PAPER

RETHINKING DEPARTMENT OF DEFENSE PUBLIC KEY INFRASTRUCTURE

Abstract

This paper identifies and examines problems with implementing a Public Key Infrastructure (PKI) in the Department of Defense (DOD). Some of these issues, such as certificate revocation list (CRL) distribution and technological immaturity have been noted in official DOD documentation and other studies without solution. Others such as directory management and operations in a tactical environment have been examined elsewhere. Here the problems of e-mail encryption and access control are discussed. Several solutions to these problems such as policy changes and PKI architecture modification are suggested. Yet the overall conclusion is that DOD PKI in its planned configuration will not achieve its operational objectives and may introduce security vulnerabilities and bureaucratic confusion where none existed before.

Author

Jason X. Hackerson

Organizational Affiliation

None

E-mail

hackerson@earthlink.net
hackersonjx@hqmc.usmc.mil

RETHINKING DEPARTMENT OF DEFENSE PUBLIC KEY INFRASTRUCTURE

Abstract

This paper identifies and examines problems with implementing a Public Key Infrastructure (PKI) in the Department of Defense (DOD). Some of these issues, such as certificate revocation list (CRL) distribution and technological immaturity have been noted in official DOD documentation and other studies without solution. Others such as directory management and operations in a tactical environment have been examined elsewhere. Here the problems of e-mail encryption and access control are discussed. Several solutions to these problems such as policy changes and PKI architecture modification are suggested. Yet the overall conclusion is that DOD PKI in its planned configuration will not achieve its operational objectives and may introduce security vulnerabilities and bureaucratic confusion where none existed before.

Keywords: Access Control, Certificates, Certificate Revocation, Public Key Infrastructure

I. INTRODUCTION

In a memorandum dated 06 May 1999, the Deputy Secretary of Defense directed that the Department of Defense (DOD) implement a Public Key Infrastructure (PKI) by October 2001. As it is currently envisioned and directed by the policy, the DOD PKI will not achieve full operational status and may introduce numerous bureaucratic and security problems.

The problems begin with the proposed scope of the PKI. As stated in the DOD Certificate Policy (Version 5.0 13 December 1999):

> "The DOD PKI must support five primary security services: *access control*, *confidentiality*, *integrity*, *authentication* and *technical non-repudiation.* The PKI supports these security services by providing Identification and Authentication (I&A), integrity, and technical non-repudiation through digital signatures, and confidentiality through key exchange. "

Identification, integrity, and non-repudiation are strengths of PKI. Public key infrastructure technology is not truly designed to effectively support access control and encryption. Access control requires additional overhead that is not effectively built into or considered a part of PKI. The scope of services denotes a desire to eliminate multiple user names and passwords that are required to access applications and systems without examining other ways to achieve "single-sign-on" capability. The public key encryption algorithm is considered a "strong" algorithm, but due to its ability to allow personal encryption that capability can lead to abuse of the network. In addition to the scope of the services, the policy directs that all Department of Defense users be issued a certificate. Eventually, this policy will apply to millions of DOD personnel. The difficulties of managing a single system with that many participants spread across a disparate enterprise have yet to be fully examined.

Another problem is the relative immaturity of Public Key Infrastructure technology [1]. While the theoretical concepts of public key based systems have been known for years, the technical components of a PKI for an entity the size of the DOD have yet to be produced or fully

researched. Unlike the commercial sector, the Department of Defense must conduct tactical military operations where communications bandwidth may be severely limited. The large size of the certificate revocation lists (CRL) that must be either accessed or distributed for the PKI to work and its potential for communications bottlenecks has been noted [2]. The solution is to wait for technology to catch up to the requirement. While there have been models developed that attempt to mitigate this problem, that research has yet be applied in the real world to see if the results can be implemented. Moreover, some of these research models were not published until October 1999, five months after the memorandum, and one month before the latest version of the DOD PKI implementation plan was produced. Therefore these latest ideals were not incorporated in that plan.

In addition, there are no set implementation standards in the commercial world. Incompatibilities exist between commercial providers of PKI products despite broad standards for certificate-based systems. The DOD conducted numerous PKI pilots on a small scale. Yet these pilots focused on one application or web site, not the combination of activities that are a part of the DOD network.

Finally, the proposed operational architecture places the onus of privacy and security on the individual when the DOD is an organizational and role-based enterprise where security is provided by unit-based components and services. In order to achieve this architecture, the DOD will have to develop whole new methods of managing and securing information and processes that formerly did not need the requirement. This transformation may open up the possibility of increased abuse and adversarial subversion of Department of Defense networks

This paper will provide a brief overview of Public Key Infrastructures. It will then examine some of the problems of the proposed DOD PKI and offer several possible solutions.

II. PUBLIC KEY INFRASTRUCTURE DESCRIPTION

A Public Key Infrastructure is made up of the supporting services needed to implement public key technologies on a wide scale [2]. These supporting services consist of certification authorities and certificate management facilities.

The Certification Authority (CA) is the entity responsible for issuing and maintaining certificates. A certificate is a “collection of information to which a digital signature has been affixed by some authority who is recognized and trusted by some community of certificate users”[2]. A certificate contains the owner’s identification information and public key value. It also holds the certification authority’s name, digital signature and a validity period.

The CA provides the assurance that the person or organization using a particular certificate is in fact the owner of the certificate. Additionally, the CA generates the user’s public-private key pair and maintains the certificate revocation list. The certificate revocation list contains the certificates that have been revoked for reasons other than the certificate’s expiration date. The entire CRL can be distributed to certificate users for status checking, but that method takes up too much communications capacity as the list grows in size. In order to reduce their size, CRL’s can also be partitioned into population groups. The CRL for each population group resides at a set of CRL distribution points. Each certificate then contains a marker telling which distribution point to check for revocation of that certificate. Unfortunately, in the DOD the potential population groups (each service and agency) may also be large, and synchronization across groups may prove difficult. Another way to reduce the size of CRL’s is to issue delta-certificate revocation lists. A delta-CRL is simply the list of added revocations and other modifications to the baseline CRL. Since the list of changes tend to be smaller than the

base CRL, this method takes up less bandwidth. A problem with delta-CRL's is that the user must ensure that the stored CRL's are secured in a trusted facility. Additionally, the transient nature of DOD personnel and operations may require constant delta-CRL dissemination. An alternative to CRL's is on-line verification, where the central directory of certificates is checked for validity of a certificate.

Certificate management involves the use of methods needed for distributing certificates and certificate revocation lists. Certificates can be distributed physically or electronically through an X.500 standard directory service. Positive identification is needed before a CA will issue a certificate. In a highly distributed environment (where a person requesting a certificate is remotely identified to the CA) it may be impractical for the certification authority to ensure the identity of its certificate users, the Local Registration Authority (LRA) acts for the CA. Specific functions of the LRA include changing attributes of subscribers, identifying subscribers, and accepting and authorizing requests for certificate suspension or revocation [2]. The LRA, which does not have the ability to issue certificates, is a conduit between a person eligible for a certificate and the CA.

The structure of the relationships between certification authorities determines the certification path. The certification path is the model of connectivity between the CA's that allows the user of a system to know that a particular subscriber to another certification authority is valid. The path is the chain of certificates and CA's that begins with the user and links him to a subscriber of another CA. In structures where there is only one certification authority, finding the certification path is simple. In a system where there are multiple CA's, electronic commerce for example, finding a path between subscribers to different certification authorities becomes more difficult [2].

The simplest means to achieve the required relationships between CA's is to establish a certification path using a general hierarchical structure. In this structure, a CA, accepted as root, issues certificates to subordinate certification authorities, who in turn issue a certificate to the parent. The subordinates, acting as dominant nodes, then issue certificates to other subordinate certification authorities, who in turn issue certificates to their parent and to more subordinate certification authorities. The structure expands until the subscribers are reached. Finding a certification path is straightforward since there is a path leading from each certification authority to a "root" certification authority, which is connected to other certification authorities. Since this structure is based on the mathematical tree, it can expand and serve a number of subscribers. If each certification authority certified fifty subordinate CA's and there were four levels of CA's, then 50^4 , or over 6,000,000 CA's can be used. The longest certification path this structure would be seven certificates [2]. The problem with this model is that each CA must trust that all other certification authorities are issuing certificates properly.

Another structure is the top-down hierarchical structure. In this system there is only one root certification authority. Subordinate CA's issue certificates to two or more subordinate authorities, but they do not issue certificates to their parent certification authority. There is only one certification path from the root certification authority to any end user. This model applies to naturally hierarchical organizations such as the Department of Defense. However, the root certificate authority must be entirely trustworthy since all certification paths include the root [2].

The DOD has adopted a variation of the Top-Down Hierarchy for its Multilevel Information Systems Security Initiative for the Defense Messaging System and its proposed Public Key Infrastructure. There will be a root CA operated by the National Security Agency and DOD CA's managed by the Defense Information Systems Agency (DISA). Subordinate

CA's will be located in regions in the continental United States, and possibly in Pacific and European sites. It is also "envisioned that the DOD will require separate CAs on each of its networks (e.g. Top Secret, SIPRNET, NIPRNET) similar to the current implementations today where identical PKIs are replicated on each network" [1]. The services and DOD agencies will handle registration of users.

III. PROBLEM DESCRIPTIONS

A. DIRECTORY MANAGEMENT

Many of the problems with the DOD Public Key Infrastructure stem from the requirement that everyone in the DOD, as well as contractors and vendors who conduct business with the department, have certificates [3]. This requirement results in a potentially huge centralized certificate directory that includes 1.4 million military personnel, 300,000 civilian employees, numerous servers and applications, as well as countless numbers of defense contractor and vendor personnel. There are numerous ways to technically mitigate this problem, especially given that the database could require less than 2 gigabytes of storage space. One method would be to make the central directory a metadirectory, a directory of directories that oversees several smaller directories [4]. A metadirectory allows for replication techniques that provide prompt and automatic updates of all directories across the enterprise, greatly simplifying management issues.

The problems will originate from the management policies of the directory. For example, if the central directory is a metadirectory, there is no guarantee that the external organizations, the contractors and allies, will allow the unfettered access to their own directories necessary for the advantages of the metadirectory to be fully realized. If the central directory is a single global directory, how individual organizations make changes to the directory is not clear. Methods and technology may not be available to ensure that one organization cannot access and modify another organization's data in the directory.

Certificate revocation lists are generated from the directories. The size of the lists will be addressed later. On-line verification of certificate validity is an alternative to CRL's, yet there have been no studies to determine the scalability and cost of this method when applied to a large directory. The DOD will establish policy for frequency of CRL distribution, but it may not be able to enforce that policy on the External Certification Authorities (ECA) of our allies and contractors. If an ECA does not publish a CRL as quickly as required by the DOD, there exists the potential for abuse by entities and individuals with revoked certificates.

B. ATTRIBUTES AND ACCESS CONTROL

The size of the proposed population considerably complicates the access control issue. To have effective access control, each individual must have a set of attributes that denote the applications and services the person has admittance to. Standard attributes of DOD certificates include a unique directory name, country, and organization, but these do not constitute enough information to determine access.

If attributes are included with the certificate they will need to be standardized across the Department. Role-based assignment of attributes requires enterprise-wide reconciliation. A quartermaster in the Army may be the equivalent of a supply clerk in the Marine Corps but it definitely is not the equivalent of a quartermaster in the Navy nor should every Navy

quartermaster have access to the same information. Another way to assign attributes is to base them on the individual's "need to know" the information accessed. How an individual's need to know is determined will be important if attributes are centrally managed. Within a large organization, need to know may be highly granular and the development of a secure mechanism for updating attribute information in the directory will be significant.

If the attribute extensions are used on X.509 Version 3 certificates, then any attribute modifications, such as those required by personnel transfer, requires revocation of that certificate and reissue of a new one. Given the mobility and rate of turnover in the military, managing attributes via a single certificate poses a severe certificate management problem. An alternative is attribute certificates (AC). An attribute certificate is a separate certificate that is bound to the user's identity certificate [5]. Attribute certificates have the advantage that only the attribute certificate is modified if attributes change. However, they also have several drawbacks. First, the assignment of attributes to AC's pose problems similar to those previously discussed when it comes to role or need to know attributes. Additionally, there is no single precise definition of how an AC is bound to an identity certificate. If attribute certificates are bound locally, at the application site, then there really is no difference from access control lists (ACL), locally controlled lists that bind identities to authorized access attributes.

Furthermore there are two procedures for presenting AC's to an application. The first method, the push model, requires the user to provide the application with the AC. This puts the attribute certificate under the control of the user. As with capabilities, procedures for revoking such AC's are unclear. If there is a list of revoked attribute certificates at the application site, then there is no real difference between an AC and an access control list. Under the second method, the pull method, where the AC's reside in the accessed application's local directory, there is again no significant difference from an ACL.

Finally, if a user needs different attribute certificates for each site or system, then nothing has actually been saved over the present, ACL-based system. An argument can be made that the problem of different passwords and user names will have been eliminated, but that feature has not been balanced against the added costs of managing multiple certificates for each user, particularly if all of these certificates are centrally located and administered.

An alternative to attribute certificates is trusted directories. Trusted directories are directories where attributes are securely bound to identity certificates. This alternative does not address the directory management problems detailed earlier, nor do they seem any different from well-managed locally administered access control lists.

Access control lists have been mentioned repeatedly in this section. Binding identities to the access control list for each system using certificates is a way to simplify entry for users. However, there remains the risk that administrators will relax their vigilance given their reliance on the certification authorities to revoke certificates of unauthorized users. Even under a PKI, administrators of the lists will need to maintain procedures to ensure that users who no longer require access to their systems are removed from the ACL, similar to the way systems are operated now. So while users gain some advantage from a single-sign-on capability provided by the PKI, the owners of the accessed systems will remain wedded to the procedures in place before PKI was introduced, while picking up the cost for PKI-enabling their applications and web sites.

C. UNIVERSAL SECURITY CLEARANCE POLICY

A security clearance denotes the level of information a user has access to. Under a system whereby everybody has a certificate, security in the context of network access control has not been greatly increased. For example, if one person has a key to a safe, then that safe is secure. If two people have a key to the safe, then the safe is still secure but less so. If 500,000 people have a key to the safe then it is no longer secure. Hence, if everybody now has a certificate that allows access, then it is a *universal security clearance*. Thus access to information is no longer granularly controlled: effectively, there are no security controls.

Use of certificates will reduce the ability of external entities to enter DOD networks; it will not eliminate access. Instead of focusing on usernames and passwords for access to applications or websites, crackers will target the certificates and the passwords required to utilize a certificate's private key. If certificates are used as universal access tokens, then once a certificate is secretly compromised, intruders will be able to gain access to even more information than before.

Under the universal security clearance policy, new classifications of information have been created. Under the common name of "sensitive but unclassified (SBU)", health, personnel, and financial information have been put on a level similar to secret and top-secret security clearances. By their nature the SBU materials are similar to the more restrictive sensitive compartmented information (SCI) data. SCI data is divided into compartments where access to that information is restricted by need to know. For example, two people may have top secret clearances, but only Alice has need to know access to information in compartment 1 whereas Bob only has need to know the data that resides in compartment 2. Alice never sees compartment 2 information, and the same goes for Bob with compartment 1. In the SCI realm access to a compartment is generally determined by billet.

The SBU environment is similar in that only medical personnel have access to health information. The SCI sector consists of separate policy, guidance, procedures, and organizations than the standard secret and top-secret world. The Department of Defense may not be willing to institute the same level of effort for SBU data which involves personal privacy not national security. Procedures will be needed for those individuals who have their certificates revoked for administrative or criminal reasons but still require on-line access to their health, personnel, and financial records. Other procedures will be required to allow military commanders access to their subordinates' personnel records: an access control issue. The systems that manage sensitive personal data cannot be blocked off from the PKI lest they lose the advantage of individual authentication that PKI brings. Yet in order to return granularity of access control they will have to institute additional discretionary and mandatory access control policies as well as add the complexity and cost of PKI.

D. SOPHISTICATED ENEMY

A sophisticated enemy poses numerous concerns that are not eliminated by a public key infrastructure. While low-level, relatively unsophisticated hackers may be blocked out of a PKI enabled network, well-financed crackers and malicious organizations will still find their way into the network.

Building a trust relationship for certification paths between the DOD CA and external certification authorities is a vulnerability. As mentioned earlier, there is no mechanism that will allow the DOD to enforce the frequency of CRL distribution of the ECA's. Other concerns include the disgruntled or criminal employee at one of the services, agencies, or at one of the countless vendors, contractors, and subcontractors that interact with the DOD network on a daily basis. This group will now expand to include not only the weapons systems manufacturers and traditional defense contractors, but also the food supply vendors and commercial shippers who must be a part of the network in order to gain the cost efficiencies of electronic transactions. Many of these companies' employees will have "trusted" DOD compatible certificates. It is not inconceivable for a compromise or sale of certificates to occur. Under the universal security clearance policy, non-vigilant components of the DOD network will allow unintended access to information. Moreover, the breach will be harder to identify and classify, since all accesses have the stamp of PKI "trust" on them.

Another avenue that unscrupulous people can use to take advantage of PKI is the use of e-mail encryption. The DOD PKI Policy [6] calls for encryption certificates to be issued along with the identity certificates. Additionally, it encourages the encryption of e-mail using the certificates. If individuals are allowed to encrypt their e-mail, then they can send inappropriate or stolen material across the network. Electronic mail monitors will not be able to scan the messages since they would be encrypted with the public key of the recipient. The messages could be decrypted if the e-mail servers have access to each individual's private key. Although key recovery is a part of the PKI, the procedures usually require legal action and/or the presence of multiple individuals to unlock the private keys. Private key recovery is a cumbersome process not designed for the constant monitoring of e-mail. Any procedure less complicated may weaken the security of the keys and the confidence in the key recovery process.

A more insidious vulnerability may be viruses hidden within encrypted e-mail that will be undetected by the monitors. Such viruses could be easily spread using the public keys found in the user's e-mail address book. It would not be surprising to find that a virus has been developed whose sole purpose was to copy and disseminate the private keys of each person who decrypts the e-mail.

E. TACTICAL OPERATIONS CONSIDERATION

The DOD Public Key Infrastructure must support the deployment of forces. Unfortunately, the amount of communications bandwidth available to deployed forces is limited. It is hard for fixed wireline communications to move with forward units. Bandwidth, range, cost, and host nation regulation limit radio communications paths. Distributing the CRL over these communications path will compete with the need to transmit operational and intelligence data. Even if the bandwidth of the paths were to increase, it is more likely the additional capacity will be taken up by real time imagery feeds and the ability to videoconference.

To understand the size of a certificate revocation list, the Marine Corps, the smallest service, will be used to illustrate an example. A recent estimate for the Navy/Marine Corps Intranet states that the Marine Corps has 68,000 regular computer users. David Cooper, of the National Institute of Standards and Technology has developed a "sliding window method" of delta CRL distribution [Appendix A] that could theoretically reduce the peak bandwidth required by 99.5%. Using his calculations and the 68,000 users as a baseline it is determined that the Marine Corps portion of the PKI will need 2.03 Mb/s of bandwidth under the traditional method

of CRL distribution [Appendix A]. Using the delta-CRL method, that requirement drops to 1.89 Mb/s per second. If it were possible to build an optimal system using the sliding window methodology, then that number would drop to 9.4 Kb/s.

These peak bandwidth numbers are not unreasonable and easily met on shore and probably in a tactical environment. But those numbers are low. 68,000 is an approximation of the number of computer users in the Marine Corps. Under the PKI policy everybody in the Marine Corps will have certificates, and as technology spreads, more of them will require constant access to a network. There are currently 172,000 Marines. Using those numbers, the bandwidth requirement goes up almost 300%. With civilian Marines and reservists another 50,000 certificates are added. Include a joint operation, and the CRL's for the Army, Navy, and Air Force are added: 1.2 million more certificates. Finally, bring in the supporting establishment, agencies, and contractors, and the bandwidth requirement goes up by approximately 2000%. These calculations do not include the requirements of our allies. While the bandwidth numbers are not unreasonable for a fixed plant establishment with access to high bandwidth communications facilities, for a deployed force that number may become overwhelming. The requirement to communicate with rear forces and vendors will continue to increase as defense logisticians push to reduce the amount of material piled up in a theater and rely more on vendors and contractors to provide supplies and services directly.

The preceding paragraphs only examine the variance in population. They also assume an upper limit of the revocation ratio (percentage of certificates revoked) of 5% a year. Mitre in its 1994 study of PKI estimated a revocation ratio of 10% [7]. The numbers used to approximate the requests per second of the CRL could also rise similarly.

Tactical Certification Authorities (TCA) deploy with the forces during long duration operations. A tactical CA that managed only the certificates in theater would reduce "reach-back", the need to access garrison based CRL's and other certificate management functions [4]. There are several drawbacks to TCA's. First, they duplicate functionality. While that is the point of a TCA, the cost of duplicating the equipment and personnel has yet to be evaluated. Second, TCA's must ensure that they manage the right certificates of the personnel and equipment in the theater of operations. In a fast paced environment where units and personnel constantly transfer and people are killed and captured, it would become easy to lose control over what certificates were present. Nor may the system be prepared to handle the increased pace of revocations. Finally, TCA's do not adequately address the tactical-to-garrison support communications requirement of the combat service support community. Once communications with United States-based organizations are needed, the certificate revocation lists can grow again into an unmanageable size.

Another concern with the tactical employment of PKI capability is the loss or capture of personnel and equipment. In addition to the normal delay when first attempting to identify the loss, there is no automatic link from personnel management systems to the PKI. A person or device's certificate could be in an adversary's hands for several days before being revoked. Only a password protects an individual's private key. As most passwords are eight characters or less, it is easy to imagine a captured certificate device allowing access to the DOD network in less than a day if not in hours. And, as future conflicts move to urban environments, getting outside access to the network during a conflict may not be that difficult.

IV. POSSIBLE SOLUTIONS

Several of the problems described here, particularly CRL size and distribution may solve themselves as technology advances. Faster processors and new methods of wireless communications will make the size of the directories and the CRL's irrelevant. However, the technology is not available yet. There is no guarantee that when the technology becomes available it will be distributed in a timely manner to DOD components. The cost of such capability may be too high. There are other ways to attack the PKI problem that do not rely solely on technology.

A. POLICY CHANGES

The quickest and easiest solution may be to change certain policy declarations from the Deputy Secretary of Defense's memorandum. Eliminating the requirement that all DOD personnel receive DOD certificates is a logical step. A significant portion of the DOD does not have the authority to conduct official business requiring digital signatures and authentication over the network. Unfortunately as technology progresses, solid identification and authentication will be necessary as more people gain access to the network.

However, even with that situation, there is no need for everyone to have encryption capability. Eliminating individual encryption certificates will reduce the problem of electronic mail abuse. There are ways to monitor encrypted e-mail. But they would require severe relaxation of key recovery rules as well as additional computer servers to conduct the operation in a timely manner. Even under such a system, private keys would remain vulnerable to viruses activated upon decryption of a message.

Another policy change would be to decrease the life of the certificates to two years or eighteen months. Reduced lifespans allow revoked certificates to be removed sooner from the CRL. This has a direct impact on the size of the certificate revocation list. While shorter certificate lifespans will increase the amount of time spent administering the central directory and certificates, this expense should be evaluated against the cost of CRL communications bandwidth required to support tactically deployed forces.

B. PKI ARCHITECTURE

Another way to alleviate some of the problems of PKI is to reorient the architecture. The present PKI architecture is based on centralized certificate management and decentralized user registration [1]. If the architecture were built around a centralized policy and audit system and decentralized certificate management and registration, directory management and certificate revocation list size can be better managed. This architecture is similar to many PKI's constructed in the commercial sector [8].

The decentralized architecture would continue to have a root CA managed by the NSA and the Defense Information Systems Agency, but each service and agency would also have CA capability. The services and agencies can then designate subordinate CA's to the lowest level where they have control over individual data transmission. These individual control nodes (ICN) are the points where PKI is enforced. In many instances the ICN's will equate to individual commands in the military that control the flow of data out of their organizations. For example at

Unit X, when Bob transmits e-mail or visits the web on official business, the data travels first through Unit X's web and e-mail servers. Unit X is an ICN. It maintains its own CRL. If Bob tries to transmit without a valid certificate, his request is denied at the server. Likewise, if Bob is to receive data without a valid public certificate that transmission is also stopped at Unit X's server. While Bob can attempt to send official traffic via an outside service, the receiving server would not allow the request because it originated from an unauthorized domain. Server certificates can validate authentication of official domain servers. Individual certificate holders would receive certificates capable only of identification and digital signature. The organization's servers would handle decryption and encryption of information.

A "common" DOD directory [1] with public key values and other public information can still be maintained. But even if a certificate has been revoked, the intended recipient still will not be able to receive any data, since local revocation is relatively instantaneous. People who work from remote locations would be required to access the network through approved organizational access servers.

This system places the responsibility of security on the local units. Trained individuals in local organizations will be more concerned about security and bring a higher level of competency to routine information handling. Local units already are the ones to initiate revocation procedures under current network architectures. Updates to the directory can be sent on a leisurely basis, without fear of abuse from revoked certificates. The central certificate directory can focus on maintaining the certificates of the servers on the network, a number that will be in the thousands, as opposed to the millions of individuals. Since users of the certificates are now required to access the network via ICN servers, more security is added under this architecture.

There are some disadvantages to this architecture. Most importantly, the services and agencies may not want the extra responsibility and its attendant costs. The current PKI architecture is based on this premise of minimal additional expenditure of manpower and resources [1]. Costs include providing thousands of approved CA equipment to the ICN's. This expenditure may be reduced if the equipment is bought in volume. Other expenses include personnel training. This is alleviated somewhat by the fact that there are personnel in these organizations already trained in personnel and communications security. Software can and has been developed to simplify certificate management.

Other issues include the level of the subordinate CA's. A reasonable evaluation would equate the lowest level ICN's with individual naval ships, army and Marine Corps battalions, and Air Force squadrons. Agencies that currently have their own domain such as hqmc.usmc.mil, and spawar.navy.mil would be other candidates. A common network across the major services and agencies will facilitate oversight and enforcement of PKI policies, but this architecture will also work in an environment where individual entities control their own networks.

C. SINGLE-SIGN-ON CAPABILITY

The decentralized architecture does not solve access control problems. A system based on identity certificates and locally managed access control lists for the applications and web sites with streamlined registration procedures may be the simplest and most cost effective solution. Commercial vendors also provide directory based single-sign-on services [9].

Another potential resolution is to develop client applications that facilitate

single-sign-on for users. It could be easier to develop a browser plug-in that allows access control servers at the agencies, applications, and websites to securely download name, username, password and attributes to an authorized user on his own client system. There it can be secured via the user's own password and encryption.

When a user attempts to access a website or application the access control server at the site prompts the client browser for the correct information. The browser (the single-sign-on application) then asks the user for permission to release that information for the particular site or application. The user authorizes this by entering in the one password he needs to access his own data. Hence single-sign-on is achieved. The single-sign-on application would operate on a push basis. Instead of the requesting site searching through the file for the correct information, the application matches the site name with the correct name in the file. Once the match is achieved, the single-sign-on application pushes that information to the site.

This type of application can be moved to a smart card based system. Security is essentially the same as that achieved by PKI, as the user still needs a password to access the private key used for authentication. Here there is no individual user PKI overhead. Most of the current non-Public Key infrastructure remains the same. The handshake procedures can use the same encryption techniques that are used now such as SSL or the servers can use certificates to establish a secure link. The plug-in can be updated rather easily across entire networks.

Drawbacks to this system include weak security protections on client systems. Moving the application to a self-contained smart card can partially mitigate this problem. Different types of encryption algorithms can be used on the client system, so that the user isn't subject to the standard encryption available with the client operating system. Procedures will still be needed to get the initial access, but these procedures are in place now, and will remain even with the adoption of a PKI. A decentralized PKI architecture can assist in this process. The ICN servers can establish trust; the individual only has to present the certificate to the local ICN in order to verify the user's identity across the link.

V. SUMMARY

The Department of Defense will institute a public key infrastructure. Despite its shortcomings, "conventional wisdom" believes that having the PKI in place is ultimately better than nothing. The problems that have been identified in the DOD PKI Roadmap and PKI Implementation plan, such as technology immaturity and CRL distribution, will not be solved by a target architecture designed to "minimize the investment as well as the manpower required to manage and operate the PKI" [1]. Due to complexity and cost, it is foreseeable to see exceptions to the PKI implementation requirement granted. Exceptions will weaken the public key infrastructure and ultimately reduce its utility in the Department to a few specialized transactions between relatively small numbers of DOD personnel.

This paper has described additional problem areas with the proposed DOD PKI that have not been previously addressed. Future technology will solve some, but not all of these problems. The solutions suggested here are at best imperfect but they attempt to mitigate some of the more critical security and implementation implications. Serious discussion should begin on whether the scope of DOD PKI is appropriate and if there is a more effective way to achieve its security aims.

REFERENCES

[1] *Public Key Infrastructure Roadmap for the Department of Defense Version 3.0*, DOD Public Key Infrastructure Program Management Office, 29 October 1999.

[2] Baum, Michael S., and Warwick, Ford, *Secure Electronic Commerce*, Prentice Hall PTR, Upper Saddle River, NJ, 1997.

[3] *DOD Certificate Policy Version 5.0*, DOD Public Key Infrastructure Program Management Office, 13 December 1999.

[4] "Preliminary Roadmap For The United States Marine Corps Public Key Infrastructure", Dan E. Morris, Major, United States Marine Corps, David W. Rowe-Captain, United States Marine Corps, United States Naval Postgraduate School, September 1999.

[5] Attribute Certificates, Lisa Pretty, Baltimore Technologies, presented to the National Institute of Standards and Technology downloaded from http://csrc.nist.gov/pki/twg/presentations/twg-99- 67.pdf)

[6] Deputy Secretary of Defense, *Memorandum Subject: Department of Defense (DOD) Public Key Infrastructure (PKI),* 6 May 1999.

[7] S. Berkovits, S. Chokhani, J. A. Furlong, J. A. Geiter, and J.C. Guild. *Public Key Infrastructure Study: Final Report.* Produced by the MITRE Corporation for NIST, Apr. 1994

[8] From information downloaded from Novell, Inc. http://www.novell.com.

[9] From information downloaded from Novell, Inc http://www.novell.com/products/sso/

[10] " A more efficient use of Delta-CRL's (Draft)", David A. Cooper, Computer Security Division, National Institute of Standards and Technology, Gaithersburg, MD, 20899-8930, 25 October 1999.

APPENDIX A

David Cooper describes a method of CRL distribution using sliding window techniques with delta CRL's to reduce the bandwidth requirement during peak accesses by 99.5% [10]. Using his numbers (300,000 entities checking 10 certificates a day and 1000 certificates revoked per day with a lifetime of 365 days) , he produces a bandwidth requirement of approximately 270 Kbytes per second or 2.16 Mb/s, more than a T-1 line. Unfortunately, DOD certificates expire in 3 years, which can potentially triple the amount of bandwidth needed.
Given his original numbers, it will be hard to find that much bandwidth below a major command level. Even at that level there will be competition for bandwidth that might interfere with the ability to conduct other business over the network. Also, the work is largely theoretical, the

sliding window system has not been constructed, and the frequency of delta-CRL distribution needed in order to achieve the optimal bandwidth usage may be too infrequent for military purposes.

Calculations

Bandwidth required =
$CRL_{requests/sec}$*[CRL $Header_{bytes}$+ $Certificate_{bytes}$*$Revocation\ Rate_{per\ day}$*Certificate $Age_{\ days}$]

Numbers Used- Assuming CRL updated once a day

CRL Requests/sec = 68000 users receive 5 e-mails each day = $4_{requests/sec}$ ($68000*5/86400_{seconds\ per\ day}$)
CRL $Header_{bytes}$ = CRL Header size = 51_{bytes} from MITRE Report
$Certificate_{bytes}$ = Certificate size = 9_{bytes} from MITRE Report
$Revocation\ Rate_{per\ day}$ = Revoked Certificates/day = 5% $_{revocation\ rate/year}$ = $13_{per\ day}$ (revocation occurring on working days)
Certificate Age = Average age of certificates in days = $3_{years}/2 = 1085_{days}/2 = 542.5_{days}$
(Current life of each certificate is 3 years)

Bandwidth required using traditional distribution methods =
$4_{requests/sec}$ * [51_{bytes} + 9_{bytes} * $13_{per\ day}$ * 542.5_{days}] = $253890_{bytes/sec}$ = $2031120_{bits/sec}$ = 2.03 Mb/s
Delta-CRL = 7% reduction = 1.889 Mb/s
Sliding Window = 99.5% reduction = 9445 b/s