Hotline Support

Emergency technical assistance in response to computer security incidents is provided 24 hours per day, seven days a week. A "help desk" provides assistance during "normal business hours" (8:30 a.m. to 9:00 p.m. (EST/EDT, GMT-5/GMT-4)). Assistance is provided via telephone, e-mail, and pager hotline.

Incident response support ranges from providing agencies with direct technical support to handle computer security incidents or providing backup support to agency response teams dealing with large and complex incidents, to only providing agency response teams with information on threats, vulnerabilities and countermeasures that allow agency teams to effectively deal with incidents on their own.

Many activities are required to provide an incident response hotline. Some sample activities that provide this support include:

- problem analysis: analyze the problem, determine the magnitude of the threat, and provide technical assistance in identifying and closing vulnerabilities; •technical advice: issue advisories to the agencies warning of the problem and describing countermeasures;
- technical advice: provide guidelines on implementing vulnerability "fixes" and other security controls;
- assistance: facilitate the interaction of victims and relevant law enforcement agencies in reporting security incidents involving violations of the law;
- assistance: coordinate with other security organizations including the Forum of Incident Response and Security Teams (FIRST);
- assistance: work with vendors to provide critical security patches and work-arounds; and,
- vulnerability analysis: perform vulnerability analysis to identify a vulnerability's root-cause in order to identify other potential problems before they occur.

The types of operating systems for which FedCIRC will handle incidents include: UNIX, VMS, MVS, DOS, Windows, Mac, NT, VM, and MPE-XE. The types of protocols covered include: TCP/IP, IPX, Ethernet, LAT, DECnet, Token Ring, and FDDI.

Around-the-clock availability of the hotline is partially supported through the shared cost of the substructure services. In addition, a specific number of hours have been allocated for each of the three subscription levels.

3.5.4.6 Annual FedCIRC Incident Handling Conference

An annual conference on the current state of security threats and security improvement practices will be held in the Washington, D.C. area. The conference will be an opportunity for federal agencies to share lessons learned from security incidents and the results of security improvement efforts, as well as an opportunity to review commercially available security products.

The annual conference is a two day event that focuses on the lessons learned from each year's service to the FedCIRC members. There will be training sessions addressing those

countermeasures identified as reducing the current risks to federal information systems. Some of the proposed agenda topics or tracks for the annual conference are:

- status of the FedCIRC effort;
- updates from various FedCIRC members, e.g., successes, lessons learned;
- sessions on forming and sustaining an organic incident response capability;
- session on incident escalation and when to contact the FedCIRC hotline; and
- trade show exhibit space for vendors.

3.5.4.1 Semi-Annual "State of the Threat" Subscriber Meeting

In a given year, two subscriber meetings will be conducted. Meeting agendas will consist of two and a half days of briefings on current incident trends, recent vulnerabilities, latest viruses, and concentrated training. Detailed descriptions, impacts, fixes and work-arounds will be disseminated at subscriber meetings. FedCIRC subscribers will share related experiences, current practices, policies and procedures during these meetings. Seminar topics covered at the meeting

- "Internet Security for System and Network Administrators,"
- "Connecting to the Internet Securely,"
- "LAN Security for Desktop Systems - Novell, Windows 95/NT," and,
- "Virus Detection, Eradication and Prevention."

The training that is provided during the semi-annual subscriber meetings will evolve as required to address current computer and information risks. Current information will be available on the FedCIRC web page.

3.5.4.2 On-Site Information Security Evaluation (ISE)

One security evaluation of a single agency program will be conducted. The program to be evaluated will be identified prior to signing the Interagency Agreement. The evaluation will be performed over a four month time frame and will consist of two on-site visits. The scope of the ISE will be tailored to the subscriber's needs. The ISE assessment includes a review of selected components of the agency's network policy, infrastructure and network topology. The assessment identifies high leverage improvement opportunities and the most serious network vulnerabilities.

Recommendations on improving the security of the program will be presented in a report. A one day training session designed for the agency and addressing the report's recommendations will be presented on-site for all network and security administrators.

3.5.4.3 Assistance in Establishing an On-site Incident Response Capability

Five days of mentoring and hands-on training at either CERT/CC and/or CIAC will be conducted for a maximum of five students. Subscriber agencies are responsible for students' travel and accommodations. Training will be tailored to the needs of the subscriber. The training consists of working with incident handling specialists to answer hotline calls, handle incidents, and prepare alerts. All the issues, such as dealing with the press, law enforcement, security and network experts, and vendors, will be explored through first-hand experience.

This training covers topics such as:

- volunteer Incident Response Capability (IRC) option;
- incident escalation to FedCIRC;
- a budget and sustaining it over fiscal years;
- space, equipment and personnel;
- operational issues - incident handling procedures, physical, and electronic security;
- press issues;
- confidential versus public information;
- developing technical documents and standard replies;
- developing a FAQ (Frequently Asked Question);
- training constituency on reporting incidents, encryption, and other computer-related security topics;
- types of incidents and types of responses;
- importance of reference numbers;
- staffing requirements; and,
- information requests.

This training will be performed only for the platinum subscription level of \$250,000. It has been calculated that \$10,000 of the total subscription cost will be spent on this training. The subscription agency account will be debited for this service.

3.5.7 *What's Missing?*

In spite of increased publicity surrounding security breaches on federal networks, IT managers, as a whole, still do not consider security a priority. In this climate of shrinking budgets, an incident response capability is not seen as a necessity until after an embarrassing or costly intrusion has occurred. FedCIRC has provided incident response capability for two years to agencies that did not have internal expertise to manage security incidents. Though the need for incident response has not diminished, the FedCIRC service will be discontinued due to lack of funding.

Therefore, the two key missing elements in the federal government in this area are:

1. Broad awareness of both the value and the costs associated with and incident response capability.
2. Mechanisms for continuous funding of an incident response capability.

3.5.8 *Next Steps*

For the FedCIRC capability to be available and useful the agencies (those involved in the CIS effort and others), additional organization (to support a larger scale operations) and, of course, a clear and stable funding source will be needed. This may or may not be part of subsequent CIS activities.

3.6 Security Testing and Certification Testbed

As part of the CIS project, NIST undertook to involve the recently established National Information Assurance Partnership (NIAP) as a testbed for testing and certification of security technology and products. The NIAP is based on the National Voluntary Laboratory Accreditation Program (NVLAP), which provides an accreditation mechanism for commercial testing laboratories. The NIAP itself is a NIST / NSA joint initiative designed to meet the security testing needs of both information technology producers and users. The program is intended to foster the availability of objective measures and test methods for evaluating the quality of IT security products. In addition, it is designed to foster the development of commercial testing laboratories that can provide the types of testing and evaluation services which will meet the demands of both producers and users.

The program should help producers increase the value and competitiveness of their products (in the U.S. and abroad) through the availability of formal, independent testing and certification. NIAP efforts will help users - in both the public and private sectors - by providing a sound and reliable basis for the evaluation, comparison, and selection of security products.

3.6. 1 Background

In today's world of electronic commerce, organizations in both private and public sectors must place increased trust in the products, systems, and technologies they use to create, process, transmit and store valuable information. That trust is a measure of the confidence or assurance that a component or product - or an entire system - will perform reliably and to specifications, even in the face of intentional or directed "attacks." In other words, we demand security as part of that trust.

Formal testing has long been a principle method of assuring conformance to functional, performance, reliability, or interoperability specifications. However, if a component or product is intended to provide security services, i.e., to protect the functionality, performance or reliability of a system, then confidence demands become greater and the testing requirements become more complex and difficult. For example, in addition to testing that certain things happen according to specifications, security testing must also help assure unwanted things don't happen. Trust in a product can be enhanced further when the product has been tested and certified by a competent, independent third party.

Competent, independent security testing is needed by anyone who wants to design and build, market, procure, or employ products or systems requiring any level of security or trust. Designers and builders need effective tests and test methods for their products before sending them out the door. Vendors seek independent testing and certification as a way to demonstrate compliance with user requirements or to increase the value and marketability of their products to would-be users. Finally, users depend on independent testing and certification as a "mark of quality," as a way to distinguish between competing products, and as a way to assure conformance to their security requirements - that is, to enhance their trust.

Testing clearly adds value - but testing is not easy, and it's not cheap. In today's highly competitive, fast changing world of IT, neither producers nor users will tolerate extensive costs or delays to achieve that value. Today, there are few organizations able to perform competent, independent security testing - and even fewer effective methods for conducting security testing. In addition, because of the global nature of the information technology marketplace, vendors and users are demanding global recognition of product certification so they do not have to undertake new testing in each national or regional market.

As information technology continues to change rapidly, as products and systems become increasingly complex, and as our dependence on them becomes vital, effective and economic testing becomes critical. Security products must change to stay ahead of evolving threats; so must the tests, test methods, and metrics used to evaluate those products.

3.6.2 Purpose

With respect to the CIS project, it was assumed that as the security technology used in the other testbeds becomes more commercially available and widely used, there would

should be a demand for third-party testing or certification of products containing that technology. Thus, a number of those technologies were examined for possible development of formal testing criteria. In some cases, such criteria were developed for other technologies (e.g., Internet firewalls) as well.

The long-term goal of NIAP is to help ensure the security of information technology systems and networks through cost-effective testing, evaluation, and certification programs. For the U.S. this should help maintain the nation's leadership in information technology. Perhaps more importantly, NIAP should improve the trust of citizens, private sector organizations, and the government in the security and reliability of the constantly increasing elements of our economy and daily lives that are dependent on information technology. To achieve these goals, the partnership seeks to:

- Promote the development and use of security-enhanced IT products and systems.
- Demonstrate and increase the value of independent testing and certification as a measure of security and trust in information technology.
- Foster research and development to advance the state-of-the-art in security test methods and metrics.
- Move current government-conducted evaluation and testing efforts to accredited, private sector laboratories.
- Help establish the elements of a robust commercial security testing industry.
- Establish the basis for international mutual recognition and acceptance of tested and certified security products.

3.6.3 Relationship to Other Testbeds

Although there was no direct interaction between this activity and the other CIS testbeds, as was mentioned earlier, the technologies used in those testbeds will, at some point, need formal testing and/or certification when incorporated into commercial products. In at least one of the other testbeds (PKI), the principal products of the effort included development of specifications that would form the basis for subsequent testing (e.g., the MISPC and the security requirements for certificate authorities).

3.6.4 Activities of the Testing and Certification Testbed

NIAP is developing tools, test methods and tests for specification-based information technology security products. This means that the security functionality and assurance requirements of a product or system must be formally described or specified. These specifications then form the basis for the development and conduct of test for the product or for a class of product (e.g., for a firewall, an access control device, or a network router).

The internationally developed Common Criteria (CC) is the focus of much of NIAP's work. The CC provides a comprehensive, rigorous method for specifying security functionality and assurance requirements for products (or classes or products), usually in the form of protection profiles (PPs). The CC provides an internationally recognized basis for specifying and testing a wide range of security technology, from components to products and systems.

Although not part of this phase of the CIS effort, Internet firewalls are a recognized, effective technology for the isolation and protection of local networks or subnets from the Internet as a whole. The NIAP team produced a formal protection profile for network firewalls that may be used to test products used in subsequent CIS work.

It is important to note, however, that the tests and methods that NIAP develops and promotes are not limited to those based on the CC. Rather, any type of specification-based product testing program can be accommodated in the NIAP program. Therefore, it will be possible, for example, to develop security specifications for any of the technologies studied in the CIS project and develop tests, which may then be conducted by accredited laboratories.

3.6.5 What's Missing?

The concept of commercial accredited testing laboratories for security product testing is not a new one, but only in the last few years have efforts been undertaken to develop the basic concept into working documents, specifications, and processes. However, except for testing of cryptographic module implementations (FIPS 140-1), there are a few such accredited labs. There are *no* accredited CC labs (in the US) because of the relative newness of the program. If the security technology used in various CIS testbeds is to become commercially available and, most importantly, trusted by the user community, such testing processes may be demanded. It would also be helpful if PPs for more of the CIS testbed technologies were developed. (e.g., secure messaging, Kerberos, even FORTEZZA).

3.6.6 Next Steps

The NIAP is only just beginning its operations, and much is yet to be done to establish its programs and activities. However, the following are currently planned or underway:

- Development of tools for use by developers and test labs.
- Assistance to organizations wishing to establish testing and certification programs for standards, products, or classes of products.
- Accreditation of testing laboratories.
- Collaborative research or testing programs with industry or laboratory organizations.
- Research into advanced techniques of formal program specification and testing.

- Establishment of a private sector CC Testing Program (CCTP) laboratory accreditation program.
- Development of CC based protection profiles and associated test sets for selected classes of security products (e.g., firewalls).
- Establishment and promotion of a formal, international mutual recognition scheme based on a CC based evaluation.
- Maintenance of accredited laboratory, certification program, and evaluated product lists.
- Serving as a general center of expertise and resources for the security testing community.

These may or may not be undertaken as part of subsequent CIS efforts.

3.7 Privacy, Digital Signature & Secure Messaging Testbeds

Three separate testbeds (Privacy, Digital Signature and Secure Messaging) were established under NASA's responsibility. The intent was to smoothly transition them into a single entity using a phased approach. The common threads for the testbeds are:

- Confidentiality - controlling access to information so that only those for whom the information was intended will be able to retrieve it
- Integrity- protecting information so that the accuracy of the data can be accepted with a high degree of confidence
- Authenticity - providing the recipient with a level of trust that the information was actually sent by the person claiming to have sent it.
- Non-repudiation - assuring the origination of the information and that the originator can not deny sending the information.
- Public/Private Keys - management of the issuance, use, certification, revocation, and control of keys and key certificates.

3.7.1 Background

The federal Government's use of computers and networking has enhanced the speed and efficiency of communications among agencies and to the general public. However, this increased reliance on the open Internet and the World Wide Web has increased the chances for the data to be compromised. Potential compromises exist in the areas of accuracy, confidentiality, and authenticity. To safeguard these vital channels for disseminating information, improvements in authenticity, confidentiality, and security of communications must be developed.

3.7.2 Purpose

The purposes of the three testbeds under NASA's responsibility were:

- The Privacy Testbed focused on the issues of authentication, security and privacy using a system of public keys and certificates based on the MIT freeware version 2.6.2 of Pretty Good Privacy (PGP).
- The Digital Signature Testbed investigated the issues of authentication, security and privacy using a Certificate Authority System based on the Northern Telecom Entrust system.
- The Secure Messaging Testbed looked at the issues of the use, management and maintenance of a secure messaging system based on public keys and digital signatures.

3.7.3 Relationship to Other Testbeds

Because digital signatures can be used to prove the authenticity of information resources of web-accessible information, as well as authentication of users and the provision of access control to resources on the web, the work of this testbed is of interest to the Secure Web Testbed.

The PKI Testbed is helping to develop a formal method for managing digital certificates. This infrastructure element provides the necessary functions to manage the certificates necessary for validating the encryption keys used in both the Digital Signature and Secure Messaging Testbeds. Certificates can be simple, PGP type public keys and certifications like those being tested in the Privacy Testbed, or they can be much more strictly handled like the Certificate Authority System ones from the Digital Signature Testbed. A useful PKI must have a flexible means for controlling the public/private keys, as well as an ordered way for managing the certification, issuing, rescinding and changing of these keys and the accompanying certificates of authentication. These are precisely the reasons why testbed participants are looking into Simple Distributed Security Infrastructure (SDSI) with Professor Ron Rivest of MIT.

The Secure Messaging Testbed is exploring areas of interest to both the Secure Web and AA Testbeds because here too the use of encryption and authentication requires a flexible means for controlling the public/private keys and the accompanying certificates of authentication. Thus, the Secure Messaging Testbed is the application that ties the PKI Testbed, the Digital Signature Testbed, the Privacy Testbed, and the AA Testbed together.

3.7.4 Activities of the Privacy, Digital Signature & Secure Messaging Testbed

3.7.4.1 Privacy Testbed

The freeware 2.6.2 version of PGP was determined to be the best application for use within the Privacy Testbed. Copies of a CD-ROM with the freeware version 2.6.2 of PGP for Macintoshes, PCs and UNIX platforms were obtained and given to all CIS participants who wanted them. Assistance was provided to the National Science Foundation (NSF) and the National Institutes of Health (NIH) in installing PGP, creating public/private key pairs, managing a PGP public key ring, and in certifying user keys. Procedures for certifying PGP keys were established and disseminated to CIS participants. Levels of authenticity were defined for use within the Privacy Testbed. A "Kerberized," web-based, PGP key server has been identified and steps will be taken to incorporate this into the testbed.

3.7.4.2 Digital Signature Testbed

Northern Telecom's Entrust products were used as part of the Digital Signature Testbed. At the beginning, participants created a CA using Entrust. As a way to demonstrate the use of digital signatures, pilot projects were formed to use the digital signature in conjunction with Informed Electronic Forms. Currently, a pilot is set up to do a NASA Research Announcement (NRA) for the Mission to Planet Earth (MTPE) where the grant announcement, the proposal submittal, and the peer-review processes are done through the web with encryption and digital signature using commercial-off-the-shelf (COTS) software. These pilots will not only provide valuable experience on the Digital Signature Testbed, they will also allow for insight on how to set up an operational CA.

A collaboration agreement has also been set up with Professor Ron Rivest of MIT investigating the use of SDSI. We will investigate a simpler and less rigid public-key infrastructure. SDSI does not depend on a hierarchical global name space, which makes the implementation of PKI so difficult.

3.7.4.3 Secure Messaging Testbed

The Secure Messaging Testbed looked at the issues of the use, management, and maintenance of a secure messaging system based on public keys and digital signatures. As part of this testing, a PKI pilot was created using Entrust. The Secure Messaging Testbed was thought of as one of the many security services that could make use of this PKI. An attempt was made to create two Secure Multipurpose Internet Mail Extensions (S/MIME) plug-ins, one for Qualcomm Eudora and one for Microsoft Exchange, so that the secure messaging solution would be completely integrated along with a PKI. The reason that these two mail packages were selected is because they are the current official mail packages for NASA.

In the beginning, Ames Research Center (ARC) worked with Leif Nelson of Lawrence Livermore National Laboratories (LLNL). At the time, Nelson developed a Privacy Enhanced Mail (PEM) messaging solution for the 68K Macintosh for Eudora. ARC

started off with the code. Since then, ARC has completely rewritten most of the codes and made many changes. Some of the enhancements are:

- Developed a PEM based secure messaging solutions for Eudora on the PC (Windows 95 and NT).
- Ported the PEM-based secure messaging solutions for Eudora to both Macintosh 68K and PowerPC.
- Developed the first S/MIME-based secure messaging solution for Eudora that makes use of Entrust as the automated key management system on the PC (Windows 95 and NT) and the Macintosh platforms (68K and PowerPC).
- Improved error handling routines.
- Used built-in Directory Service for other secured partners.

Currently ARC is testing the S/MIME plug-in for Eudora with the NASA Postmaster Group and the Agency's Intranet Prototype Team. ARC is also actively working with Entrust and Worldtalk to come up with a S/MIME plug-in for MS Exchange.

3.7.5 What's Missing?

The technologies in the Privacy, Digital Signature and Secure Messaging Testbeds are not widely accepted or used. They tend to be restricted to small pilots and/or minimal rollout. There needs to be a concerted effort to increase understanding and acceptance of these technologies as integral parts of the agencies' information technology infrastructures.

Management has not acknowledged the fact that security technologies must be included from the initial design phases, rather than layered on or retrofitted later. There needs to be a concerted effort here to increase understanding and their acceptance of these technologies by management as an integral part of system design and development.

Interoperability between different CA systems is currently lacking. Users are struggling with issues of certificate import and export from browsers, interoperability among different vendor certificates, the lack of widely available software for issuing and managing client certificates, interoperability with certificates for secure -e-mail, etc. Industry acceptance of the NIST generated Minimum Interoperability Specifications for PKI Components (MISPC) would go a long way to improving this situation.

3.7.6 Next Steps

The integration of the three Testbeds (Privacy, Digital Signature and Secure Messaging) into a single Secure Messaging Testbed needs to be completed with support for S/MIME standard included. ARC will continue looking into alternate PKI solutions such as SDSI.

There is much work still to be done to exploit the relationships among the combined Secure Messaging Testbed and the other CIS Testbeds. As each of these activities progresses, the need for collaboration increases in importance. For example, collaborating on the automated issuance of client certificates based on the authentication mechanisms in the Kerberos, FORTEZZA, and Secure Web Testbeds should prove particularly fruitful.

3.7.7 References and Notes

To learn more about NASA Ames Research Center, see <http://www.arc.nasa.gov>.

To learn more about Entrust, see <http://www.entrust.com>.

To learn more about S/MIME and other interesting stuff related to S/MIME, see <http://www.rsa.com>.

To learn more about SDSI, see <http://theory.lcs.mit.edu/~rivest/sdsi10.html>.

To learn more about PGP, see <http://www.pgp.com>.

For an excellent compilation of web security issues, known bugs, and common questions, see <http://www.w3.org/Security/Faq/>.

For a good discussion of SSL, client and server certificates, etc., see the paper *Securing Communications* by Netscape at <http://home.netscape.com/newsref/ref/128bit.html>.

VeriSign has an excellent discussion of Digital Id's at <http://www.verisign.com>.

3.8 Secure Web Testbed

3.8.1 Background

The World Wide Web has exploded as a major vehicle for the dissemination of information among and within organizations. The Web is a standards-based, highly flexible system that is already a key element of the government's evolving information infrastructure. For the federal Government, the Web represents a way to communicate with the general public, as well as an opportunity to increase intra- and inter-agency information sharing and collaboration. As in the commercial world, federal business processes change quickly. This rapid change has tended to create environments where security is considered only after web-based systems are deployed. With the maturation of web development tools, server software, and web-enabled applications, web-based applications can be fielded with security in mind.

At the same time that web security tools have been developed, the evolving role and high visibility of the web has increased the requirements for security. In the context of the web, security concerns usually involve:

Deleted:

Inserted:

- **Authentication.** Because users are "hidden" behind their computers, dependable methods to authenticate a user's (or server's) identity is required.
- **Access Control.** Once a user's identity can be authenticated, access to information can be selectively allowed or denied, making the appropriate information available to the appropriate people.
- **Secure Transmission.** When sending information over the public Internet, the data passed between the web browser and web server must be secured to prevent snooping and modification by third parties.
- **Non-repudiation.** Once the originator of a message (via e-mail, web page, etc.) can be authenticated and the message transmitted securely, the message cannot be repudiated.

Due to recent, highly publicized, federal web site intrusions wherein documents were modified, securing information content on a server has become a high priority. While web-based authentication, access control, and non-repudiation are important aspects of securing information content, the techniques for securing the information assets were not explored as part of this testbed. While "server-side includes" and CGI programs can allow manipulation of data by intruders, most of the security concerns fall into the domain of host security.

3.8.2 Purpose

The Secure Web Testbed focused on the issues of authentication, access control, and trusted transactions. The testbed looked at client-based and server-based approaches as well as issues regarding interoperability between various software products.

3.8.3 Relationship to Other Testbeds

Because the web is an application that can use a variety of security mechanisms, the activities of nearly all the other testbeds are related to securing the web. Some web security mechanisms depend on digital IDs or certificates. Since a PKI is the wider context in which such certificates are proposed to be managed, the activities of the PKI Testbed are of great importance to web security. To the extent that Web pages can contain signed and encrypted objects, the Secure Messaging Testbed is exploring areas of interest. Authentication is an important step in Web security, and thus the work of the AA Testbed, and the Kerberos Testbed, are relevant. Because FORTEZZA can be used for both authentication and encryption, the FORTEZZA Testbed is also relevant. FORTEZZA mechanisms for both have been added to versions of the Netscape browser and server.

3.8.4 Activities of the Secure Web Testbed

Of all the CIS areas, those that the Secure Web Testbed investigated were the most rapidly changing. The tremendous proliferation of commercial web servers, web clients,

and web-enabled applications made the testing of specific solutions difficult as the target technologies quickly evolved.

3.8.4.1 Client Technologies

When this testbed began in July 1996, there were still a many different brands of web browsers in use worldwide. Early on, considerable work was put into testing recent alpha and beta releases of client and server software developed by the National Center for Supercomputer Applications (NCSA). They were available at no charge in source code form, and NCSA had put considerable work into experimental security mechanisms. They had a working implementation of S-HTTP, and experimental PGP and S/MIME mechanisms. In addition, licensing a Secure Socket Layer (SSL) toolkit enabled use of SSL.

During the latter part of 1996, it became clear that two dominant web clients had emerged: Netscape Navigator (now available separately or as a component of Netscape Communicator), and Microsoft IE. Netscape claimed 85 percent of the browser market a year ago, with others sharing the rest. Today, Netscape's market share has decreased significantly while Microsoft's IE has been increasing. The remaining browsers comprise less than 5 percent of the market. At the time of this report, both Netscape and Microsoft have released version 4.x of their respective browsers. Netscape Navigator is available on a wide variety of platforms for a modest price, although "demos" can be downloaded free. Microsoft IE is still available at no charge for Windows and Macintosh.

The emergence of Navigator and IE as the dominant browsers, coupled with NCSA's decision to cease browser development in 1996, led the Secure Web Testbed to quit experimentation with unique client security mechanisms. It was then decided that any security mechanisms that could be employed had to interoperate with Navigator and IE. Where the features of these browsers differ, an organization could choose to standardize on one or the other, but from a federal multi-agency perspective that was not practical. The software recommendation became (and remains): use Navigator or IE as your client/browser.

There was a downside of this market change however. Security mechanisms that were not supported by Navigator or IE could no longer be pursued. This means that S-HTTP is effectively dead (giving way to SSL). Further, interoperability between servers and these two clients is inconsistent at best. It is also harder now to work with Kerberos authentication on the web, because the clients don't have built-in support. PGP and S/MIME authentication, while still possible using custom plug-ins and CGI programs, is probably no longer worth pursuing. Because client certificates for these browsers are not interchangeable, individual users should standardize on one of the two browsers.

3.8.4.2 Server Software

The web server market is much more varied than the client market. As discussed under client software, this testbed originally focused on the NCSA server code, because it incorporated a number of security mechanisms. When NCSA ceased development, a

switch was made to the Apache server. Apache is available at no cost in source code form. With freely available patches and RSA crypto libraries (which require no license for government use), it is possible to build a freely available SSL capable web server. Apache also has a dominant share of the public web server market (while still true, it has been dropping over the past year). Recently, a Windows NT version of the Apache Server entered beta testing.

For building a freely available web server from source code for a UNIX platform, the Secure Web Testbed recommends Apache. Several things happened in June/July of 1997 that amounted to excellent news for the world of free, source code available, secure web servers:

- 05 June 97, Apache 1.2.0 was released;
- 25 June 97, SSLeay 0.8.0 was released;
- 29 June 97, ssl_1.7 for Apache 1.2.0 was released;
- 05 July 97, ssl_1.8 for Apache 1.2.0 was released (bug fix);
- 05 July 97, Apache 1.2.1 was released (bug fix);
- 18 July 98, SSLeay 0.8.1 was released (bug fix).

There has also been a great deal of activity in the commercial web server market. Over the past two years, increased competition encouraged vendors to drop their prices and quickly add new features and functionality. For example, most vendors were relatively quick to add support for SSL, x.509 digital certificate and the Lightweight Directory Access Protocol (LDAP). An added benefit is that commercial web servers are generally easier to install, configure and maintain. Over the long run, the decreased support costs make them well worth the purchase price. Today's premier commercial servers are Netscape Enterprise Server 3.x, Microsoft Internet Information Server 3.x, and Stronghold 2.x. (Stronghold is based on Apache) Any of these can be used to implement a tightly secured web server. The specific features, as well as pricing, change frequently. Details are available at each vendor's web site:

- <http://www.microsoft.com/iis> (Microsoft Internet Information Server 3.0);
- <http://merchant.netscape.com> (Netscape Enterprise 3.0);
- <http://www.c2.net> (Stronghold 2.0).

3.8.4.3 Client and Server Certificates

Certificates or Digital ID's for web clients and servers are central to the use of SSL. There are several commercial sources of these certificates, e.g. VeriSign and Thawte. VeriSign claims that over 16,000 web servers have Digital IDs issued by them. As of

August 14, 1997, only 302 organizations had purchased "Class 3" Digital IDs, their highest level of assurance. None of those organizations were federal agencies.

VeriSign also claims to have issued over 750,000 client certificates. This number is large because their "Class 1" certificates are available over the Internet at no cost. This shows that perhaps 3 percent of web users have gotten client certificates. Use of this technology is still very limited, though this may begin to change as Digital ID's get used for S/MIME secure e-mail.

Jeff Schiller of MIT demonstrated a system of issuing client certificates based on a different form of authentication. In his case it was Kerberos. If the user can present a Kerberos principal name and password to a certificate server, it will return a signed client certificate that can be used for web browsing for a certain period of time. This scheme could be used with any form of authentication, and may be attractive at sites where large deployments of technologies like Kerberos or FORTEZZA are in use. However, because of their temporary nature, these certificates are unsuitable for digital signatures and encrypted e-mail.

3.8.4.4 Privacy

Web server sites can usually learn a lot about visitors. For an example, visit <http://www.digicrime.com>. Methods of browsing anonymously have been devised, but are not yet widely available. The best known such service, the Anonymizer, can be found at <http://www.anonymizer.com>.

3.8.5 *What's Missing*

Encrypted SSL sessions are not yet widely used. They tend to be restricted to specialized information, such as the submission of billing or other sensitive information. There is a penalty for its use, in terms of additional computing cycles for the encryption on the client and server sides, which may be an issue for a server under high load. It is not an issue for clients, which usually spend more time waiting for data transmission than they do in processing the results. Far more use of encrypted sessions for all web access would be beneficial because it provides some privacy to the end user (the contents of their web transaction cannot be snooped), and authenticates the information provider (server).

The use of client certificates is still in its infancy. We believe that client certificate access control is the most promising technology for securing the web, while still making it easy to use. To date however, people are still struggling with issues of certificate import and export from browsers, interoperability between Netscape and Internet Explorer certificates, the lack of widely available software for issuing and managing client certificates, interoperability with certificates for secure e-mail, etc. Not all web servers do a good job of supporting client certificate access control mechanisms. More maturity of PKI would help in these areas.

3.8.6 Next Steps

Over the next several months, testbed participants expect to see the maturity of SSL-based client certificate authentication, both in the Apache-SSL and other servers. Support for client certificates is also improving in the latest browser releases. While all of the participants in this testbed believe that this is the most promising technology to pursue, none have yet experienced wide-scale deployment of this technology within the represented agencies, nor has it been used on an interagency basis. Such a deployment is viewed as a valuable next phase in this testbed. There are lessons to be learned about issuing and managing large numbers of client certificates (including across agencies), and about user education and reaction to this technology.

The other promising area to pursue is the automated issuance of client certificates based on other, possibly already existing, agency authentication mechanisms, e.g., Kerberos and FORTEZZA. Demonstrations of this type of capability by Jeff Schiller of MIT and others are exciting. An early version of this code was used in this testbed, but the creation of a portable system for doing this is not yet complete. Finally, in conjunction with the FORTEZZA testbed, the Secure Web Testbed will help test an SSL-enabled Netscape server that uses FORTEZZA for encryption and authentication.

3.8.6.1 References and Notes

The following URLs provide a good start for web-based security. This list is by no means exhaustive.

For a study on the growth of the web, see <http://ww.useit.com/alertbox.9509.html>

The WWW Security FAQ provides an excellent compilation of web security issues, know bugs, and common questions. This can be found at:

<http://www-genome.wi.mit.edu/WWW/faqs/www-security-faq.html>

For a good discussion of SSL, client and server certificates, etc., see the paper *Securing Communications* by Netscape (<http://home.netscape.com/newsref/ref/128bit.html>).

Information about the Apache web server can be found at:

<http://www.apache.org> (The Apache Project)

<http://www.algroup.co.uk/Apache-SSL> (SSL patch information)

<ftp://ftp.ox.ac.uk/pub/crypto/SSL> (SSL patch download site)

<http://www.psy.uq.oz.au/~ftp/Crypto> (SSLeay FAQ)

<http://www.psy.uq.oz.au/pub/Crypto/SSL> (SSLeay download site)

<http://www.apacheweek.com> (Apache Week Newsletter)

For an excellent discussion of the digital certificates and their relation to web technologies, see :

<http://www.verisign.com/repository/clientauth/clientauth.htm>

4 Conclusions

The experience of security professionals collaborating among agencies with differing security policies, computing environments, and security requirements yielded a varied set of conclusions. The following conclusions are among the more important discoveries of the CIS.

4.1 Security Solutions

4.1.1 Feasible Alternatives to Passwords

The CIS project identified a large number of alternatives to clear-text passwords including:

- biometrics (voice recognition, hand geometry, fingerprints, handwriting, retinal scans, face recognition);
- tokens (smart cards, SecureID cards, FORTEZZA cards,);
- software solutions (Kerberos, SSH, S-key, certificates of identity).

All of these solutions work and are available, so why aren't they universally used?

4.2 Obstacles

4.2.1 User reluctance

Most users are still not aware of the need for computer security. Many scientists have been heard to utter statements such as "if they get my data it will set them back five years." The other class of users is unwilling to tolerate anything that will make it less easy to access their computer resources. Both of these objections can be solved if the security mechanism is transparent (i.e., requires no action on the part of the user).

Most of the biometric solutions require user participation, and thus might be classified as "invasive." However, facial recognition software can be triggered automatically when the user sits at the computer console.

All of the token technologies require that the user carry the token around with him. However, some of the newer tokens (e.g., Dallas Semiconductor) are quite small and for example, can be incorporated into a finger ring.

All the software solutions require that the user "login" at least once in a session. However users are used to doing this, so it may not be construed as being too burdensome. In fact, solutions such as Kerberos can offer the user an improvement because the single logon will allow the user to securely access any Kerberized resource for the entire day. Other software solutions such as S-Key are quite burdensome.

Certificates of identity, especially Web-based client certificate authentication offers a secure, transparent mechanism of authentication and at the moment appear to offer the quickest, cheapest method of improving security.

4.2.2 Cost

Security is not cheap. Almost all of the token technologies cost about \$200/user including costs of administration and training. Technologies such as the FORTEZZA Card can be significantly more expensive because a reader must also be purchased for each workstation. There is an additional ongoing cost burden to maintain the integrity of the security databases, to administer the secure servers, and to keep the client machines well patched when government budgets are being cut. It is difficult to convince the organization's administration that these costs are high priority

It would be easier to make the case for spending the necessary money on computer security if government sites would calculate the true costs of computer security incidents, especially time lost by the user community. Unfortunately, these calculations are difficult, rarely performed, and the results are usually kept private.

4.2.3 Cross-Platform Support

The extremely heterogeneous computing environment of government agencies makes it essential that any security solution work on almost all of the computer systems being used. These systems range from hand-held PDAs to Crays. Obtaining a solution that is available on all of these platforms proved to be the biggest stumbling block for the CIS project.

- Most biometric solutions are available on Windows platforms only.
- Software solutions such as Kerberos and DCE are available for UNIX, partially available on Windows, and available with difficulty for Macintoshes. Kerberos is free on UNIX, but commercial versions cost about \$100 a copy on Windows and Macintoshes.
- Tokens that require readers usually only have support on Windows systems.

One solution to this problem is to standardize on the computing platforms that are deployed at each site. Making such a decision can create a storm of protest from users, especially at research laboratories where individual users may have a real need to run unpopular operating systems (e.g., OS/2). Other sites are stuck with old PCs that cannot run 32-bit Windows. Unfortunately, manufacturers create software and hardware first for the biggest markets. Users who try to buck these trends may find it difficult to participate in security solutions.

4.2.4 Lack of Ideal Solutions

The most fundamental problem is that today, there is no ideal security solution that can be deployed to protect all computing resources in a secure, but yet user-friendly way. CIS participants have combined Kerberos, DCE, and SSH to provide a pretty good solution in a UNIX environment, but this combination secures transmission and logins but does not address the problem of non-file-based authorization control.

Because there is no ideal solution for site security, and most existing solutions are expensive, a site runs a risk in deploying an existing solution because things are changing so rapidly. Hardware can become obsolete overnight, and software requires constant upgrades and maintenance.

Therefore, none of the CIS sites has succeeded in fully deploying any of the security systems mentioned in this report. Security-sensitive personnel and areas have been secured because the cost benefit of risk avoidance is clear.

4.3 Deployment Strategies

Despite the above-mentioned hurdles to the deployment of security solutions, the CIS believes that it is becoming more urgent for Government agencies to increase security by removing clear-text passwords and encrypting network traffic. It is hard (impossible) to construct "one size fits all" solutions; but it is possible to assist people and agencies to help themselves.

A recommended approach is to develop, advertise, and educate the network community on a set of security technologies (a "toolbox") that can be applied, or combined for suitable application to an agency. An agency must understand its network environment, execute a risk assessment, and develop requirements for a security infrastructure. An agency should probably define its environment as a collection of smaller environments (based on system architectures, user communities, etc.) that can be addressed separately. Flexibility, interoperability requirements, and growth should then be factored into the design and construction of a security system.

The best assistance comes in the form of education and understanding of network security technologies available to address the set of threats in an agency's specific environments. CIS has tested various approaches, and has created Web-based user assistance for the deployment of those technologies deemed to be ready for more widespread use, namely Kerberos and DCE.

Incremental deployment appears to be an attractive method for getting a security infrastructure launched. For example, requiring that all new computer systems have strong security, encourages users to implement it on older systems in order to access the new resources. If agencies decide that hardware security solutions (e.g., FORTEZZA cards) are appropriate, they should require that all new hardware be compatible with the chosen solution (e.g., have appropriate card readers).

Computer costs are dropping rapidly, but personnel costs are not. Consequently, to solve the legacy problem, CIS believes that in many cases it is cost effective for an agency to replace obsolete hardware and operating systems to be able to implement more of the recommended security solutions.

4.4 Issues

4.4.1 Encryption Export Controls

There is an ongoing national discussion on encryption centered on how to accommodate the private interest of individuals and businesses with the public interests of law enforcement and national security. On the one hand civil libertarian and business groups are concerned about government intrusion into areas they believe should be in the public domain. Areas of libertarian concern include computer security, cryptography policy, free speech and privacy. Areas of business concern include those, plus concerns about being placed at a competitive disadvantage with foreign businesses involved in IT.

On the other hand national security and law enforcement agencies are concerned about losing tools they maintain they need to investigate and sometimes prevent serious crimes and terrorist acts. They believe our national security may be jeopardized. Law enforcement strongly believes that the widespread use of robust non-key recovery encryption ultimately will devastate their ability to fight crime and prevent terrorism. They are concerned about drug lords and gangs being able to communicate about their crimes with impunity.

Balancing legitimate law enforcement and national security interests in preserving public safety and national defense against the public's desire to be free of unwarranted or excessive surveillance is a significant challenge facing any discussion of security technologies. Recent Clinton Administration efforts to strike a balance between the two camps by relaxing export controls, creating an international market for cryptography, and fostering key management infrastructure are all part of a sound investment strategy. Specifically, the policy requires key/data recovery but it transcends that issue and focuses on the more fundamental question of key management infrastructures (KMIs). The policy is designed to:

- Enable trustworthy encryption to be used widely to protect important U.S. information,
- Ensure that U.S. exporters can compete favorably in foreign markets, and
- Protect public safety, personal privacy and national security.

Additional steps should be taken to build public confidence in this policy initiative and the use of encryption products. Acceptance will enable the United States to continue its economic success while also maintaining public safety and security. KMIs will need to support key/data recovery for compelling business reasons such as accessing data when

the encryption key is lost or not available, and to support law enforcement's authorized access to evidence.

The lack of a PKI, a type of KMI, impedes cryptography use and, therefore, the full potential of electronic commerce. The CIS effort quickly came to appreciate the lack of commercial products in this area. Users need a secure and standardized mechanism to generate, store, authenticate, and transfer keys between users. Industry and government must work together to develop a PKI and the attendant products and services to ensure users can transmit, receive, and confidently use information electronically while also allowing timely and lawful government access. The development of a trustworthy key management infrastructure(s) supporting key recovery is a cornerstone of the Clinton Administration's encryption policy initiative. Certainly users should have the ability to choose their own responsible agents to generate and store their keys, but the government's public safety responsibilities will require that law enforcement, with proper authorization, be able to gain access to such keys. The Administration proposes to use privately-operated KMI data recovery features to support authorized law enforcement investigations, rather than creating a separate infrastructure that solely supports those investigations.

The Administration believes its policy enables industry and government to work together to develop and build the infrastructures for managing encryption keys. The fundamental issue is "how" industry will build key management infrastructures to support mass products with encryption. If infrastructures are built that support key recovery, then the export control debate can be concluded. With numerous Congressional committees studying the issue, certainly more discussion can be expected especially since export control policies complicate the implementation of solutions and collaborations. However, for federal agencies, a key-recovery mechanism should be a requirement to insure that the government maintains access to its own data.

4.4.2 Collaboration with Industry Leaders MS and Netscape

The two key computer industry leaders today are Microsoft and Netscape, and they are competing to control the World Wide Web market and your desktop. As the competition heats up, security issues become one of the top selling points. Security for Internet commerce and secure access to local resources are both very important, and can be handled in different ways.

Microsoft is approaching these two targets using different tools. For Internet commerce in general they are using public key based solutions, which are for the most part compatible with the Netscape solutions. For the access to local resources, NT version 5.0 will be using a Kerberos based solution, which should interoperate with other Kerberos based security products.

It is important that we interact with these industry leaders to facilitate the distribution and implementation of new security technologies. Our technologies must interoperate with their new products to gain strong market support. The name recognition that comes with these two companies will be an essential tool in introducing new technologies. Their

ability to harness the support of such a large number of users will be very effective in alpha-beta testing and collecting feedback.

4.4.3 Secure Web Technologies are Cost Effective and Provide a High Payoff

The World Wide Web is becoming the interface of choice for a wide variety of information resources. Agencies are using the web as a primary interface with the public, interagency collaboration over the web is increasing, and greater numbers of business applications are being web-enabled. This trend, along with the much-publicized cracking of federal web sites and advances in web security technologies and products, has led the CIS to conclude that web security should be a primary area of focus for government agencies.

CIS participants have successfully tested strong, biometrically-based user authentication and one-time password schemes using RADIUS with Netscape. Other authentication systems are also available with APIs that allow integration with browser services. CIS has also found that SSL encryption provides significantly greater protection for browser-based web management functions.

One significant problem that must still be solved is the matter of the infrastructure that is needed to support certificates for authentication and public key cryptography. While several commercial products have been developed that provide these services, they fall short of addressing the needs of the general public.

4.4.4 Leverage Other's Experience

It is not only the federal government and its agencies and contractors who must face security issues. The private sector and the educational community are facing the same problems. Often these other organizations face security challenges that closely match those faced by portions of the federal community. Those implementing security solutions can often save considerable time and expense by leveraging the expertise of others in the federal, commercial, and academic sectors. One example is the Big Ten universities who are struggling with authentication of thousands of students, both within their respective universities, and for some functions (such as library access), across universities. They have a joint project underway to use OSF/DCE that is based on Kerberos. The campus environment of the Big Ten, in some ways, resembles the environment of the DoE laboratories that face similar security considerations. By leveraging the experiences of these universities, DoE (and others) may be able to expedite the implementation of cross-realm authentication.

4.4.5 Barriers to Effective Interagency Collaboration and Cooperation

One interesting aspect of the CIS effort has been the mechanisms of interagency collaboration. Simply put, there usually are none. Each agency is a realm unto itself and other agencies are often equated to private industry. In addition, transferring funds and sharing resources between agencies has been, at best, difficult. Filling out forms and

getting the necessary approvals added significant cost and delays to all CIS interagency efforts.

For example, the Department of Energy classifies money coming from the GSA "work for others." The reason for this is that there are many rules and regulations in place that restrict the use of government facilities for non-official purposes. This stricture is usually extended to apply to other agencies as well. As a result, it took about three months, and required almost a month of person time, (spread over about six people) to fill out the necessary forms, obtain the necessary waiver of a DoE tax on work for others, and to get approval from the DoE Operations Office. GSA also placed restrictions on the use of CIS funds. We later learned that they could not be used to buy anything other than person time.

4.4.6 Levels of Trust Between Government Agencies

The work of the CIS dealt primarily with the interoperability and capability of security technologies in a testbed environment. In reality, there are many issues to be addressed in a real operational environment. Policies and procedures for ensuring trust between agencies need to be developed for an operational security infrastructure. Issues include:

- What is the suitable trust architecture to use — a single, consistent trust hierarchy or more distributed trust architecture?
- How should infrastructure CAs distribute information with the community of end-users?

As part of this FNC-CIS activity, NASA ARC has cross certified with Lawrence Livermore National Laboratories (LLNL). This means that the ARC CA and the LLNL CA recognize each other. In an operational environment there are certificate policy issues that should be addressed by a Certificate Policy Statement and its derivative Certification Practice Statements. These should cover all aspects of security policy from facilities and personnel requirements to technical security requirements and include a definition of the contents of the public key certificates issued by the infrastructure. The details of these statements will be dictated by the nature of the business applications, the organization's existing procedures, facilities, security environment, and organizational structures. These documents will provide the framework to process certificates issued by the PKI in a cross-certification agreement.

4.4.7 No "One Size Fits All" Solutions

Security requirements and solutions are highly dependent on environments, applications used, resources, and risk assessments. The CIS effort served to emphasize the diversity that exists across agencies, within single agencies, and even within individual organizations.

Not only are the hardware and networking environments varied and often incompatible, but also the needs of each group within an organization can be dramatically different. As

a result, CIS has reached the conclusion that there is no solution available that fits all users, and there are very few solutions that can be deployed across a large organization. Further, it is unlikely that a "one-size-fits-all" solution will ever exist. Possible solutions for this dilemma were discussed in section 4.2.4.

4.4.8 Keep Personnel and Security Information Resources up to Date

Many of the member agencies of CIS have already deployed pockets of expertise at their individual sites. Like many agencies, the expertise often resides in the skills and knowledge of particular individuals. It is important that personnel with responsibilities over information security and information security resources be current with evolving threats and computing environments.

For example, the ARL has a large effort underway to Kerberize their High Performance Computing Center. The DoE ESNet Applications Working Group has set up a series of Web pages to provide information on topics such as video conferencing. In addition, The CIS effort has generated its own help resource pages. Each of these organizations has made a commitment to keep their skills and information current.

The CIS Kerberos Task Force has created a page

<http://www.epm.ornl.gov/~jar/HowToKerb.html>.

To work, these pages must be regularly updated. The best pages seem to be under the control of a single individual who is willing and able to accept this responsibility.

4.5 Security Management Infrastructure (SMI)

SMI is the aggregate of supporting security services that enables a system to provide information security to electronic commerce or other enterprise computing. It consists of security components, as well as policies, procedures, and people associated with the operation of the system. Examples include creation and distribution of certificates that are associated with individual users in a trusted way, trusted back-up services, compromise recovery mechanisms, public key directories and handling of hardware tokens. SMI pilots should address the key delivery thread, proofing of customers, proposed standards, recovery from lost or compromised certificates, and security of directories. The long term goal is to have an SMI that is compatible and interoperable across Government and the private sector.

Public and business awareness of network security issues is and will continue increasing. The only technology that currently offers hope for large-scale, interoperable, robust information security is public key cryptography. However, to be most useful, this technology requires a service infrastructure that supports the enrollment and management of security participants.

Companies that provide PKI, a subset of a SMI, services will grow rapidly over the next decade. Competition will be intense, but the availability of usable standards will prevent the market from coalescing to a single provider.

In a sense, the services of PKI providers underpin the entire security model of our future information systems. The issuance of digital credentials is a supremely serious step in a world where large financial transactions, personal business, and government activities are beginning to be conducted online. The ultimate costs and overhead involved in providing this service are not yet clear, but only with a good, robust, secure and trusted SMI will our country be able to reap the full economic benefits of global networking.

Lastly, over the next decade, PKI service may enter more sectors of our society. If conditions are favorable, PKI may offer another path for disparate elements of our society to enter the economic mainstream. This desirable end can be achieved only if the costs associated with SMI can be held to a manageable level. A high assurance, interoperable, SMI is necessary to provide the means to enable the next step in the evolution of the National Information Infrastructure. It will provide the security management services needed for further development of secure information related applications.

5 Recommendations

In addition to the overall conclusions discussed in the previous section, a number of general recommendations were developed from the work in each of the testbeds. These recommendations are divided below into two major groupings: 1) recommendations for management involved in the use of the Internet, and 2) recommendations for others in the Internet community, i.e., outside the direct influence of most Internet users but still needing attention.

5.1 Recommendations for Internet Managers

These recommendations are directed primarily to management responsible for implementing security controls in an agency's Internet-based applications. They are meant to convey lessons learned through the experiences of testbed participants.

5.1.1 *Don't Forget the Basics*

While new products and technologies are constantly being developed to assist in the process of implementing security on computer networks, focus should not be shifted from sound, basic principles and techniques. Turning off unused network services, installing the latest vendor patches to the OS and network software, backing up important systems regularly, removing unused accounts, and fixing newly discovered vulnerabilities are representative examples of these security basics.

5.1.2 *Focus on Requirements before Technology*

The CIS testbed demonstrated (as should have been expected) that security technology (or any technology, for that matter) should not be implemented for its own sake but rather based on well-articulated need. Technologies should fit the security goals of the enterprise. The function and limits of a particular technology or product should not shape enterprise security.

5.1.3 *One Size Does Not Fit All*

While it would be ideal to have a single set of security technologies that can be applied throughout federal networks, differences between organizational requirements, computing and networking environments, and user sophistication make universal solutions infeasible. The work of the CIS testbed participants has shown that application of security technologies on a smaller scale greatly increases the chances that a particular technology will meet agency needs without hampering its mission. For those looking to implement security technologies, interoperability should be a top priority.

5.1.4 Look for Small Victories

The government has looked to large, single, "standard" solutions to many of its IT needs, including security. Large, monolithic solutions are difficult to field and often do not meet the needs of a changing infrastructure. Rather than wait for the full implementation of such a system, it is recommended that managers and administrators work to secure what resources they can today. Setting up one-time password authentication for mobile users, installing the latest patches on a busy server, or installing the latest version of sendmail all add to the overall security of an enterprise and can greatly reduce an organization's vulnerability.

5.1.5 Focus First on High Payoff Applications

While budgets for implementing security measures on government systems have increased over the last few years, funding is still very tight for managers and administrators. Therefore, it makes sense that IT managers should focus security dollars on high payoff applications. While these may vary between agencies, two applications immediately stand out as indispensable elements of today's networking infrastructure: e-mail and Web applications.

5.1.6 Make Use of Available Expertise and Resources

Rarely are an organization's security requirements unique. Usually, network and system administrators can find that other agencies have faced the very same requirements and many have implemented technologies to meet those needs. Those individuals in other organizations who have been working toward similar goals in their own networks can be an invaluable source of information. Finding common pitfalls of a particular technology, weeding out technologies that do not meet expectations, and gleaned lessons learned from experienced administrators will help reduce personnel costs, increase the likelihood of meeting schedules, and reduce the risk of stranded investment in technology. Expertise in other agencies as well as network accessible resources can greatly increase an agency's chances of successful deployment of security technologies.

5.2 Recommendations for the Federal Internet Community

The recommendations for the federal Internet community are aimed at high level managers and policy makers whose decisions can effect the ability of technology managers and administrators to implement security measures on agency networks and systems.

5.2.1 Encourage Working Level Collaborative Projects

The CIS project has demonstrated clearly that interagency collaboration at the working/technical level works and provides value to more than just the collaborating agencies. We believe, therefore, that other types of inter-agency collaborations (i.e., not just in the area of Internet security) should be encouraged.

Although the specific mission of each agency is different, the overall business processes, information technology and security needs, and the technology, products, and services available to meet those needs are essentially the same. This is true not only among government agencies, but also among government and private sector organizations as a whole. Therefore, more collaboration is sure to result in more efficient use of resources of all types - information, technology, financial, physical, and human. The CIS project has demonstrated this on a small scale, and the government has the opportunity to demonstrate it on a much larger scale.

5.2.2 Remove Barriers to Collaborations

The GITS Innovation Fund concept has proven to be an excellent way to support and encourage new ideas in the use of information technology in general and the Internet in particular. However, the CIS project, as a multi-agency collaboration faced some special hurdles that should not be there. In particular, because the GITS project required the distribution of funds among several participating agencies and what amounted to multi-lateral memoranda of understanding (MOUs), the process was exceedingly difficult. Personnel from GSA, the NIST comptroller's office, and our support contractor, DynCorp, were very helpful in getting this sorted out. Nevertheless, this took an inordinate amount of time to get the project off the ground, at least legally and financially. Moreover, several of the participating agencies found it easier to use their own funds to support their participation rather than go through the billing and reconciliation process of interagency funds transfers.

This process needs to be simplified — both when special funding arrangements are involved and even when collaborating agencies are going to provide their own support funding. If potential collaborators must deal with all these administrative matters, they will simply avoid the collaboration.

This should not be all that difficult, and it may be as simple as providing easy-to-obtain administrative (i.e., both financial and legal) support for proposed interagency collaboration projects. Standard Memorandum Of Understanding (bilateral and multilateral) and standard funding and funds-sharing mechanisms need to be available, along with information (understandable even by technical personnel) for their use.

5.2.3 Encourage Multiple Solutions and Interoperability

The government has in the past often relied upon large, single, "standard" solutions to many of its IT needs. However, when combined with the legal and procurement constraints often imposed on government (i.e., the "level playing field"), this has often resulted in massive, top-down "architecting" efforts that have had the unintended effect of drawing out the process and hobbling improvement efforts. This had been especially true in the security area.

In recent years, efforts to develop reusable software components and the focus on standards for the purpose of interoperability have resulted in increased flexibility on the part of users and developers and less dependence on single solutions—either technical or

supplier. This same concept needs to be employed to help solve some of our IT security problems. While the government has (as argued earlier) largely the same business and IT needs as the private sector, it is not monolithic, and a "one size fits all" approach to security needs is neither feasible nor desirable.

We believe that the CIS project has demonstrated the feasibility of using multiple solutions across agencies needing to work together or share information. This has *not* shown a need, as some have argued, that every agency must therefore support all possible standards and all possible types of security policy and technology.

5.2.4 Institutionalizing the CIS Concept

We believe that the basic CIS concept—open and active collaboration among agencies in the use of "grass roots" Internet security technology—provides sufficient process that there should be some examination of how the collaboration model could be "institutionalized" in and among U.S. Government agencies. Far too often, agencies (especially in the area of security) tend to "do it themselves," thus duplicating efforts and failing to learn from the experiences of others. Institutionalizing such technical collaboration efforts will probably entail senior-level commitment (e.g., through existing interagency bodies such as the CIO Council) and perhaps the support and cooperation of OMB.

6 Follow on Activities

The following activities are some of the possible technologies that the CIS is looking to investigate in its continued activities in interagency security.

6.1 Kerberos Testbed

The Kerberos testbed has found support from the private sector. Kerberos V5 will be used in the next version of Windows NT. Therefore, this testbed will be examining the interoperability of this Kerberos with other Kerberos systems. Additionally, they will continue to provide up to date information and advice on using Kerberos through their web site. The steps for the next phase are listed below.

1. Begin testing on Kerberos V5 when it comes with the next version of Windows NT
2. Keep the "How to Kerberize your Site" website updated with all of the latest information on Kerberos

6.2 FORTEZZA Testbed

The FORTEZZA testbed has made plans for the Second Phase of our CIS efforts. Walt Hinson and Mike Green of NSA lead this testbed in collaboration with ARL, DoE and NIH. In moving forward with their FORTEZZA testbed they have listed these below as steps for the next phase.

1. Finish our current effort to provide MISSI products to testbed members (PCMIA Card Readers and FORTEZZA enabled applications).
2. Provide new software applications as they become available as well as "smart card based FORTEZZA and software based FORTEZZA.
3. More work must be done in the "Security Management Infrastructure"(SMI) area. The lack of a SMI is a major impediment to applications such as electronic commerce and the FNC P&SWG CIS initiative can help focus this critical security need.
4. The NSA approach to Information Security is continuing to evolve. A new initiative, Commercial INFOSEC Solutions (CIS), is being formulated and our strategy and findings could be shared with the P&SWG CIS.

6.3 Advanced Authentication Testbed

Participants within this testbed have moved on to other positions outside of the interest of this working group. Therefore this specific testbed will not operate in the second phase. It has made a large contribution to the FORTEZZA and Kerberos testbeds that will continue to use what it has gained from the AA testbed in creating a better solution.

6.7 Privacy, Digital Signature and Secure Messaging Testbeds

In pursuit of improving and creating an infrastructure that can support secure messaging, privacy and digital signatures this testbed is moving forward into the study of a PKI. This testbed is now interested in the both theoretical limits and actual implementation issues involved with a full deployment. They are not moving away from secure messaging, privacy, and digital signatures, but are trying to build up the infrastructure to allow the deployment of these security services.

A description of the plans for the second phase is as follows:

First, the existing technology in PKI will be studied. These include the traditional X.509 standard, the PGP approach, and the merging of Simple Distributed Secure Infrastructure (SDSI) and Simple Public Key Infrastructure (SPKI). The advantages and issues of these three models will be investigated.

The SDSI/SPKI design is invented by Professor Ronald Rivest of MIT and others. Prof. Rivest is the founder of RSA Data Security, Inc. and is the world leader in cryptography and network security. We are currently conducting a cooperative research with him on this design. Second, the reliability, scalability and performance of PKI will be analyzed by using advanced mathematical modeling. Specifically, queuing theory and graph theory will be utilized to model the PKI.

The parallel channels and truncation (M/M/C/K) model and Open Jackson Networks in queuing theory will be used to analyze the scalability and performance of the PKI. These queuing models will be designed and implemented by using MatLab numerical software.

The maximum flow, minimum cut algorithm in graph theory will be used to investigate the reliability of the PKI. A refinement of this algorithm will be utilized for modeling. This algorithm provides a systematic method for finding an augmenting path in a network with a given flow.

From these models, the maximum throughput, the minimum queuing delay, the maximum reliability, and the Mean Time Between Failure (MTBF) of a node to access a specific CA will be found. The boundary conditions of each variable will also be studied.

Performance simulation of PKI:

We will study the scalability and performance of a PKI using Entrust 3.0 product. Simulation will be based on the current existing configuration of NASA. It will assume a centralized CA with distributed Registration Authorities for each NASA Center. Simulation will done under assumed normal usage, different time zones, normal revocation, and current network configuration. Test results will be recorded and analyzed for performance and scalability.

6.8 Certificate Management Testbed

The first phase of CIS identified the use of certificates as one of the highest cost-benefit ratio security infrastructures, and recommended that identity certificates be adopted to provide for secure, non authenticated Web connections. If one decides to adopt certificates (both identity and authorization certificates), there are several issues that must be faced. Because the use of certificates is just getting underway, tools to solve these problems are just emerging.

The basic problem is that a certificate represents a person's electronic identity. Most of us require more than one identity. At a minimum, we have an "at home" identity (used for conducting personal finance) and an "at work" identity (used for representing our employer in transactions). Unfortunately, there are also instances when current technology forces us to have more identities than we want. As a result, a user may end up with a dozen certificates issued for different purposes, and managing them poses a severe problem.

The same identity certificate does not work in both Netscape and Microsoft Internet Explorer (IE). Aside from the fact that Verisign, for example will not issue two certificates with the same e-mail address, there is no way to transfer a key pair from Netscape to IE. Thus, if I send encrypted e-mail from my Netscape mail client, and I receive the encrypted reply on a machine with IE, I will be unable to read it.

Even if a user restricts himself to using, say, Netscape, he will have to export his certificate and key pair from the machine that generated it to the browsers on his other machine. This is fairly straight forward with Netscape, but tedious. Netscape exports the key and certificate in a PKCS #12 file format which is a documented standard.; MS IE also has key import/export capabilities. The key file has a .PFX extension and is incompatible with the Netscape format.

Financial institutions will soon start issuing certificates (e.g., SET) for financial transactions. Initially, every institution will issue its own certificates.

Users will have to maintain their certificates in a secure manner on all the machines that they want to use to transact business. The certificate files are protected by a password that is probably much less secure than the key used in the certificate, and if these files are stolen, they can be cracked off line.

The Certificate Management Testbed will be concerned with the interoperability and certificate management problems outlined above, and will try to determine a rational policy for determining how many certificates a user should require.

Items such as electronic wallets will be tested.

6.9 Electronic Laboratory Notebook Testbed

The concept of a laboratory notebook is familiar to most people. Laboratory notebooks serve two important functions:

- Serve as a record of activity for a researcher.
- Provide legal proof of the time of an invention via numbered and witnessed pages.

Both of these functions are still vital for research. However, the tools a scientist uses, the form of the information that must be recorded, and the sheer volume of information are incompatible with the traditional paper notebook. For example, how does one record video or audio in a paper notebook? To cope with the new scientific venue, electronic versions of laboratory notebooks have been introduced. They have many advantages over their paper counterparts:

- Entries can be made electronically, thus eliminating transcription errors.
- They provide full multimedia support.
- Most scientists type faster than they write, so it is easier to write in an electronic notebook.
- They facilitate remote access. Because the notebook is hosted on a web server, it can be accessed remotely (securely if it is an https URL), or can be hosted on an always present laptop computer.
- They facilitate collaboration. Multiple researchers on a single project can share the same notebook and annotate other's entries.

There are philosophical arguments about the scope of electronic notebooks. One faction feels that an electronic notebook should encompass a scientist's whole world. It should be able to contain all data, launch applications, and be the scientist's window onto his research. The other faction feels that other tools are better suited to these purposes, and the electronic notebook should be a more general version of the traditional laboratory notebook.

DoE is sponsoring a significant effort on electronic laboratory notebooks (<http://www.epm.ornl.gov/enote/>) at three sites, ORNL, LBNL, and PNL. While the flavor of each version is different, they have agreed to a common architecture so that information can be transferred among the notebooks. The ORNL version of the notebook is being used by over 50 groups around the world (<http://www.epm.ornl.gov/~geist/java/applets/enote/users.html>). The electronic laboratory notebook runs on any platform with Perl 5 and a web server. You can try the notebook at the URL <http://www.epm.ornl.gov/~santa/enote.html>. The username is *guest*, and the password is *Demo*.

The biggest (unsolved) problem is a legally acceptable method for certifying a page as having been created on a certain date and having been unchanged since then. The ORNL group will introduce digital signature and external CA methods into their notebook in the near future. The problems to be solved here include:

- Pages that have already been signed should be able to be annotated without invalidating the original signature.
- How do you sign a page with dynamic content? The Materials Microcharacterization Collaboratory (MMC) DoE2000 project has a live video feed on a page of their e-notebook.
- How do you keep signature costs to a reasonable level? Organizations such as the Post Office will offer signature services, but if every scientist in a laboratory signs his pages every day, the costs for this mount up. Can the government create a signing body such as NIST?
- The legal issues involved in patents of work recorded in an electronic notebook are as yet unresolved.

For the CIS testbed, a secure electronic laboratory notebook will be set up for members of the CIS to use in exchanging information. Access will be controlled via certificates. Because the notebook accepts full html pages, the electronic laboratory notebook can replace the private CIS web pages. One advantage of this is that CIS members can change the content of (annotate) these pages as they read them rather than sending an e-mail to everyone.

6.10 Intrusion Detection (ID) Testbed

Organizational internets in the federal government are primarily relying on firewalls to protect their networked resources. While this provides some level of security, firewalls are insufficient in today's environment. Firewalls can only protect an organization from attacks that cross the firewall and do nothing to guard against unauthorized connections being established or protect resources from internal threats. It has been estimated by some law enforcement agencies that 80 percent of security breaches are carried out by internal threats.

The Intrusion Detection (ID) Testbed will

- Discover the current uses of both commercial and emerging research ID tools in the federal Government.
- Evaluate ID technologies and products in various environments.
- Develop guidelines for use and deployment of ID technology within the federal government.

6.11 DNSSEC Testbed

The Domain Name System (DNS) has become a critical operational part of the Internet infrastructure yet it has no strong security mechanisms to assure data integrity or authentication. Without such mechanisms, name resolution cannot be fully trusted. Secure implementations of DNS (DNSSEC) are being introduced.

The DNSSEC Testbed will focus on:

- Interaction between DNSSEC compliant servers and current non-compliant resolvers;
- Key management issues related to the .gov DNS hierarchy;
- Integrating DNSSEC-compliant resolution requests into current applications.

6.12 Push Technologies/Software Distributions

Push Technologies will be an effective tool in creating a more efficient method of securely distributing patches, applications, and updates within an organization. In order to keep security risks low, all workstations within a network should be using the latest version of their security applications. With the exponential growth of technology today, using Push technology in updating workstations will be one of the most effective ways in reducing security risks within security applications. Push technology has already been implemented in smaller networks, but this is still an open area for work. DoD may be looking into this area, with participation from another agency, they could run a testbed on Push technology.

6.13 Network Scanning

Network Scanning could prove to be a useful security tool. Using this tool we will be able to do an intelligent analysis of a local network. This analysis will check for holes in the network and alert staff of possible security risks within the system. This type of technology would be an excellent tool in creating a more efficient way to monitor networks and save valuable time.

Appendix A. Examples of Security Basics

In addition to all of the advanced security technologies discussed in this report, it is still necessary to follow sound basic security practices. Far too often these have been neglected, and have proven to be the source of a security incident.

- Security in many cases is a line management responsibility that far too often has been assigned to the system administrator who is overworked. Since good work by the system administrator is rarely ever seen, or rewarded, it tends to become a low priority item, and does not get the attention it deserves. Management need to allocate appropriate resources to maintain good security.
- Apply security-related software patches in a timely fashion. You are most vulnerable to a security-related bug after its vulnerability has been discovered and you have not installed the fix. Many hackers are not very sophisticated, but use well-known bugs to access systems that have not been patched.
- Don't run unnecessary services on a machine that does not use them. For example, turn off inbound sendmail if the machine does not need to receive mail. Turn off TFTP and FTP if not needed.
- Restrict access to services to only those who need access. For example this might entail running *wrapper* to limit from where connections can be made. Require additional authentication if access is attempted from off site,
- Don't mix general applications with special applications on the same machine. For example don't run your web server on a machine where you allow outside users to login.
- Use good passwords. This also applies to the pass-phrase used to protect your private keys, or your Kerberos password.
- Avoid using the same password for different systems, especial if they have different exposures. For example, if you need a password for a web application, don't use the same password that you use for your login.
- Remove old accounts. Accounts of terminated employees, students, consultants, collaborators etc. should be removed when their need for access is no longer required.
- Watch out for that old neglected system, the demo system, or the system that is part of some piece of equipment and on your network. If a hacker can break into it, they may be able to sniff passwords, or attack other systems while appearing to be an insider.

There are many other security practices that could also be applied. A good risk assessment is one component of what might need to be done. See FedCIRC at <http://csrc.nist.gov/fedCIRC> for additional information.

Appendix B. References/Links

B.1 *Participating Agencies*

The following agencies participated in the CIS effort by providing management, computing and human resources for the investigation of the CIS technology focus areas.

Defense Advanced Research Projects Agency (DARPA)

<http://www.darpa.mil>

Department of Defense (DoD) / Army Research Laboratory (ARL)

<http://www.arl.mil>

Department of Energy (DoE)

<http://www.doe.gov>

National Aeronautics and Space Administration (NASA)

<http://www.nasa.gov>

National Institutes of Health (NIH)

<http://www.nih.gov>

National Institute of Standards & Technology (NIST)

<http://www.nist.gov>

National Science Foundation (NSF)

<http://www.nsf.gov>

National Security Agency (NSA)

<http://www.nsa.gov>

Other sites for Internet privacy and security include:

B.2 Incident Response & Monitoring:

The following are references to other government organizations/projects that have been initiated in response to the need for a rapid response solution for federal network incidents.

- CERT Coordination Center
<http://www.cert.org>
- CIAC
Computer Incident Advisory Capability
<http://ciac.llnl.gov>
- FIRST
Forum of Incident Response and Security Teams
<http://www.first.org>
- NASIRC
NASA Automated Systems Incident Response Capability
<http://www-nasirc.nasa.gov/nasa/index.html>

B.3 Internet Security

To gain a quick, light-hearted view of the types of cyber crime, fraud, and attacks that exist, visit

<http://www.digicrime.com>.

This site is a particularly good inculcation for new Web users to make them aware of the dangers of web surfing. It is a service brought to you by some of the best experts in web security.

Internet security is rapidly becoming an active arena for several government agencies. Below you find references to these agency run projects and programs that are directly involved with Internet security as well as important academic and industry projects.

- federal Security Infrastructure Program
<http://gsa.gov/fsi>

The FSI Program will coordinate, operationally oversee, monitor, implement, and report on the development of an information security infrastructure to support electronic commerce, electronic messaging, other applications and support services to users.

- COAST Hotlist on Computer Security, Law and Privacy

<http://www.cs.purdue.edu/homes/spaf/hotlists/csec-plain.html>

Computer Operations, Audit, and Security Technology- a multiple project, multiple investigator effort in computer security research

- Computer Security Technology Center

<http://ciac.llnl.gov/cstc/>

Providing solutions to U.S. Government agencies facing security challenges in information technology

- IETF Security Working Groups

<http://www.ietf.cnri.reston.va.us/html.charters/wg-dir.html>

A list of active IETF Security Working Groups

- National Security Institute's Computer Security Resources

<http://www.nsi.org/compsec.html>

NSI provides professional information and security awareness services to defense contractor and industrial security executives throughout the United States

- NIH Computer Security Information

<http://www.dcert.nih.gov/security/dcertsecurity.html>

This page features security-related information of potential value to users of DCRT supported computing systems

- NIST Computer Security Resource Clearinghouse

<http://csrc.nist.gov/>

Provides an online archive of useful information

B.4 Internet Privacy

Internet Privacy has become a new issue in today's rapidly growing information age. Several organizations have been formed to represent the wide spectrum of views on Internet privacy; these can be located at the web sites indicated below:

- Center for Democracy & Technology (CDT)

<http://www.cdt.org>

An advocate of public policies that advance constitutional civil liberties and democratic values in new computer and communications technologies.

- Digital Future Coalition (DFC)

<http://www.dfc.org>.

A coalition committed to striking a balance in law and public policy between protecting intellectual property and affording public access to it

- Electronic Frontier Foundation (EFF)

<http://www.eff.org>

Electronic Frontier Foundation is dedicated to finding ways to resolve these and other conflicts while ensuring that essential civil liberties are protected.

- Electronic Privacy Information Center (EPIC)

<http://www.epic.org>

EPIC is a public interest research center focusing public attention on emerging civil liberties issues and to protect privacy, the First Amendment, and constitutional values.

- Internet Privacy Coalition

<http://www.privacy.org>

A coalition that encourages relaxation of encryption export controls and provides public availability of strong encryption and promotes privacy and security.

- OECD Privacy Guidelines

<http://www.oecd.org/dsti/sti/it/index.htm>

Guidelines on the protection of Privacy and transborder flows of personal data

- Privacy Forum Archive

<http://www.vortex.com/privacy.html>

- Society for Electronic Access

<http://www.sea.org>

An open membership society that is "making Cyberspace a better place to live, work, and visit".

B.5 Kerberos

This network authentication system is a product of the Massachusetts Institute of Technology. Below you will find pointers to acquire a free version of Kerberos and other pointers to Kerberos-related web sites.

- MIT Kerberos

<http://web.mit.edu/kerberos/www/index.html>

The MIT Kerberos initiative web site

- How to Kerberize your site

<http://www.epm.ornl.gov/~jar/HowToKerb.html>

A step-by-step guide to "Kerberizing" your site; designed by the Kerberos Testbed participants.

Getting Kerberos:

- Kerberos (UNIX)

<ftp://athena-dist.mit.edu/pub/kerberos>

- Kerberos (UNIX, Mac, Windows)

<http://www.cygnum.com/product/kerbnet-index.html>

B.6 FORTEZZA

FORTEZZA, created by the National Security Agency, is an evolving tool for implementing security services on federal networks.

- Multilevel Information Systems Security Initiative

<http://beta.missilab.com>

The MISSI Program's Beta Test Homepage

- FORTEZZA Card Information Page

<http://www.tcel.com/~sabbott/fortinfo.html>

This page aids in the setup and usage of FORTEZZA cards.

- FORTEZZA Libraries and Drivers

http://www.litronic.com/products/FORTEZZA/ft_soft.htm

The links below are linked to a site maintained by members of the FORTEZZA program office in the Department of Defense and has been established to provide an easy way to distribute FORTEZZA information to the public.

- FORTEZZA Documents

http://www.armadillo.huntsville.al.us/FORTEZZA_docs/

- FORTEZZA Crypto Cards

<http://www.armadillo.huntsville.al.us/general/fortcard.html>

- What's New at FORTEZZA

<http://www.armadillo.huntsville.al.us/whatsnew.html>

- FORTEZZA Software

<http://www.armadillo.huntsville.al.us/software/index.html>

- FORTEZZA-Related Links

<http://www.armadillo.huntsville.al.us/links.html>

- FORTEZZA General Information

<http://www.armadillo.huntsville.al.us/>

- FORTEZZA FAQs

<http://www.armadillo.huntsville.al.us/general/faq.html>

B.7 Other Links

These references point to other sites that we feel have merit and could provide additional resources to information not closely covered by our working group.

- DMS Home Page

<http://www.lmdms.com/>

An architecture and suite of products that will replace the DoD's messaging environment.

- NT Security

<http://www.ntshop.net/>

This site offers both security information and security tools.

Appendix C: Other Technologies We Did Not Address

The criteria for the selection of the CIS testbeds were somewhat arbitrary and highly influenced by the capabilities, needs, and desires of the CIS participants. Accordingly, there are other technologies that deserve further mention because they may be widely used, or because they are "hot" up-and-coming technologies.

C.1 Firewalls

Firewalls are in fairly widespread use, or perhaps misuse. While firewalls can serve as a valuable element of site security, they can be difficult to configure and can slow network performance. More alarmingly, a site may install a firewall and think that it solves all of their security problems. This is far from true. In the first place, many security problems originate from users inside the firewall, and in the second place; some attacks may masquerade as legitimate services that are allowed through the firewall. A simple example of this might be a virus that is embedded inside a mail attachment.

Nonetheless, a firewall or even a properly configured router should probably intercept all traffic (i.e., act as a gateway) between a site's internal network and the Internet. In particular, IP address filtering should be activated so that no internal IP addresses are allowed to go through the gateway from the external network interface. This configuration prevents IP spoofing attacks. In general, a properly configured firewall can effectively stop nuisance-level attacks, and that allows security personnel to concentrate on more serious problems and to spend more of their time on productive matters.

C.2 Technologies for Classified Systems

Many interesting technologies exist for high security environments; especially those designed for multilevel-secure systems that enforce both classification levels and need-to-know access restrictions. Some of the participants in the CIS study utilize these technologies, but felt that they were too expensive and narrowly focused to be included.

The CIS utilized one of the MISSI products the FORTEZZA card, but this represents only the first level of the MISSI-developed products. For example, the MISSI mail guard allows a secure means to connect systems at two classification levels.

Compartmented mode workstations (CMW) have reached a level where they can be used in place of the corresponding non-CMW UNIX flavor (e.g., HP-UX CMW 10.09 instead of HP-UX 10.20). By this statement, we mean that most of the COTS software products can be used at one security level without activating any special privileges. In particular, all of the large relational databases are available in a true CMW implementation.

For most users, the military security aspects of these systems may not be of great importance. However, there are many non-military uses for these systems, and government agencies should know about them because they are the correct solution for many high-security situations. For example, the Security First Network Bank

(<http://www.snfb.com>) uses an HP CMW system with two network interfaces as the interface between the bank and the internet.

C.3 Hot New Technologies

Several very promising technologies were not sufficiently developed when the CIS study started to be included in the testbeds, but we feel that these technologies are very promising, and point them out for consideration.

C.4 Authorization Certificates

In the past year, the SPKI (simple public key infrastructure) effort has submitted a draft standards document to the IETF. Most existing security technologies relate to file or resource access permissions. However, many security policies cannot be easily expressed in terms of file access permissions. Unlike ordinary PKI X.509 certificates that bind a public key to an identity, SPKI certificates bind a public key to an action.

Multiple SPKI certificates can be combined to express very complicated security policies. For example, to control an online facility, a user might have to present certificates certifying that he has received training (and passed a test), has paid for the session, has reserved a time slot, and that the information must be handled as proprietary. SPKI certificates can be used to enforce transparently many site policies. For example, if a user has not completed an annual computer security training refresher, his certificate for access to the computer system expires, and logon access is automatically denied.

C.5 IPV6

The new IP (Internet protocol) version 6 networking protocols allow for much larger amounts of security information to be included in the IP header than do previous protocols. This information can be used by routers to segregate information onto different sections of LAN. Different classes of service can be used to help protect against denial of service attacks launched by flooding the network with spurious packets.

C.6 Custom Applications

While we encourage the use of commercial off-the-shelf (COTS) software wherever possible, many government applications are "one of a kind" and must be created by those intimately associated with these projects and their requirements. In these cases, we encourage developers to implement strong security features in their software. Such implementations are made easier by the availability of software development kits (SDKs) for many major software applications and environments. SDKs are available for:

- All Netscape clients and servers
- Microsoft NetMeeting
- DCE, CORBA and DCOM (distributed object-oriented access to resources)

- Entrust certificate engines
- Java
- Databases