

Experiences Establishing an Experimental International Coalition Public Key Infrastructure

By Glenn Fink (*Naval Surface Warfare Center, Dahlgren VA, finkga@nswc.navy.mil*),
Shawn Raiszadeh (*Lockheed Martin Corporation, Fairfax VA, US, shawn.s.raiszadeh@lmco.com*), and Timothy Dean (*QinetiQ, Ltd., Malvern, Worcestershire, UK, tbdean@qinetiq.com*)

Abstract

Research and testing teams from the US and UK participated in joint design and testing of a Public Key Infrastructure (PKI) for international military coalition operations. We planned the design and testing in five phases from an initial PKI interoperability study through design of a second-generation PKI based on web services. Each design phase is followed by a testing and demonstration event to verify and recommend improvements to the system designed.

The paper opens with a description of the unique set of requirements an international military coalition must levy on its PKI. Next, we briefly describe each design and testing phase to give the reader a sense of context. This paper documents experiences with PKI technology that our research group had during the two most recent testing phases, II and III. We have included design and test-structure information for these two phases and highlighted our lessons-learned. We conclude with our current plans for future phases of the study. The intended audience for this paper is experienced PKI users, vendors, and researchers. We hope our findings and recommendations will be useful to the scientific community as we attempt to enable solutions complex problems through technology.

Keywords: Public Key Infrastructure, PKI, Security, International Military Coalition, Authentication, Nonrepudiation.

1.0 Introduction

The Virtual Operations Network (VON) project is an international military effort to facilitate management of naval coalitions involving forces from many nations. Teams of researchers from the UK (QinetiQ in Malvern and Portsmouth West) and the US (Lockheed Martin (LM) in Fairfax, Virginia, and Naval Surface Warfare Center (NSWC) in Dahlgren, Virginia)¹ joined together to form our VON PKI research group.

1.1 Unique Requirements of Coalitions

Coalitions in operations like Desert Storm and East Timor have demonstrated that traditional solutions for communications among a diverse group of coalition partners require an unsatisfactory amount of time and effort to establish and maintain. Some of the communication problems arise from equipment and software incompatibilities. Other communication issues come from the inability to trust once communications are established. Part of the VON effort involves establishing a degree of trust to facilitate information-sharing among coalition partners that are not traditional allies or may even be traditional adversaries. This project is complicated by the dynamic nature of modern coalitions where members may join for a relatively short period of the overall operation and may change roles during the operation. Nations participating in international coalitions come from a broad spectrum of technological ability—from low-tech, third world nation-states to technological superpowers. To level the playing field, nations with technology advantages may have to provide “throw-away” PKI components and services to their disadvantaged partners. While it is likely that the US or one of its high-tech allies would host some of the coalition PKI, it is essential that any nation, including the PKI hosting nation, be able to walk away from the coalition at any time without leaving indispensable personnel or sensitive equipment behind to maintain the PKI. Any equipment that must be left behind must be highly tamper-resistant to prevent technological espionage.

The coalition PKI should be accreditable by various nations. This implies that nations can be assured that none of their national secrets will be released into any associated coalition without the nation’s explicit consent. Accreditation generally requires presenting evidence that the risk is sufficiently low to make it worth the information gained. Accreditation also influences the amount of time taken to establish a coalition. The coalition PKI may be able to reduce this

delay by selecting standard or pre-approved hardware and software packages.

Because the partner nations are quite independent, a coalition PKI must have decentralized management of trust. Some partners may already have national PKIs, and most have national secret networks. Each of these partner-nations will want full access to information from the coalition PKI but tightly control the flow of national data into the coalition.

In military operations of all sorts, timely authentication and nonrepudiation are mission-critical requirements. PKI clients must be able to determine the validity of digital signatures quickly with a high degree of certainty that the status is up to date. Hardware tokens are envisioned for this application so that nonrepudiation may be more reliably achieved. For timeliness, we plan to require revocation windows of less than an hour.

The planned coalition PKI will run on shipboard platforms communicating over High Frequency (HF) or Ultra-High Frequency (UHF) radio links with extremely limited bandwidth and intermittent connectivity. In each battle group there will be one or more "gateway" ship(s) with satellite communications (SatCom) capability that will connect battle groups to the shore-based Network Operations Centers (NOCs). The NOCs may be nationally or internationally owned and will interconnect via secure, fixed links. PKI applications that cannot operate correctly under circumstances of intermittent connectivity and low bandwidth need not apply. Figure 1 is an overview of the communications concept used by VON.




Figure 1: VON Communications Concept

1.2 The Experiments

Our research group conducted experiments with two separate PKIs during Fall 2000 and Summer 2001. During the test periods, laboratories at four

geographically separated sites hosted simulated tactical platforms (ships and NOCs). The platforms were interconnected by dial-up ISDN lines simulating radio frequency (RF) transmission speeds. This effort was a step towards deploying PKI technology in a multinational at-sea trial in 2002.

The VON PKI effort can be divided into five phases of experimentation leading toward eventual deployment in operational environments:

- Phase I Lockheed PKI Interoperability Study
- Phase II UK-US Joint PKI Interoperability Test
- Phase III UK-US Joint Proof of Concept
- Phase IV Multinational At-Sea Prototype Trial
- Phase V 2nd Generation PKI: Web Services

Currently the project has completed Phase III and some initial testing for Phase IV. We will complete Phase IV during Summer 2002. This report documents experiences our joint research group had while fielding experimental PKIs during Phases II and III and will outline plans for the following phases.

Phase I was conducted by the LM team according to requirements defined by NSWC and the Office of Naval Research (ONR). LM evaluated five PKI certificate management systems (CMSs) and two Lightweight Directory Access Protocol (LDAP) directory-server products, given the requirements we had defined. The team simulated a three-nation coalition PKI using three different PKI vendors. This study is documented in [1] and is not further expanded upon here, but findings from it form the basis for the phases that followed.

Phase II testing occurred in the Fall of 2000 (September through early December), with the focal testing events conducted 13-17 November. The purpose of Phase II was to test the work done in [1] in a truly international setting. This was our first bilateral experiment in the PKI school of hard knocks. Two Certification Authorities (CAs) were set up, Netscape Certificate Management System in the US and Baltimore UniCert in the UK. We achieved limited PKI interoperability by maintaining a trusted lists of CAs in the clients. Parties successfully exchanged and verified signed and encrypted e-mail (sans attachments), and, with mixed success, visited each others' SSL-secured web pages. We also established secure network tunnels (via Internet Protocol Security (IPSec)) but used only static keying without automated enrollment via PKI.

Phase III testing was conducted in late summer of 2001 (July through August). The purpose was threefold:

1. To centralize trust management at the national level (as opposed to each user managing trust lists individually),

2. To reduce risk of component or system failure during the planned at-sea trial during Phase IV, and
3. To incorporate hardware tokens (smartcards, etc.) for end user credential storage.

Phase III testing was focused on cross-certification, exchange of S/MIME e-mail with attachments, and revocation testing. Both nations setup their own root CA (the US used Entrust, and the UK used Baltimore) and the teams cross-certified the two PKI domains. The 2001 testing period is believed to have been the first time that government/military organizations from different countries successfully established trust between independent national PKI domains using different vendor products. Participants at four separate sites exchanged, validated, and read digitally signed and encrypted email messages, proving the interoperability afforded by the coalition PKI.

The Phase IV at-sea trial will exercise the PKI configuration established and refined in earlier phases. This phase may involve more nations and will be on actual rather than simulated shipboard platforms. This phase should be completed by the end of this summer.

Phase V will incorporate the knowledge gained during the at-sea trial and attempt to define a middleware prototype that will standardize the application program interface to the coalition PKI regardless of the underlying PKI structure. This phase will rely heavily on Extensible Markup Language (XML) technologies, especially XML Key Management Specification (XKMS) and Security Assertion Markup Language (SAML). Work beyond this phase will probably involve further interfacing the coalition PKI with national PKIs and the multitudes of policy issues that arise from these interfaces.

2.0 Phase II Experiments

2.1 Objectives of Phase II

The overall objective was to set up a simulated coalition communications infrastructure and PKI to test interoperability results obtained during the study done in the previous phase. The supporting objectives of this experiment were:

1. Build a simulated RF shipboard network using ISDN links and RF simulators.
2. Establish TCP/IP (e-mail) connectivity.
3. Standup national PKIs and establish coalition trust via trust list.
4. Exchange signed and encrypted e-mail.
5. Test mutual web-server access and SSL.
6. Test publishing certificates to an LDAP directory and test remote LDAP replication.

7. Experiment with certificate issue, revocation, reissue, and CRL distribution.

2.2 Testbed Configuration for Phase II

The testbed for Phase II consisted of a wide area network (WAN) of computers using ISDN as the backbone. Figure 2 shows the coalition communication concept for this phase. Four simulated ship platforms from fictional countries: Green (San Francisco, CA, US), Red (Portsmouth, UK), Blue (Dahlgren, VA, US), and Orange (Malvern, UK) communicated over simulated radio links.




Figure 2: Phase II Coalition Structure

No NOC was used although network operations were concentrated in Red and Green. Blue and Orange were the PKI providers for the exercise. The US hosted the following services:

- CA/RA: Netscape Certificate Management System v4.1.5 (NT)
- LDAP Directory: Netscape Directory Server v4.1.5 (NT)
- SSL-compliant Web Server: iPlanet Web Server v1.0 (Solaris)
- Web Clients: Netscape Navigator Clients v4.7.5 (NT/Solaris)
- Mail Server: Netscape Messaging Server v4.1 (Solaris)
- S/MIME-compliant Mail Client: Netscape Communicator Messenger v4.7.5 (NT/Solaris)

The UK hosted the following services:

- CA/RA: Baltimore UniCERT Certificate Management System v3.0.5 (NT)
- LDAP Directory: ISOCOR Directory Server v2.3r1 (LDAP)

- SSL-compliant Web Server: MS Internet Information server v4.0 (NT)
- Role-Based Access Control: WebMACE v1.1 (NT)
- Web Clients: MS Internet Explorer v5.5 (NT)
- Mail Server: MS Exchange v5.5sp2 (NT)
- S/MIME-compliant Mail Client: MS Outlook 98 (NT)
- Firewall/Mail Guard: SWIPSY (Trusted Solaris)

2.3 Testing Conducted and Results from Phase II

We followed a pseudo-military scenario that involved a coalition forming, performing a mission, evolving, and disbanding. From the scenario and the objectives we derived the following our technical PKI events of interest. Each national CA sent the other CA its self-signed certificate for the end users to add to their trusted list. Then US CA issued “coalition” certificates to UK users and vice versa. We tested these certificates by exchanging signed and/or encrypted email and by visiting each other’s secure web sites (via Secure Sockets Layer (SSL) v2.0 using both server-side and client-side authentication). After using the certificates we revoked them and attempted the same tests with the revoked certificates to make sure that revocation ended the trust relationship.

Overall success was achieved in most areas. The most notable deficiencies were caused by incomplete implementation of PKI awareness in the client applications.

2.3.1 Problems Encountered in Phase II

There were numerous bumps along the way and a few failures of minor objectives. This section is a collection of our problems grouped according to the software unit where the problems were manifested.

iPlanet Directory Server—We learned that the Directory Information Tree (DIT) structure is tightly coupled with working of the CA and other PKI servers. We originally underestimated the degree of coupling and could not publish certificates to the directory. We were forced to do several directory naming scheme reworks to make certificate publishing work.

Even after fixing the directory problems, we were unable to publish certificates from Netscape CMS via SSL to the iPlanet directory consistently. Either the directory or the CA seemed very buggy on this point. Once we got it working we dared not touch it. This behavior would not be acceptable in an operational environment.

Netscape Certificate Management System (CMS)—

SSL server-to-server communications never worked for the Netscape Messaging server. Although certificate enrollment for the Messaging Server seemed to work well, the subsequent use of the certificates in SSL communications did not work. This implied no secure transfer of e-mail from one mail server to another, no secure Internet Mail Access Protocol (IMAP), and no secure access to directory data from the directory server. We were however using IPSec to bulk-encrypt all traffic so these issues were not immediate problems.

Netscape CMS seemed to be quite brittle requiring reinstallation numerous times. Simple changes (e.g., IP addresses of servers, etc) could render CMS useless until it was re-installed.

Netscape Communicator clients in general—The inability of client software to reliably check certificate status was a major problem. In Netscape Communicator’s web and e-mail clients, revoked SSL server certificates would not raise any alarm until a CRL was explicitly downloaded into Communicator from Netscape CMS’s end-user web interface. The button used to download a CRL to Communicator apparently is only available when visiting the client web portal of Netscape CMS. Once a CRL was downloaded into the client, revoked web site certificates generated the appropriate warning, and mail users could not use expired certificates for signing messages. All this was expected and proper, but after downloading a CRL, the client must manually reload a new CRL before the old one expires or be unable to use any SSL or S/MIME facilities. This then prevents the user from downloading a new CRL! This behavior is clearly counter-productive. Some flexibility to allow a user to participate in SSL transactions even if the local CRL has expired would be helpful. Another possibility would be to automate the CRL download process. For our application, CRL lifetimes were very short (fifteen minutes) so we were forced to ignore CRLs altogether to avoid the continual annoyance of downloading new CRLs manually.

Netscape did provide a Personal Security Manager (PSM) plug-in for its Communicator 4.73 client. This plug-in would allow the use of Online Certificate Status Protocol (OCSP) to verify certificates presented to the client. However, PSM was so buggy and caused so many crashes that we decided not to use it. Since there were at that time no other freely available OCSP-aware clients we elected not to use OCSP.

Netscape Navigator web client—Users of both national PKIs were able to register for and receive certificates from the web portal of the foreign CA, but US users were inexplicably unable to import the UK’s CA chain into Navigator’s trust list. Numerous creative attempts failed, although the UK was unable to duplicate the incompatibility. The reason for this

problem was never discovered and may have been caused by influences outside either the Baltimore CA or Navigator.

By default, Navigator expects the Distinguished Name (DN) of an SSL server's certificate to follow a specific format. A certificate's DN must have the common name (CN) of the server as its first element, and the CN must match the server's Domain Name System (DNS) name exactly. Using a more human-readable CN (e.g., "CN=Stanleys Web SSL Cert") in the certificate generated name mismatch errors in the browser every time the web site was visited. This makes maintaining a large number of certificates unwieldy because they are not readily identifiable by humans. Supporting the Subject Unique Identifier field or allowing the CN to be free form would help.

The UK certificates generated by Baltimore CA and issued to US users could not be used to sign messages or validate signatures. The problem appeared to stem from the inability of the US's Netscape clients to import the UK's trust chain. Reasons for this inability are unknown.

Role-Based Access Control (RBAC)—Hosts at all sites were able to access the native web interface of NSWC's Netscape CMS CA using SSL with mutual authentication. US users with certificates issued by the UK were able to access UK home-grown websites requiring presentation of a client certificate. But US Netscape users were unable to properly access UK pages controlled by the RBAC software, WebMACE. The reasons for this are not known. The US did not attempt to protect any of its home-grown websites via PKI because it was not immediately apparent how to implement this and testing time was limited.

Firewall and Guards—The UK deployed a coalition guard on the periphery of its national network. The purpose of the guard was to prevent leakage of sensitive information from the national network into the coalition. Unfortunately, the guard did leak e-mail addresses with names that revealed the underlying structure of the UK network (e.g., the domain name indicated which platform the user was located on). Eventually this guard would also be a PKI signature proxy. The guard would replace the signatures of individual UK users with the guard's signature so that the internals of the national PKI would be shielded from the coalition. This feature has not yet been implemented.

General PKI Instability—The US lab at SPAWAR Systems Center—SanDiego, California (SSC-SD) provided Radio Frequency (RF) link simulation for the exercise via AdTech SX-12 RF simulators installed at their site. The RF simulators were intended to provide realistic bandwidth limitations and error characteristics to emulate the HF radio and Satellite communication

links that will be used in at-sea scenarios. Unfortunately, we were unable to simulate RF links in Phase II because the PKI was never stable enough to be stress tested.

2.3.2 Accomplishments of Phase II

Out-of-band resources were established for exchange of administrative data among experimenters. These resources included ftp, web, and chat servers, Voice-over-IP (both in the clear and over IPSec), and teleconference phone calls. The latter two were indispensable in overcoming the PKI and networking obstacles we encountered.

We used an IPSec encryption mesh between each of the four sites using pre-shared keys and 56 bit DES. This allowed us to assure the security of the experiments without relying exclusively on PKI.

We published certificate and user information to US and UK LDAP directories accessible to all. There were no problems with users registering or retrieving certificates, except for the US's problem attempting to import the UK's trust chain. Thus, users at all sites were able to exchange signed and encrypted email using at least US-issued certificates.

The US deployed a Network Time Protocol (NTP) server for eventual use as a trusted time server for non-repudiation. The NTP server was, however, only used to synchronize clocks in order to preserve the correct order of receipt of mail messages from all sites.

2.4 Lessons Learned in Phase II

Many general lessons were learned about the issues of PKI deployment:

- PKI interoperability was, at that time, an afterthought among vendor products we tested.
- PKI-enabled applications were rare and limited in their implementation of PKI features such as certificate status checking.
- PKI was much harder than we thought, and implementations were not at all robust. The brittleness of all the PKI implementations tested meant that they could not be relied upon for operational use at that time. We learned that the foundation of workable PKI is the directory. The format of information stored in national border directories is crucial for all parties to agree upon.
- Constant coordination was required to bring up a coalition PKI.

The state of PKI technology did improve over time as did our understanding of it. We had much more success in the next phase of experimentation.

3.0 Phase III Experiments

3.1 Objectives of Phase III

The goals of VON Phase III were threefold:

1. To centralize trust management at the national level,
2. To reduce risk of PKI component or system failure during the at-sea trial (Phase IV) by defining common minimum architecture requirements and baselining the configuration for the at-sea trial, and
3. To incorporate hardware tokens for end entities' certificate storage and presentation.

Testing was focused on cross-certification, exchange of S/MIME e-mail with attachments, and revocation testing (both end-entity and cross-certificate). Web and other services were de-emphasized in favor of solidifying the PKI itself. As the PKI evolves, we anticipate adding other services.

3.2 Testbed Configuration for Phase III

The testbed for Phase III (shown in Figure 3) simulated five platforms located at four geographically separate sites: two national NOCs, one in the US and the other in the UK, a US gateway ship and two US leaf nodes. The US NOC was physically split between two locations. The LM site provided the PKI servers in its half and NSWC provided DNS and mail servers and served as a network hub. Both US sites hosted LDAP servers for performance, redundancy, and fail-over reasons.

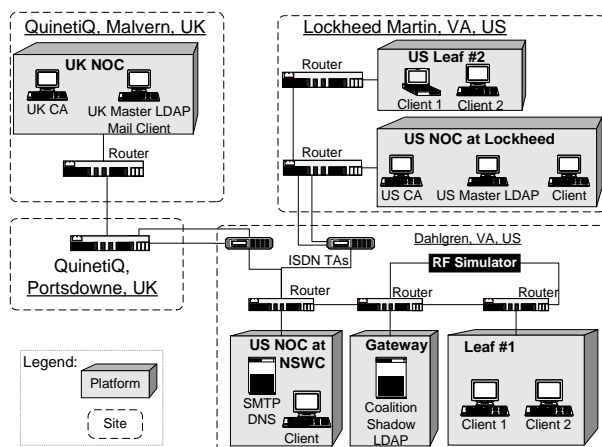


Figure 3: Testbed Configuration for Phase III

The US team developed a proposed coalition PKI architecture document [2] that specified interface standards that PKI products used in the demonstration must support to achieve the minimum acceptable level of interoperability. Only commercial PKI products were used in the demonstration. The proposal was

accepted by the UK with minor changes. In particular, it was agreed that secure and trusted collaboration would be achieved by cross-certification between the US and UK CAs over a single ISDN 64 Kbps channel that emulated throughput expected during the at-sea trial in the following phase.

The configuration below was outlined in the proposal to achieve secure communications and mutual trust between US and UK systems. The boldface items represent changes from the Phase II configuration. The results of the testing confirmed this as the baseline configuration for Phase IV.

The US hosted the following services:

- CA/RA: **Entrust v5.1.1** (NT)
- LDAP Directory: Netscape Directory Server v4.1.5 (NT)
- Mail Server: Netscape Messaging Server v4.1 (Solaris)
- S/MIME-compliant Mail Client: **MS Outlook 2000 (NT) with Entrust Express plug-in**

The UK hosted the following services:

- CA/RA: Baltimore UniCERT Certificate Management System v3.5 (NT)
- LDAP Directory: **Border: iPlanet Directory Server v4.1.5 (NT); Internal: Novell DirXML 1.0 and eDirectory.**
- Mail Server: MS Exchange v5.5sp2 (NT)
- S/MIME-compliant Mail Client: MS Outlook 2000 (NT) **with Baltimore MailSecure.**
- Mail Guard: SWIPSY (Trusted Solaris)

3.2.1 Certification Authorities

Both fielded CA products supported cross-certification as defined in RFC 2587 [3]. To ensure the security of the certificate exchange, an “out-of-band” process (voice telephone) was used to verify the thumbprint of a cross-certificate request.

Scalability problems arise when establishing and maintaining trust relationships solely via cross-certification. A total cross-certification trust model implies a mesh topology with $O(n^2)$ cross-certificates to be issued and maintained. However, we assumed that the number of relationships is manageable given our small demonstration coalition. We chose cross-certification as a potential step toward a bridge CA trust model that would require only $O(n)$ cross-certificates.

To avoid the undesirable side-effects of transitive trust, we specified that the *pathLenConstraint* field of the Basic Constraints extension would be set to zero as described in RFC 2459 [3]. Transitive trust is indirect

trust between PKI domains that can be established either knowingly or inadvertently. For example, suppose CA₁ trusts CA₂ and CA₂ trusts CA₃. If after this CA₁ now trusts CA₃ then transitive trust exists. Transitive trust management via name constraints, etc. was not used.

Risk reduction tests conducted prior to Phase III found that a number of CA configuration options had to be agreed upon in order to ensure client application interoperability. Therefore, the CA products for both countries were required to support the following configuration:

- 160-bit SHA-1 hash for authority and subject key identifiers
- X.509v3 certificates with the following standard extensions:
 - *keyUsage*
 - *authorityKeyIdentifier*
 - *subjectKeyIdentifier*
 - *cRLDistributionPoints*
 - *subjectAltName* (containing the subject's email address per RFC 822), and
 - *basicConstraints*.
- All other extensions marked as non-critical.

The US installed its CA at the LM NOC site and published CA information including CRLs, CDPs, ARLs, and certificates to the collocated US master directory server. Likewise, the UK installed its CA at the UK NOC site and published CA information including CRLs, CDPs, ARLs, and certificates to its master directory server

The US issued two identity certificates to each US users one for encryption and another for signing. Private keys for the signing and encryption certificates were generated on smart cards; but only encryption private keys were escrowed at the CA. The UK issued certificates to its users similarly, except that they used soft tokens and did not escrow any keys.

The UK and US then exchanged copies of their respective Root CA certificate both in native format and in a PKCS #10 signing request via in-band e-mail. Once exchanged, both parties verified the thumbprints of the PKCS #10s over the telephone. These tasks helped us to understand the impact of the following problem-domain issues: the effort involved in using a secure method of exchanging the PKCS#10 requests, the amount of work needed to configure cross-certification, and the time required to set-up a root CA for coalition operations.

3.2.2 Directory Service

The US and the UK agreed to standardize on the iPlanet Directory Server v4.1.5 as the border directory service implementation. The agreement to use a common

directory product avoided several technical and implementation issues, most notably directory replication. Surprisingly, although iPlanet directory server v5.0 was available to us, its replication function is not compatible with version 4.x of the same product. Since the US did not have the resources to test interoperability between Entrust and the v5.0 directory, the UK decided to use the older directory server for its border directory. Directory interoperability is certainly an area where standards are lacking. Emerging standards and products for directory-to-directory interoperability such as LDAP Duplication/Replication/Update Protocols (LDUP), Directory Services Markup Language (DSML) and Novell's DirXML are possible solutions. The UK demonstrated the use of Novell's DirXML internally as an automated directory synchronization agent between iPlanet Directory Server v4.1.5, Microsoft Exchange and Novell eDirectory.

We used centralized-partitioned (a.k.a. hub and spoke directory) topology for our directory replication scheme. Communication between the UK and US directories occurred through the US hub and its UK replica. In a coalition environment where connectivity is sporadic and throughput limited, the hub and spoke topology was best for scalability, redundancy and manageability. Each coalition member provided a read-only directory replica of local security information to the hub directory. The hub directory provided a complete read-only replica to each spoke, thus allowing each coalition member a complete local view of the coalition. Figure 4 depicts an idealized hub and spoke directory topology in a coalition environment.

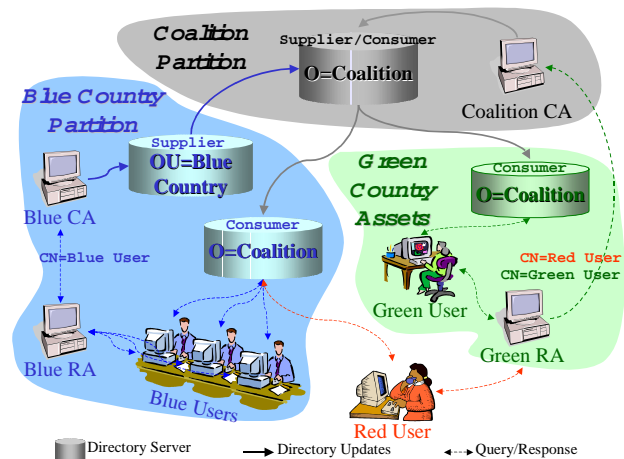


Figure 4: Hub and spoke directory topology

In the figure, Blue country supplies its own master directory information to the coalition and receives back a re-mastered copy of the entire coalition directory (including entries for Green country and the Red user).

This model allows for countries to participate without supplying a master directory or a CA/RA. Replication agreements are minimized while redundancy is preserved. Any country providing a master directory server and a coalition shadow may take over as the coalition hub in case the original hub is damaged or lost. Note that the Coalition CA in the diagram need not exist at all and the coalition directory may be hosted by any partner nation.

Our implementation of hub and spoke topology is shown in Figure 5. Both parties agreed on a directory schema including DIT, added PKI attributes, etc. The US configured two directory servers: one as a US Replication Hub (US-1), one as a US master replica (US-2). Then, the US configured a simulated gateway ship computer (US-3) as a read-only replica of the US Replication Hub (US-1). The US set up replication from US-2 to US-1 (replication path RP1); and from US-1 to US-3 (RP2)




Figure 5 Coalition directory replication topology

The UK also configured two LDAP servers: one as a UK master replica (UK-1), and one as a read-only replica (UK-2) of the US Replication Hub (US-1). The UK collaborated with the US to set up replication from UK-1 to US-1 (RP3). Finally, the US collaborated with the UK to set up replication from US-1 to UK-2 (RP4).

Replication path RP2 demonstrated replication over intermittent links or unreliable connections as may happen between the Gateway ships and NOCs on the shore. Replication paths RP3 and RP4 demonstrated replication over a reliable link, as expected between the two NOCs in the following phase and in deployment.

Replication was achieved using LDAP bind IDs and passwords, rather than certificates for this phase. Replication over SSL will be used in later phases. All replication was server-initiated (push) rather than consumer initiated (pull).

3.2.3 Applications

Secure (S/MIME) email was the touchstone application used to test the Phase III coalition PKI architecture. S/MIME provides authentication and integrity via digital signatures over message hashes, and data confidentiality via encryption. Both the US and UK used Microsoft Outlook 2000 for encoding and decoding of S/MIME messages. We used plug-ins for Microsoft Outlook 2000 to provide trusted exchange of messages leveraging coalition cross-certificates. The US used the Entrust Express plug-in and the UK used Baltimore's MailSecure product for verifying trust between the cross-certified PKI domains. The plug-ins enabled Microsoft Outlook 2000 to check user certificate status by downloading Certificate Revocation Lists (CRLs) from a local directory replica.

3.3 Testing Conducted and Results from Phase III

As detailed above, before beginning testing in this phase we took pains to define minimum interoperability standards. This precaution resulted in a much smoother testing period. We tested by transmitting unsigned, signed, encrypted, and signed-encrypted e-mail messages both with and without attachments during the test phase. Our results demonstrated working path validation and discovery. We also tested revocation by sending signed e-mail between realms after revocation of a user certificate or a cross-certificate.

We used network analyzers to record email and LDAP traffic and verify system correctness. The recorded traffic was analyzed to ensure email messages were indeed digitally signed and/or encrypted when applicable. The recorded traffic was also used to ensure proper workflow for certificate validation. Figure 6 depicts the certificate validation logic the US Entrust Express client used to validate a digitally signed email message from a UK user.




Figure 6: Client-Side Signed Email Validation

3.3.1 Problems Encountered in Phase III

This section presents major problems encountered during Phase III testing. The problems are organized according to the products where they manifested themselves. We have explained each problem to the extent of our forensic abilities, but because of the inherent complexity of PKI, formal attribution of problems is not possible. We hope these records will be useful to the vendors and to new PKI users as they field their own PKIs.

Entrust CA—Insufficient fields were present in the PKCS #10 cross-certification request from the US’s Entrust CA for a correctly formatted cross-certificate to be produced by Baltimore’s UniCERT CA. In particular the Subject and Authority Key Identifier fields appeared to be missing. These fields are essential in correct trust path building. Cross certification was successfully achieved using the US root self-signed certificate instead of the PKCS #10 message. The missing fields were manually added to the cross-certificate by the UK’s CA operators.

Baltimore UniCERT CA—The UK found it difficult to achieve reinstallation of Baltimore’s UniCERT CA

without reinstalling the machine’s entire operating system. It is very important to establish correct CA configuration throughout the coalition at install time.

A few other minor incompatibilities were also discovered between MailSecure and UniCERT in trust path building using cross-certificates.

Entrust RA – After revoking a user through Entrust RA, the CRL must be manually created and pushed to the directory via the Entrust RA interface in order for the latest CRL to be immediately published to the directory. Once again, a publish-and-subscribe CRL mechanism would be ideal.

Entrust Express Outlook 2000 plug-in – Outlook 2000 must be installed in “Corporate Mode” in order to support Entrust Express. Installing in “Internet Mode” produced inconsistent results and strange errors when doing signature validation.

When trying to add a user to the Entrust Address Book from a Directory Search, Entrust Express generated an errorⁱⁱ. Entrust assumes that the certificate being added to the Entrust Address Book from the LDAP Directory is an encryption certificate (e.g. the *keyUsage* value is “Key Encipherment”). Entrust does not publish digital signature certificates to the directory because they are sent in every S/MIME of digitally signed message. If the *userCertificate* attribute for a user in the directory contains multiple certificates, the first or only certificate must be the user’s encryption certificate. For Entrust Express, the ideal would be for each user entry of the directory to contain only one certificate: the users’ current encryption certificate. To avoid problems, any revoked certificates must be manually remove from the directory and the first certificate entry must be a valid encryption certificate.

Entrust Express was unable to validate the certificate chains with heterogeneous signature algorithms. VON’s policy specified DSA key pairs, since DSA was the preferred US and UK Government algorithm. When the RSA algorithm became public VON’s requirement changed to using RSA key pairs since RSA has wider usage. The UK had installed its Baltimore CA using a DSA self-signed certificate prior to the policy change and preferred not to reinstall the CA in order to comply. Instead, the UK team decided to issue all end-entity certificates with RSA key pairs and leave the self-signed root certificate alone. Unfortunately, we found during testing that Entrust Express displayed an errorⁱⁱⁱ when opening digitally signed messages received from the UK since the sender’s CA certificate public key algorithm was different from the public key algorithm used by end-entity certificates. The work-around was to ensure the same public key algorithm is used for CA and end-entity certificates. To fix this problem during the testing events, UK had to reinstall its entire CA to change the CA’s self-signed certificate to use the RSA

algorithm. All UK user certificates were then issued with the RSA public key algorithm. In general we determined that the Entrust plug-in could handle homogeneous RSA or DSA algorithms all the way up the chain, but cannot validate certificates whose validation paths use mixtures of DSA and RSA signing algorithms.

iPlanet Directory Server – Occasionally, replication agreements did not result in automatic replications when the directory service in question functioned as both a supplier and a consumer of the same tree (e.g., coalition mirror directories that also replicated themselves to other directories).

iPlanet Messaging Server – The Messaging Server must be able to write to the directory root organization (e.g. “o=coalition.mil”) where it pulls email-related information. Otherwise the Messaging Server will fail to start Simple Mail Transfer Protocol (SMTP) services. The Messaging server uses the root entry to store certain administrative data. If the root entry is not writeable, the SMTP service cannot start, but other services may. The US had to constrain directory replication to its Messaging Server to the “ou=United States, o=coalition.mil” subtree to work around this limitation.

Entrust & Baltimore Mail Client Plug-ins—By default, Entrust and Baltimore cache Certificate Revocation Lists (CRLs) and Authority Revocation Lists (ARLs). It was therefore necessary to restart the clients to download the latest CRL from the directory when conducting revocation tests. This is not a shortcoming; both retrieve CRLs from the directory when the most recent CRL expires. Unfortunately, we found no way to push an interim CRL containing newly revoked certificates to the clients before the next update time. Turning caching off produced excessive CRL network traffic, and caching time could not be set below four hours for Entrust because that is the minimal CRL lifetime allowed in the version of Entrust CA we were using. Our requirement for timely revocation drove this testing, and no suitable alternative could be found. OCSP was not supported by either client, and even with OCSP, our requirement to tolerate intermittent network connectivity would have limited OCSP’s utility. The most satisfactory arrangement would be if there were some way to set up a CRL publish-and-subscribe mechanism where CRLs could be pushed asynchronously to clients.

Problems were encountered when sending e-mail messages between Entrust Express software and Baltimore’s MailSecure software. Entrust Express includes the entire certificate chain with each signed message. MailSecure used the chain included in the message to perform validation instead of consulting the directory. Therefore all Entrust Express-signed

messages failed to validate in MailSecure because the US-signed-by-UK cross-certificate found in the directory was never seen. Since the UK’s trust of the US was documented in the cross-certificate, the US root self-signed certificate was not trusted directly. Individual user certificates could be validated after opening the messages by manually resolving trust paths back to the cross-certificate. Since it would be impossible for Entrust Express to include the correct validation chain for a UK user, a straight-forward solution would be to no longer include the validation chain in messages at all. Unfortunately, Entrust Express did not provide such a facility. Inclusion of a proper validation chain would help satisfy the intermittent network connectivity requirement, but the amount of additional data sent with each message could pose a bandwidth problem under the strain of operational use.

A number of attributes needed to be added to the UK’s directory entries that were mandated by Entrust: First Name, Last Name, Common Name, User ID, Password, *mailrecipient*, *nsmessagingserveruser*, *mailbox*, *Maildeliver*, *Mailhost* to correctly process them. These were not strictly needed by the UK, but were added for compatibility reasons.

Baltimore MailSecure—MailSecure did not recognize the cross-certificates we used to establish trust because it did not use the *crossCertificatePair* attribute of the US CA’s directory entry. As a workaround, the UK obtained the US’s cross-certificate signed by the UK (labeled <<US signed by UK>> in Figure 7) and copied it into the *cACertificate* attribute of the US CA’s directory entry. They did this in a “stub” directory copied from the real directory so as not to modify the original. They then pointed MailSecure to the stub directory as the first source for certificate path validation. When a certificate’s trust chain led MailSecure to the US CA’s certificate in the stub directory, the *cACertificate* attribute further referenced the UK’s own CA as a superior in the trust chain. We believe this work-around does not impact the trust hierarchy. However, if the same modifications were made in the master (US) directory, all PKI enabled applications under the US’s CA that use the *cACertificate* attribute would work incorrectly. Therefore the stub directories are a necessary part of the approach. Fortunately, MailSecure does allow the use of multiple directories to build validation paths. Without this capability the UK users would have had to copy their entire directory into the stub directory to make the process work.

Figure 7 shows how MailSecure searches the stub directory first to find certificates. When it needs to find the US CA certificate, it finds the appropriate entry and looks at the *cACertificate* attribute. The first value in the attribute is the <<US signed by UK>> certificate

that points to the UK CA. This feature allows MailSecure to automatically trust all US-issued certificates. The self-signed certificate remains as the second value of the attribute for compatibility purposes. However, we have found, in general, that PKI path-building clients do not look beyond the first value of an attribute.

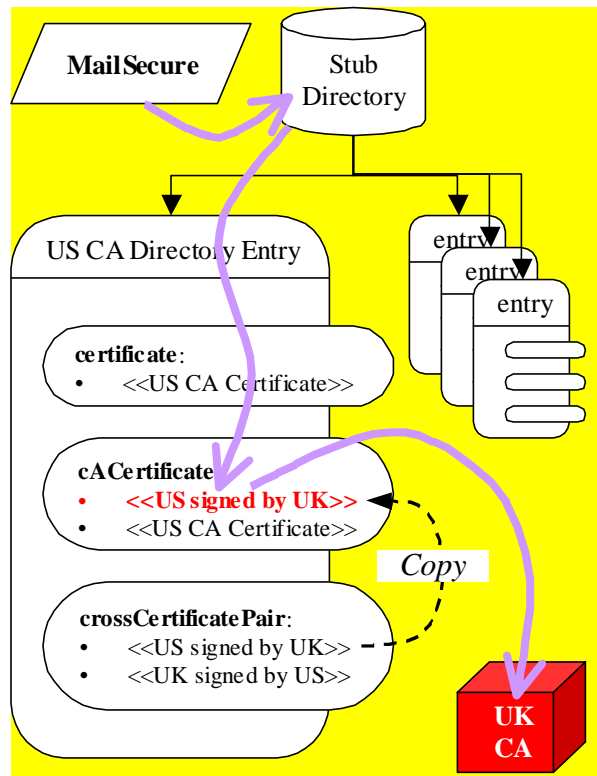


Figure 7: Using a Stub Directory with MailSecure

General PKI Problems—Some PKI products expect a country (C=) code to be the root element of all coalition DNs (after the fashion of X.500). Since VON uses the Organization (O=) code as the root, we encountered several PKI problems. For example, with no country code in the DN, the UK users were unable to generate their own keys and request certification via MailSecure. However, this was achieved at the local Registration Authority (RA) using face-to-face certification resulting in the manual transfer of user certificates to client machines. This DN restriction also means that each user may need a set of certificates for the coalition and another set for national use. The practicality of this must be considered.

RFC 2459 [3] is ambiguous in its specification of CRL Distribution Points (CDPs). Although all PKI products we used follow the standard, different legitimate interpretations resulted in incompatibilities between compliant products. We discovered that the UK's Baltimore UniCERT CDPs could be configured in a way that made them incompatible with Entrust Express,

although both appeared to be following the standard. The ambiguity allowed directory locations to be resolved from UK and US certificates in incompatible ways. In order to resolve this incompatibility, we found the Issuing Distribution Point (IDP) must be set to non-critical and fully qualified CDPs must be used.

We found it necessary to use third-party utilities to confirm the correct configuration of certain pieces of software used in the trials. For example, certificate viewers and Base-64 decoding tools from the OpenSSL distribution were needed to debug problems with certificates issued by foreign CAs. We suggest that vendors include such tools in debugging suites to increase the interoperability of their software with others.

3.3.2 Accomplishments of Phase III

All participants accomplished the following during the summer 2001 test period:

1. Established network infrastructure over a private ISDN link.
 - Simulated platforms included national NOCs and several simulated ship platforms.
 - Infrastructure included both nations providing coalition e-mail and DNS servers.
2. Established nationally supplied directory services interconnected into a unified coalition directory with automatic replication between sites.
3. Set up national PKIs and cross-certified them yielding a unified coalition PKI including:
 - Directory Servers
 - Certification Authorities (CAs)
 - Registration Authorities (RAs)
 - PKI enabled e-mail clients
4. Verified the functionality of the coalition PKI via e-mail tests.
 - Conducted 48 e-mail tests (including digitally signed and/or encrypted e-mail both with and without attachments) with no unqualified failures.
 - Discovery of encryption certificates via unified coalition LDAP directory worked consistently.
5. Tested revocation of individual coalition users and cross-certificates.

3.4 Lessons Learned in Phase III

The Phase III testing identified a number of issues with the vendor products used. While all 48 email-exchange tests were successfully performed, a few of the exchanges required workarounds deemed unsuitable for

a tactical environment. These workarounds were due to PKI vendor incompatibilities. In addition a number of issues were discovered concerning the underlying network infrastructure (e.g. DNS, routing, etc), which must be resolved prior to at-sea trials. The teams will perform additional work in 2002 to get the demonstration testbed ready for at-sea trials in 2002.

Following are some logistical lessons we learned during the testing process:

- The conference telephone call was an invaluable tool that allowed problems to be solved in an efficient and timely manner. It also allowed out-of-band verification of certificate fingerprints during the cross certification process. We found using an out-of-band channel for verifying certificates and PKCS #10s to be simpler and more cost-effective than face-to-face certificate exchange.
- Detailed configuration planning in advance avoids unnecessary, lengthy reinstallations of software.
- Separating key server machines among several sites makes it more difficult to locate and rectify network configuration and other problems.
- The US found it useful to have several administrative user accounts for each nation: echo, record, and revocable. The echo user is configured so that e-mail to this user is automatically echoed to the sender. This account is useful in testing basic e-mail connectivity so that one nation can verify that another's e-mail server is responding without further coordination or specialized knowledge. The record user was used as a repository for CCs of all mail messages sent during the testing. This user's mailbox formed a complete record of all e-mail sent during the test and often served as verification that a nation actually sent a message when network congestion caused delayed delivery to the recipient. The revocable user accounts are useful for conducting revocation testing. These user's certificates are intended to be revoked for testing purposes so that other users' accounts need not be disturbed and no one's feelings get hurt!

Following are some lessons we learned about planning and managing LDAP directory servers for PKI:

- Hub and spoke replication topology worked well, allowing access to the complete coalition directory even when remote links were down. Further experiments may be needed to check that this strategy will work with high volumes of data and/or low bandwidth links.
- The e-mail address book is often separate and disconnected from the coalition directory because the directories are used for different purposes. Manually copying e-mail information into the

coalition directory is a slow and error-prone method. Automatic replication between the e-mail and the coalition directories is highly desirable. . The UK successfully demonstrated Novell's dirXML product for this purpose in their testbed.

- Each nation needs to ensure that its users' entries are fully completed in the directory so that the PKI-enabled client software in use for other nations can process all users' certificates.
- The directory must be a robust product. Restarting the directory and rebooting the directory server regularly will not be satisfactory in real-time operations.

4.0 Conclusions and Future Work

The Phase IV at-sea trial will exercise the PKI configuration established and refined in earlier phases. Work is ongoing now to refine the configuration in preparation for the testing event. Several more e-mail exchange tests have been conducted, and the testing methodology has been refined to a high degree of precision. Since the test will be shipboard, a great deal of logistical matters must be considered. It normally takes over a year to determine the ships where an installation will be done, schedule a time for the ship to be in port, find a place for the installation, and verify that the installation works without negatively impacting any mission-critical systems. At this time, the logistics dominate the preparation process and the exact venue is still uncertain. This phase may involve more nations and will involve untrained users for the first time. We are prepared to collect data on both the functionality and the usability of our design from a user perspective.

In Phase V, we will seek to overcome the problems of PKI by using Extensible Markup Language (XML) and its child technologies. XML is quickly becoming the de facto standard for providing interoperability between disparate systems. XML's meteoric rise together with the momentum of Web Services may finally push PKI to deliver on its promise of universally defined trust and usability. In particular, XML standards that may be leveraged to make PKI easier to use and implement include XML Digital Signatures, XML Encryption, XML Key Management System (XKMS) and Security Assertion Markup Language (SAML). These standards may help solve problems inherent to the design of a Coalition PKI. For example, providing PKI services for nations that do not have a pre-existing PKI or the technology to establish one. With a standard web-based interface to the coalition PKI, the coalition would be able to meet nations at their level of technology and, with minimal provision, make it accessible. The coalition PKI should have a common interface that is usable in the same way by all partners regardless of the

underlying PKI provider. PKI should be a transparent part of the network infrastructure and should be usable over low-bandwidth links and on low-end workstations or mobile devices. It should allow for considerable mobility by low-end clients and be easy to set up and tear down dynamically as coalition partners come and go. Different coalition members need different access to the coalition PKI for the various roles they may play. Particularly useful is the offloading of CPU intensive PKI processes from the client to the server and making developers job of integrating PKI into applications easier. As a result, thin clients can take advantage of the strong security a full-fledged PKI provides. Using XML as a fundamental technology for PKI may allow machines to communicate in a language they already understand without a complex rollout of customized hardware and software. The issues of the online nature of these follow-on technologies will be a subject of considerable concern in this phase. We plan to contribute to the development of the standards to the benefit of all those who cannot depend on continual availability of the internet or high-bandwidth connections.

Beyond managing a single coalition, one of VON's future aims is to manage interactions among multiple, simultaneous coalitions. Each coalition must be treated as a separate "community of interest" with administrative and policy structures that are somewhat independent from those of the member nations. Additionally, there are usually multiple security levels and compartments within each community. Given n nations the potential number of communities is bounded by the expression, $2^n - 1$. The number of security levels and compartments is completely arbitrary and may be as complex as the coalition administration finds useful. The picture is further complicated when one considers the existence of informal ties and covert channels between nations. The rules for controlled interchange among such communities are necessarily complex and should be enabled/enforced by a coalition PKI. This very difficult problem may not be addressable by any technological solution at all, but the goal of the VON project is to identify and implement technology that will enable at least a partial solution to problems of this sort.

In conclusion, we observe that military coalitions are often formed between partners with complex political relationships and data sharing requirements. These requirements must be underpinned by technologies that support individual identification, encryption of content for privacy purposes, data separation and access control, and non-repudiation. These will all be essential services for future network-enabled warfare operations between military allies. PKI has been shown to provide the technical underpinning for such services, and is

likely to be an important part of future coalition operations. The technologies have been demonstrated practically, and are found to be reaching the state of maturity where they can be used for such purposes. Nevertheless, there are some areas where further work is required if the military is to reap maximum benefit from this young technology. In particular, policies on the use of PKI must be refined, the robustness of the technology must be determined under a variety of circumstances, and network operators must be trained in its use if it is to support coalitions of the future.

References

- [1] NTA Coalition Information Technology Interoperability Final Report, 5 January 2001
- [2] "Proposed AUSCANNZUKUS+ Coalition PKI Architecture for VON 2002 Pacific Demonstration," 26 June 2001, by NCAT/Lockheed Martin M&DS and Glenn Fink, NSWC Dahlgren
- [3] RFC 2587 - Internet X.509 Public Key Infrastructure LDAPv2 Schema
- [4] RFC 2459 - Internet X.509 Public Key Infrastructure Certificate and CRL Profile

Biographical Data:

Glenn Fink (MS, Comp Sci) is a member of the Information Transfer Technology group at the Naval Surface Warfare Center (NSWC) in Dahlgren, Virginia. He specializes in computer and network security, especially PKI and Intrusion Detection. He has worked for the DoD for 14 years during which time he has been associated with a variety of projects mostly involving software development. He plans to leave the government in Fall 2002 and pursue a doctoral degree in Computer Science at Virginia Tech. Apart from work, his primary interest is in his family: his beautiful and intelligent wife and two sweet young children. He is the "principal" and occasional teacher for his children's home-school. His family participates in a house-based church fellowship.

Shawn Raiszadeh has worked for Lockheed Martin designing and developing cutting-edge security systems for the past three years. Shawn has worked on a number of research and development programs with the goal of creating new and innovative solutions to existing problems. Shawn has a Bachelor of Science degree in computer science from Virginia Tech.

Tim Dean, (BSc, MSc, MBCS) leads a research team and is a technical specialist in IT Security. His particular interest is in Public Key Infrastructures (PKI) and the issues associated with their practical deployment. He worked for the UK Ministry Of Defence for 14 years during which time he led teams in a variety of defence-related messaging and security projects. These included the design of new security protocols and architectures, including a key management scheme for a NATO communications network. For the last five years he has headed a research team studying network vulnerabilities and countermeasures in a military context. He now works for QinetiQ, where he is continuing his research interests. In his leisure time, he enjoys playing the piano and violin, which he uses as part of the worship group at his local Baptist Church.

i VON is a joint Office of Naval Research (ONR, US) and Defence Science and Technology Laboratory (DSTL, UK) In the US, ONR subcontracted its VON work to Naval Surface Warfare Center in Dahlgren, Virginia (NSWC, US); SPAWAR Systems Center in San

Diego, California (SSC-SD, US); and Lockheed Martin, Management and Data Systems, Integrated Solutions Center, eSecurity Center of Excellence in Fairfax, Virginia (LM, US). On the UK side, DSTL (formerly Defence Evaluation and Research Agency (DERA)) subcontracted work to the government-owned private company, QinetiQ Limited (QinetiQ, UK). Of these entities, our bilateral PKI research team was composed of members of QinetiQ, LM, and NSWC. To reduce confusion bred of multiple acronyms QinetiQ will be referred to as the UK team and LM and NSWC jointly will be referred to as the US team where their separate accomplishments are not significant to the context.

ⁱⁱ Entrust Express error: “(-3975) A certificate attribute for this EntrustName does not exist.”

ⁱⁱⁱ Entrust Express error:“(-4089) Signature algorithm cannot be used with given key.”