An Architecture for Flexible Multi-Security Domain Networks

Tim Gibson, Ph.D. *
Lieutenant Colonel, U.S. Army
United States Pacific Command
Camp Smith, HI 96861
tgibson@acm.org

* The views expressed in this paper are those of the author, and are not necessarily those of the United States, the Department of Defense, or the United States Pacific Command. This paper has been reviewed and released for publication by the Department of Defense Public Affairs Office, the Pacific Command's Public Affairs Office, and the Pacific Command's Computer Security Division.

Abstract – This paper briefly explains how the U.S. military currently implements secure networks internally and with multi-national alliance partners, the limitations of the current implementations, and proposes an architecture to overcome these problems. The proposed architecture provides a secure, environment that does not require all members to be treated as peers and allows different private communities. The proposed architecture is not necessarily the only or the best architecture, but is a starting point for discussion and provides the research and private computer communities with an insight into the military's unique problems. Many of the concepts or requirements discussed in the paper can be directly applied to the commercial sector.

1. Introduction

For the military, the post-Cold War era presents many different and complex problems that were unimagined during the previous four decades. These problems apply to many facets of the military-industrial complex, including how to apply cryptography to meet the needs of the military commander in the post Cold War political environment. Throughout the Cold War, there were clear definitions of which countries were our friends and allies and which countries were either neutral or not friendly. Under these Cold War alliances, if a country decided to release information to its alliance, all alliance partners had equal access to the information. These basic rules applied to both "sides" during the Cold War.[1] Today's political climate has changed dramatically from that of the Cold War. The lines between friend and foe have blurred and can shift over time or over particular issues. While this situation is complex enough for politicians and military commanders, it provides an entirely new level of complexity for those who provide secure computer and communications systems.

Many of these new security challenges are only now being grasped, and the complexities are far reaching. The advent of electronic mail and the World-Wide–Web (WWW) complicates matters even more. The United States military takes Internet-based electronic mail, web-browsing, and many other actions for granted. It is extremely difficult to exchange classified email with foreign countries using accredited and certified systems. It is more problematic to provide web browsing capability across security domains. The result is either separately encrypted bilateral networks with each nation or a large multi-lateral peer network. Bilateral connections have a high maintenance cost and reduce interoperability for multi-lateral operations. Similarly, one large peer network enhances interoperability, but the utility of a peer network is reduced because all information is implicitly available to all network members—and not every network member wants to share information equally. These problems are particularly vexing for the United States in the Pacific region. With over forty countries in the region, the Pacific Command includes countries that are old friends, new friends, neutral, old adversaries with improving relationships, and a few clearly unfriendly countries. The core problem is how to provide a system that is accessible from U.S.-only classified systems, provides different

[1] These general rules obviously have exceptions. For example, nothing keeps one alliance partner from having a special, unilateral, relationship with another partner that differs from the general purpose, multi-lateral alliance relationship.

levels of classified allied connections, and quickly allows partners to voluntarily—and involuntarily—connect and disconnect from the network.

This paper discusses these network security issues in the next four sections. Section 2 provides background information on typical data security methods developed and used by the United States during the Cold War, many of which are still in use today. Section 3 briefly discusses problems caused by our current implementation strategies. In Section 4, I propose an alliance data network that meets most or all of the requirements from Section 3. Finally, Section 5 concludes by discussing potential problems and by showing how the proposed security architecture not only applies to military alliances but has commercial applications as well.

2. Background

During the Cold War, the Soviet Union and its immediate allies posed the greatest threat to the United States and its partners. While the United States and its allies engaged in peripheral or proxy conflicts with the Soviets in Korea, Vietnam, Malaysia, Yemen, and other countries, the main focus was always on Europe. The United States closest Cold War allies were those in the North Atlantic Treaty Organization (NATO). Because of this focus on Europe and NATO, U.S. policy for release of classified information was arranged to support the NATO model. The basic assumptions for the NATO classified information release model are that information released within NATO is freely available to all NATO partners, and that partners do not connect and disconnect from the network—essentially, *the NATO network is a peer network with a constant set of known partners*.

In addition to the NATO alliance network, most NATO countries also maintained separate, private, classified network(s) for their internal use. With the advent of the Internet and the WWW, many military units also gained access to unclassified public networks. Because simple and reliable multi-level security operating systems were not readily available, having access to three networks (alliance classified, internal classified, and unclassified) meant having three or more separate networks with as many sets of user terminals. The United States was no exception to this arrangement, nor have any of these arrangements changed since the end of the Cold War. Table 1 shows the primary components of the current U.S. network security architecture.

Network	Classification	Encryption	Community	Peer Network
SIPRNET	SECRET	Transmissions bulk encrypted with military grade secret-key cryptographic equipment. LANs are unencrypted.	U.S. only	Yes
Classified Alliance Networks	Alliance SECRET	Same	Alliance	Yes
NIPRNET	Sensitive but Unclassified (SBU)	Same. LANs protected by firewalls.	U.S. only	Yes
Internet	Unclassified	TELCO encrypted.	All	Normally

Table 1—Different types of networks, SECRET and below.

In Table 1 there are four network security levels on four unique physical networks. The lowest U.S. government network security classification level is the internal U.S. government unclassified network. It uses TCP/IP as a transport and networking protocol and is called the NIPRNET (National IP Routed Network). Direct connectivity between the NIPRNET and the global Internet is achieved through several portals. The next U.S.-only network is the SECRET level SIPRNET (SECRET IP Routed Network). From a U.S. security viewpoint, most SECRET level coalition networks reside between these two U.S.-only systems. These coalition networks contain information the individual coalition members decide to release to the other coalition members.

Besides having separate networks for the different classification levels, U.S. and alliance networks are encrypted using secret key, bulk encryption devices between transmission nodes.[2] Special communications sections are responsible for maintaining the networks, the communications links, and the encryption systems. Users normally assume that any traffic on the network is safe from any outside eavesdropper. Obviously, encryption between transmission nodes does not address any insider threat because the local area networks are not encrypted, although they are physically protected.

[2] Secret key systems rely on a key that is known only to the parties sending and receiving the message; compromising a secret key allows all encrypted traffic to be read. These secret key systems are in contrast to newer public key systems that use large prime numbers to give every user a public key and a private key. Secret key systems typically have a higher throughput than public key systems. For additional information see [7, 8, and 9].

The first problem of this arrangement is the financial cost of accessing the multiple networks. Few military units or government organizations have the money to provide users with three or more separate network drops and terminals. Most units must choose which network they will use (NIPRNET, SIPRNET, or alliance), and provide computers to everyone on that network. Access to the other networks is provided through common use terminals. This arrangement causes problems because users must go to another location, wait for a free terminal, and log in. While this situation can be remedied with multi-level security (MLS) operating systems, these products are few in number, and are neither inexpensive nor easy to use.

Another difficulty of the arrangement becomes evident when a user tries to move data from one security level to the next. When a user wants to move information on a computer from one security classification level to a computer or network on another classification level, the bytes associated with the information may be moved in one of two ways. The data can be moved either directly from one level to the next via the time tested "sneaker net" or through a guard (manual or electronic) connecting the security levels. The former requires a user to copy files from one computer network to removable media (usually a floppy disk), and physically move the disk to a computer on another network with a different security classification. The manual guard mechanisms use the traditional "man in the loop" to verify traffic and pass it between the networks using a special workstation connected to both networks. The manual guard introduces the problem of the "man" being unable to place information in context. Additionally, a person can be overwhelmed during high traffic periods. The electronic guard mechanisms are more complex, have a higher throughput than manual guards, but have their own unique problems. These include strict formatting to pass through the guard and the ability for insiders to bypass the guard by removing the words and phrases the electronic guard looks for.

3. Operational Problems Caused by Existing Guard Policies

From an operational commander's viewpoint, the U.S.-only restriction on the SIPRNET can severely hamper connectivity with foreign coalition partners. Table 2 shows the connectivity a regional commander's headquarters has with alliance Y, a close U.S. ally with historic contacts. The table also shows the connectivity the regional commander's subordinate commanders (air, ground, and sea) and supporting bases in the continental U.S. (CONUS) have with alliance Y. In the example, alliance Y has access to secure telephones and has a SECRET level internal, alliance network. Only the U.S. regional headquarters has a guard connecting the alliance network with the regional headquarter's internal network.

	Regional HQ with Mail Guard	Alliance HQ with SECRET alliance network	U.S. Component HQ in different local domain	CONUS Support Units
U.S.-only SIPRNET email	X	Only with Regional HQ via Mail Guard	X	X
Alliance SECRET Network email	Via Mail Guard	X		
Web browse SIPRNET	X		X	X
Web browse Alliance network		X		
Full Internet/NIPRNET	X	X	X	X
Secure Telephone	X	X	X	X
AUTODIN Text Only Teletype System	X	X	X	X

Table 2—Interconnections of Different Headquarters and Networks.

The six shaded cells in Table 2 highlight the communications gaps between the U.S.-only and alliance networks. Secure telephone units and teletypewriters are the only areas that provide full interoperability. The lack of a web browsing capability is irksome, but can be overcome via email messages for status reports. Currently, the lack of email between military service level components, their operational subordinates, CONUS based support units, and any alliance network causes problems. Component commanders (air, ground, and sea) are the people who actually provide ships, aircraft, soldiers, and supplies to any alliance. As such, component commanders often need to coordinate directly with alliance headquarters and units. Similarly, CONUS based support units provide airlift, ground troops, supplies, and special units (*e.g.*, civil affairs units) and need to communicate directly with the alliance. The only way Component commanders or CONUS based units can

contact an alliance headquarters or unit is with secure telephones, through the 1950's era AUTODIN message system, or by establishing their own dedicated, secure connection with a separate guard.

Operational units usually have access to secure telephones and have some type of SIPRNET access, albeit without electronic guards. This means that sending long electronic messages (*e.g.*, Air Targeting Orders can be hundreds of pages long) from the alliance network to an aircraft carrier's U.S.-only network is problematic.

The lack of network connectivity shown in Table 2 is clearly a problem for U.S. operational commanders working with an alliance. The current solution is to provide some alliance level computers where they are needed, and to move information from one network to another via the "sneaker net." While this works, it is slow, inefficient, and prone to errors.

One final operational degradation is caused by the intermittent nature of most alliance networks. To understand and exploit any network's full capability users need to develop their skills through regular use. The U.S. and its allies have standing, classified, alliance networks in only a few locations world-wide. For the rest of the world, alliance networks are temporarily operational only during exercises or real world contingencies, so any lessons learned are generally lost.

The difficulties lie in solving these interoperability problems without compromising the security of internal U.S.-only networks, or any other country's internal network. The author proposes a solution to these problems in Section 4.

4. Alliance Network Architecture (ANA)

The proposed architecture, presented in Section 4.2, attempts to solve both the email and web-browsing problems by combining asymmetric encryption (*i.e.*, public key technology) for encrypting individual sessions and signing individual data objects, with symmetric encryption (*e.g.*, Virtual Private Network technology) on the transmission links. This arrangement provides strong identification of the user, strong authentication of objects, and strong encryption the transmissions.

4.1. Requirements for Alliance Networks

As mentioned earlier, most network research and implementation for multi-national networks was undertaken during the Cold War. Because of this, most current alliance and coalition network research and implementation plans continue to have an Euro-centric approach. This Euro-centric approach assumes all partners are peers, that network membership is stable, and that anything threatening the alliance threatens all partners. These assumptions are not valid in all regions of the world. Table 3 shows the requirements for an Euro-centric network and compares them with what the author believes the requirements are m to be.

The major difference between the proposed alliance network and current coalition networks is that the alliance network is not necessarily a peer network. While it can be a peer network, it does not have to be one. NATO networks and the proposed U.S. Department of Defense Coalition Wide Area Network (CWAN) assume network members always belong to the coalition and that every member has the right to access all of the data on the network. However, there are several geo-political regions where this is not practicable, the Pacific is one of these. The Pacific region has many "hot spots" and has many potential communities of interest. Military allies may not be on the same side of every conflict. There are often real-world operations or crises happening at the same time exercises are taking place and information is not allowed to flow freely between exercises and real-world—both because the exercise information may be classified and to reduce the possibility of real-world operational commanders mistakenly acting upon exercise-only information. As a result, alliance networks in the Pacific region must be able to service separate communities of interest simultaneously. Similar arguments can be made for networks in Africa, South America, and the Middle East. As a result, the basic assumption of current coalition network architectures—a stable membership of peers—is invalid for modern alliance networks.

Requirement	Euro-Centric Coalition Network	Alliance Network
Connection between national networks and coalition/alliance network	Yes	Yes
Peer network	Yes	No
Stable Network Membership	Yes	No
Provides Virtual Private Networks for different national groups	No	Yes
Allows for multiple crises or exercises with different classification groups and information requirements	No	Yes
Multiple security level and communities of interest	No	Yes, by combining asymmetric (public key) and symmetric (secret key) cryptography
Transmission links	Bulk encrypted with secret-key cryptographic device	Same
Ability to force member off network	Yes, by communication link termination only	Yes, by public key revocation or link termination

Table 3—Alliance Network Requirements.

4.2. Proposed Alliance Network Architecture (ANA)

Given the requirements in Table 3, it becomes a matter of finding a way to provide for them. The author believes this can be accomplished by combining asymmetric key advances with conventional symmetric key system. The proposed solution in this section concentrates on solving email and web-browsing, the most common shortcomings.[3]

The system uses several layers of complementing cryptography to achieve strong authentication and identification of the user and provide strong communications link encryption. The former uses asymmetric key technology to identify individual users to the network, digitally sign object being passed to the alliance network, and exchange individual session keys for bulk encryption of user packets. The communications link is separately encrypted using separate keys from the bulk encryption the user has generated for her individual session.

To enter the alliance network requires at least five separate steps, these are shown in Figure 1 on the next page. Prerequisites for a connection include Hardware Public Key Encryption Devices (HPED) and a separately encrypted communications link. The latter can use pure military grade point to point link encryption or can tunnel through unclassified networks using military grade encryption on the data portions of the transmission packets.

The first prerequisite is that all computers have a HPED. Whether the HPED is a Fortezza card, an allied manufactured equivalent, or something new is immaterial, as long as it can provide hardware based public/private keys and digital signatures.[4] The next requirement is a strongly encrypted communications link.[5] A virtual private network (VPN) is shown in the Figure 1. To enter the enclave, the user first enables the HPED with a password or biometric device. She then begins negotiating a trust relationship with the enclave guard through the VPN. The user exchanges digital certificates with the enclave guard (steps 1 and 2), and uses these to establish a secret symmetric encryption key for bulk encrypting transmissions later in the session. The current methods and algorithms used for negotiating and encrypting sessions on 128-bit Secure Socket Layer (SSL) sessions may be acceptable given the packets are being tunneled through a strongly encrypted VPN.

Once the individual user's session key is established and a bulk encryption session begins *inside* the VPN, the user can log into the enclave via the enclave guard with a userid and password. At this point the user has nominally

[3] There are undoubtedly some flaws in this proposed architecture despite the fact that it addresses all the requirements from Section 4.1. The author requests the reader's indulgence to not "throw out the baby with the bath water" and to use this architecture as a framework for discussing and solving the alliance coalition problem.

[4] Some readers may be surprised to find Fortezza cards mentioned here. Rest assured that these devices are alive and well in the U.S. military, providing a level of identification and authentication well beyond what software tokens can provide.

[5] Passing U.S. classified traffic currently requires military grade encryption. The best public domain and U.S. exportable algorithm, triple DES, is inadequate for transmitting classified data. This may change when a DES replacement is chosen.

entered the enclave. This is possible only after providing a HPED password (or biometric), having a correct public/private/digital signature on the HPED, providing a correct userid and password, and being connected to the enclave through the hardware encrypted VPN link.

Once past the enclave guard, the user contacts the arbitration server. The arbitration server decides what services and machines a user can access within the enclave. Because the enclave is not a peer network, every user will not be able to access every service or machine within the enclave. The arbitration server issues time limited certificates to users for access to the different servers and services within the enclave [2,3]. Additionally, every server in the enclave checks both the user's credentials and the authentication server issued certificate before access is granted. Figure 2 provides an architectural overview of an enclave and shows both the guard and arbitration server.

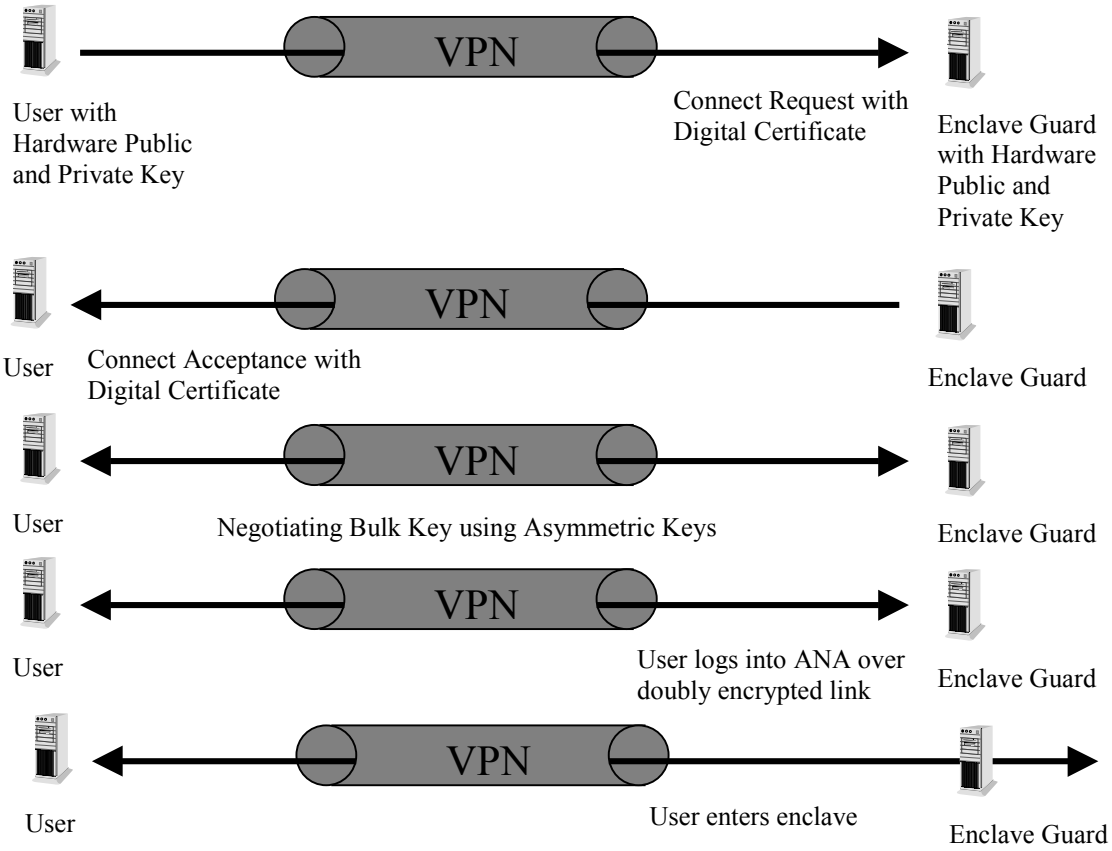

Figure 1 – Five Steps for Connecting to the Alliance Network

An additional requirement for the enclave guards is to prevent connections from one alliance partner's network directly to another partner's network (*i.e.*, no "back doors" through the enclave guards). This is shown in the figure with a U.S. user being refused a connection to the foreign network. This capability to prevent backdoor connections is required for all partners to trust the system. *Note: There is only one enclave guard per enclave; two guards are shown in the figure for illustrative purposes only* .

Systems similar in many respects to the enclave guards exist today in several different commercial products [6, 7]. Similarly, certificate issuing machines that grant certificates to authorized users also exist, or can be constructed using available commercial software [2, 3].

Email protection in the enclave is straightforward. Every email message sent to the enclave mail server is triply encrypted. First, with the recipient's public key and the user's private key (digital signature) from the HPED and a randomly generated key—using the same basic technique as PGP. Next, the message is sent through the user's encrypted session. Finally it is sent through the hardware encrypted VPN link. All email messages are stored on one mail server. Because each mail message is encrypted and contains the sender's digital signature and the recipient's public key, storing them on a single server should provide adequate security.

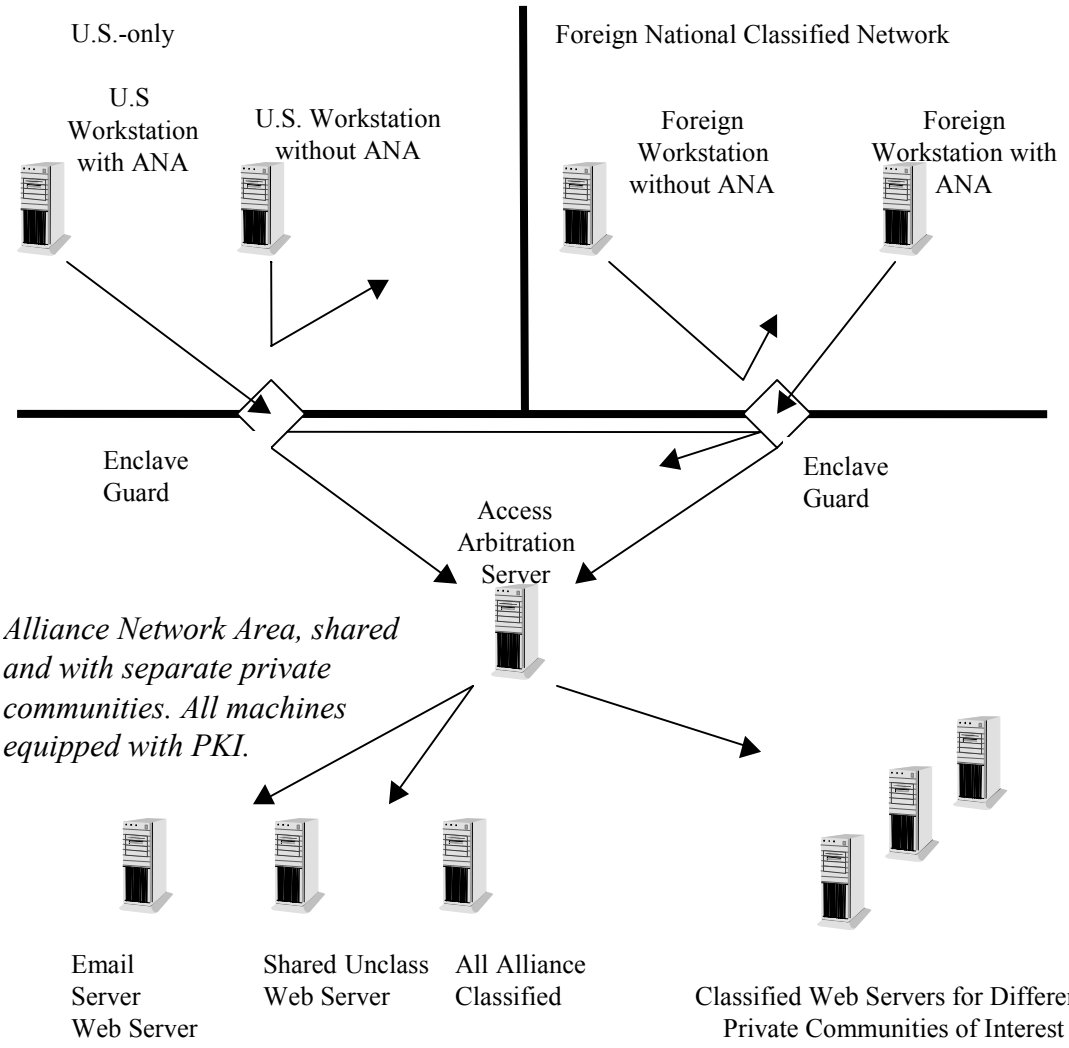

Figure 2—Proposed Alliance Network

To provide web browsing security, Web users also uses three layers of encryption, the session, the transmission link, and signing and classification marking by the person placing the object on the web server. The arbitration server allows users access only to those web servers they have permission to access. This allows communities of interest within the enclave. Placing items on web servers passes through the same two layers of encryption as browsing. Again, the arbitration server verifies access rights to post objects. Additionally, the web server itself also verifies that user has the correct permissions to post objects to the web. Items posted to the

internal enclave web servers are also digitally signed by the user posting the object. Applications accessing the enclave (web browsers, ftp, and telnet) must be public key/digital signature "aware" and verify the signatures on objects and compare them with the originator's digital signature as a "double check."

To explain the enclave access arrangement in greater detail, assume there are five states, New York (NY), Illinois (IL), West Virginia (WV), Virginia (VA), and South Carolina (SC). States *NY*, *IL*, and *WV* all currently belong to the alliance. *NY* is currently conducting an exercise with *IL*; *NY* is also providing *WV* with real-world border clash intelligence about *SC*, a non-alliance member. *VA* belongs to the alliance, but is actively is encouraging outsider *SC* against alliance member *WV*. Additionally, *VA* passed classified alliance information to *SC*. As a result, the alliance leadership decided to remove *VA's* full ANA privileges, with the exception of email because they hope to use *VA* to mediate with *SC* about the border clash. Table 4 shows the access privileges each country has within the ANA based upon this scenario.

	Valid Members	**Invalid Members**
Email (all types, classified and unclassified)	*NY, IL, WV, VA*	*SC*
Unclassified Web Sever	*NY, IL, WV*	*VA, SC*
Alliance General Purpose Classified Web Server	*NY, IL, WV*	*VA, SC*
Alliance Exercise Web Server	*NY, IL*	*WV, VA, SC*
Alliance Real-World Border Clash Web Server	*NY, WV*	*IL, VA, SC*

Table 4—Privileges Allowed to States *NY*, *IL*, *WV*, *VA*, and *SC*

Based upon the Table 4 privileges, the different guards and servers in the enclave can be configured to provide services to the appropriate users and countries. Table 5 maps the privileges in Table 4 to services and certificates issued or denied in an enclave. Services provided are identical to those privileges allowed in Table 4.

	Users Provided Service	**Users Denied Service or Access**
Allowed into ANA by PKI Guard	*NY, IL, WV, VA*	***SC***
Request Certificates from Arbitration Server	*NY, IL, WV, VA*	-
Issued certificate for mail server	*NY, IL, WV, VA*	-
Issued certificate for unclassified web server	*NY, IL, WV*	*VA*
Issued Exercise Web Server certificate	*NY, IL*	*WV, VA*
Private Exercise Server Session #1	*NY*	*IL, WV, VA*
Private Exercise Server Session #2	*IL*	*NY, WV, VA*
Issued Border Clash Server certificate	*NY, WV*	*IL, VA*
Private Border Clash Server Session #1	*NY*	*IL, WV, VA*
Private Border Clash Server Session #2	*WV*	*NY, IL, VA*
Issued certificate for general purpose classified server, each with encrypted separate sessions	*NY, IL, WV*	*VA*

Table 5—Alliance Network Area Connections Allowed by PKI Guards and Servers

U.S.-only SIPRNET users are not allowed *carte blanche* privileges in the enclave. SIPRNET users are required to have the correct authentication, assurances, and training to enter the Alliance Network, just like any other ANA user. All web-browsing and posting is conducted through the guards with the assurances and web object signatures discussed earlier. All email traffic is signed by the sender for authentication purposes and encrypted for security purposes. Given the level of protection provided by this architecture, allowing enclave guards to service multiple SIPRNET domains from more than one local area may be feasible.

The organization controlling the arbitration server obviously has complete control of all objects—both machines and files—within the enclave. While this centralization of authority can cause problems, this can be mitigated by policy. Additionally, there can be multiple enclaves (see Section 4.4), each with the arbitration server for that enclave under the control of the enclave's physical owner.

4.3. Revoking Keys and Access

One of the problems with the Euro-centric networks is removing users and organizations from the network. Currently, the only way to do this is to change the cryptographic keys or shut down the communications link. Neither method is trivial. Access to the entire ANA can be denied by simply refusing individual connections at the enclave guard by either locking the individual's ANA account or revoking the rights of their public/private key pairings. National level connections can be refused either in the same manner, by revoking the national groups' privileges, or by changing the hardware link encryption variables (in the example this was a Virtual Private Network).

The ability to disconnect individuals or national groups from portions of the ANA while continuing to allow access to other ANA components rests with the user/group accounts manager at the arbitration server. Limiting existing services with the arbitration server consists of removing access privileges to specific machines (each serving a separate community of interest) from the user or national group. The next time a newly restricted user requests a token for an "off limits" from the arbitration server, the request is denied.

An individual user in a particular country may belong to several communities of interest in the ANA. The individual must have a unique hardware token, userid, and password(s) to gain access to the enclave. Revoking some privileges may not completely remove the user from the system. Removing a country from access to the ANA will remove all of that country's users. However, once a user's or country's privileges are disabled does not mean they cannot be reinstated later using the same keys, particularly if the revocation was done at the user/national group account level.

4.4. Communications Infrastructure for the Alliance Network

The network and security architecture presented in Sections 4.2 and 4.3 detail the basic security arrangements within the Alliance Network Area and the requirements for users to gain access to the alliance area. It does not discuss how different ANA nodes communicate with one another or the user's communications path from their work area to an ANA. There are two ways to physically reach an ANA node. First, the alliance network can have a separate, encrypted, communications backbone. This is the current method, is very secure, and relatively difficult to attack from outside the network except through physical infrastructure attacks. It has the drawback of being relatively expensive. The ANA proposed in this paper does cost less than the current method because it allows a single connection to a multi-level enclave instead of requiring many bilateral connections.

The second technique is to exploit the commercial Internet. This is not commonly done now because it is difficult to control where packets pass while they traverse the Internet. For example, packets passing between two countries may conceivably pass through a non-ANA country. Internet based systems are also more vulnerable to denial of service attacks than separate backbone systems. However, by using the proper encryption techniques, the Internet may be used at a greatly reduced financial cost than a dedicated network. This was done recently during operations in East Timor [13].

5. Conclusion

This paper presents the requirements for an alliance network that does not require all alliance members to be treated as peers, and that also provides separate, private communities within the network. The network does this using a combination of symmetric and asymmetric encryption technology and existing or developing equipment. The proposed alliance network can provide a currently unachieved level of interoperability between foreign classified networks and U.S.-only classified networks.

There are several problems with the architecture proposed in Section 4.2. First. the proposed architecture proposed in Section 4.2 must be fully reviewed, accredited, and approved. These review and accreditation problems aside, there are several other technical and administrative problems that must be addressed. For example, there is the purely technical problem of writing enclave aware applications to administer and transfer files to machines located within an Alliance Network Area. For example, if a user wants to remotely transfer a file to a web server today, she uses the *File Transfer Protocol* (FTP) to log into the server and copy the file to the server. FTP programs that can negotiate through enclave guards do not currently exist. Similar problems hold for *Telnet* and other remote administrative tools. These technical problems can be solved and do not pose a significant problem.

The administration and management of a network using both symmetric and asymmetric encryption is daunting. Using HPED technology requires every person to have an individual hardware token with all of the appropriate certificates. For example, if a user is a command center watch officer and an official message release authority, both require additional certificates to be issued. Quickly expanding the number of properly equipped HPED users in a crisis may not be an option, so the enclave system needs to be established in peacetime.

Despite these problems, the alliance architecture proposed in Section 4.2 (or something similar to it) needs to be established for tomorrow's military commanders. Current systems do not provide the interoperability required between multi-national forces that we need on the battlefield. The advantage to any vendor developing these systems is that all the systems described herein have clear commercial value as well. The enclave guards, arbitration servers, and public key aware administration tools can all be used immediately in the private sector. Because of the dual utility, the author hopes researchers and commercial firms will develop commercial versions of these devices which the military can use to implement improved alliance networks.

References

[1] J. Hamre, Deputy Secretary of Defense Memorandum, Subject: Department of Defense (DoD) Public Key Infrastructure, Washington, D.C., 5 May, 1999.

[2] J. Kohl, and B. Neuman, *The Kerberos Network Authentication Service*. Network Working Group RFC 1510, 1993.

[3] J. Kohl, B. Neuman, and T. Tso, The Evolution of the Kerberos Authentication Service, *Distributed Open Systems*, IEEE Press, 1994.

[4] National Security Agency, *NSA Policy on Use of FORTEZZA for Protecting Classified Information*, Ft. Meade, Maryland, January 22, 1996.

[5] National Security Agency, *Security Concept of Operations for the Secure Network Server Guard Through Phase 2D*, Ft. Meade, Maryland, September 17, 1998.

[6] The Rainbow Corporation provides information on their cryptologic and Public Key devices at www.rainbow.com., October 25, 1999.

[7] The RSA Corporation provides information on their cryptological products at www.rsa.com., October 25, 1999.

[8] B. Schneier, *Applied Cryptography : Protocols, Algorithms, and Source Code in C*, John Wiley and Sons, New York, 1995.

[9] S. Singh, *The Code Book: the Evolution of Secrecy from Mary, Queen of Scots to Quantum Cryptography*, Doubleday, New York, 1999.

[10] United States Department of Defense, *X.509 Certificate Policy*, Washington, D.C., March 1999.

[11] United States Department of Defense, *Public Key Infrastructure Roadmap for the Department of Defense*, Washington, D.C., July 1999.

[12] United States General Accounting Office, *Report to the Secretary of Defense, Subject: DoD Information Security: Serious Weaknesses Continue to Place Defense Operations at Risk (GAO/AIMD-99-107)*, Washington, D.C., August 1999.

[13] B. Murray, *Government Computer News*, "U.S. Peacekeepers use Net to Access Classified Network," Volume 19, Number 11, 15 May 2000. (see www.gcn.com/vol19_no11/dod/1938-1.html)