Federal PKI Directory Concept of Operations

20 April 1999

CYGNACOM SOLUTIONS

Suite 100 West♦7927 Jones Branch Drive♦McLean, VA 22102-3305♦703 848-0883♦Fax 703 848-0960

TABLE OF CONTENTS

SECTION **PAGE**

TABLE OF FIGURES

1 INTRODUCTION

The Federal Public Key Infrastructure (FPKI) is intended to support security services for communication between the public and government employees and between government employees associated with different agencies and organizations. A Federal Bridge Certification Authority (BCA) has been proposed to cross-certify agency and organizational principal certification authorities (PCAs), providing the necessary mapping information to support the verification of certificates between differing trust domains.

1.1 PURPOSE AND SCOPE

This paper describes the architecture for a proposed Federal PKI repository, composed of a collection of interconnected directory servers. The paper addresses the interconnectivity of a Federal BCA directory server with a number of border directory servers providing information on behalf of interconnected trust domains. The paper provides a proposed concept of operation, examines protection issues, and describes a strategy for the evolution of the Federal PKI Directory.

1.2 BACKGROUND

The Federal Public Key Infrastructure (PKI) will be constructed using PKI components from numerous trust domains. Some of these will be initially limited to particular applications (e.g., S/MIME signature verification), but some will be multi-application PKIs. Because of the nature of its composition, the Federal PKI cannot be viewed either as a monolithic structure or a single enterprise PKI.

The approach adopted for the Federal PKI is based on the concept of a “bridge CA.” The bridge CA provides trust (or certification) paths between PCAs for the trust domain PKIs. Industry organizations and other nations are adopting similar solutions, where a designated CA cross-certifies with high level CAs in different trust domains, to create certification paths. This approach, illustrated in Figure 1, allows large-scale government, industry, national or global PKIs to be assembled from application or enterprise scale PKIs. The approach is further described in [TWG-98-29] and the Federal PKI Concept of Operations [CONOPS].

The Federal Bridge CA will provide cross-certification among the trust domain PKIs. Each trust domain will designate a single principal CA to cross-certify with the Federal Bridge CA. The combination of a principal CA and its associated PKI form a domain of trust, wherein the principal CA provides a known point of trust for the domain. Trust domains currently use certification authority (CA) products and services from different vendors, and client products from many vendors.

To verify a digital signature, a client must build a certification path from the originator's certificate back to the certificate for an authority that the client trusts. For the bridge CA approach to work, directory servers must be able to retrieve the certificates and certificate revocation information from trust domains that are cross certified via the Federal Bridge CA.

The Bridge CA is cross-certified with the principal CAs for the client's trust domain and the originator's trust domain. This enables the client's directory agent to build a certification path from the originator's certificate back to the client's known point of trust, allowing the client to verify the originator's signature.

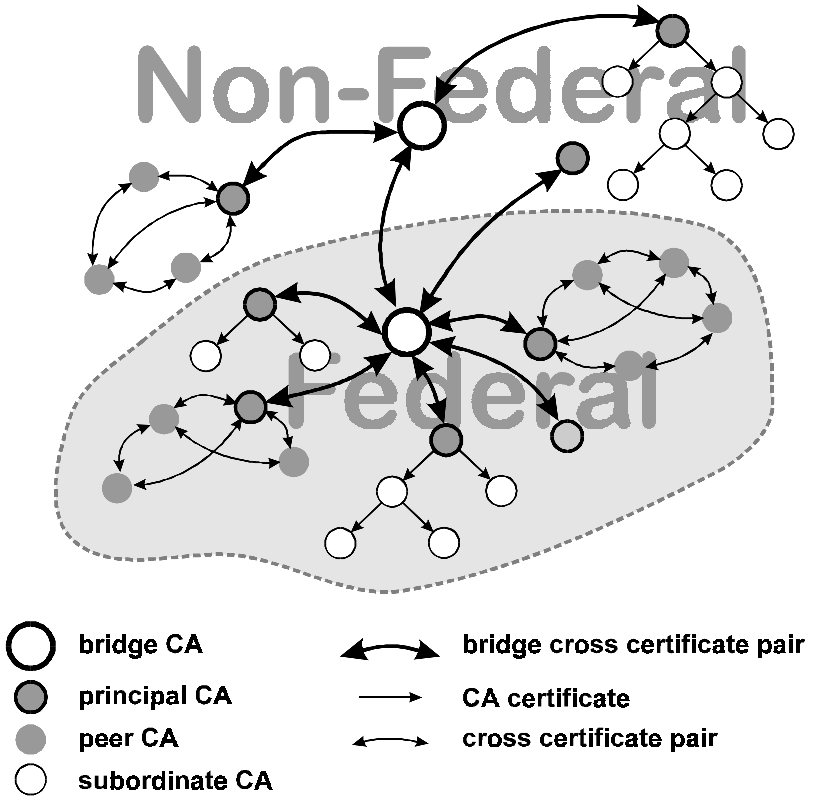

Figure 1. FPKI CERTIFICATION PATH ARCHITECTURE

The Federal BCA facilitates cross certification among trust domains within the Federal PKI. Access to the relevant certificates, certificate revocation information, certificate policies, and certification practice statements must be provided by publishing this information in one or more repositories to support access by users in other trust domains. The Federal PKI repository proposed in this paper includes a Federal Bridge directory server that is interconnected with a set of border directory servers.

2 FEDERAL PKI DIRECTORY ARCHITECTURE

The Federal PKI directory will be a distributed directory system comprising a bridge CA directory server and a number of other directory servers. Each trust domain will provide directory information via one or more of these directory servers. The directory servers are intended to provide a mechanism for clients within Federal Government trust domains to retrieve certificates and certificate revocation information from other trust domains without requiring the client user to deal with multiple directory access protocols.

Generally, a border directory server is a directory server that has been designated to provide the primary public directory system interface for a trust domain. By providing a separate border directory server, an organization can retain its existing directory infrastructure and still be able to communicate within the Federal PKI.

A separate border directory server would be a new directory server specifically created to interface with the Federal PKI. These border directory servers would implement an "external" interface to other Federal PKI directory servers and, optionally, an "internal" interface to the organization's internal directory infrastructure. One of the key reasons to implement a separate border directory server is to allow the trust domain to separate the information stored in its infrastructure into information that can be provided to the public and information that cannot be provided to the public.

2.1 FEDERAL DIRECTORY COMPONENTS

This section includes descriptions of the entities involved with the Federal PKI . The definitions and terms used are consistent with those defined in the Proposed Federal PKI Concept of Operation [Burr].

- *Trust Domains*: In the Federal context a trust domain is a portion of the Federal PKI that operates under the management of a single *policy management authority*. One or more Certification Authorities exist within the trust domain. Each trust domain has a single *principal CA*, but may have many other CAs. Each trust domain has a domain repository. [Note: By fielding a border directory server, a trust domain can provide public access to all or only a part of its repository.] In the non-Federal Context, trust domains may be more loosely organized, but consist at a minimum of a group of CAs that share trust and operate under consistent policies.

- *Federal Policy Management Authority (FPMA)*: this management authority sets the overall policies of the Federal PKI, and determines how trust domain policies map to federal policies in providing cross-certification. It operates a Federal Bridge CA and a repository (i.e., the federal bridge CA directory server).

- *Domain Policy Management Authorities (DPMA)*: a policy management authority approves the certification practice statements of the CAs within a trust domain, and monitors their operation. The DPMAs operate or supervise a domain repository. In the non-federal context, a DPMA may be an association of CAs that share trust and use consistent or comparable CA policies.

- *Certification Authorities (CA)*:

  ◊ *Bridge CA (BCA)*: the Federal Policy Management Authority operates the Federal Bridge CA, the purpose of which is to be a bridge of trust that provide trust paths between the various trust domains of the Federal PKI, as well as between the Federal PKI and non-federal trust domains. FPMA-approved trust domains designate a principal CA that is eligible to cross-certify with the Federal BCA. Note that the BCA is not a *root CA*, since it does not begin certification paths. When the BCA cross certifies with CAs it may limit the propagation of trust to other, cross-certified domains using certificate extensions such as: nameConstraints, basicConstraints or policyConstraints

  ◊ *Principal CA*: A CA within a trust domain that cross-certifies with the Federal BCA. Each trust domain has one principal CA. In the case of a domain with hierarchical certification paths, the Principal CA is the root CA of the domain. In a mesh-organized domain, the principal CA may be any CA in the domain. However it will normally be one operated by, or associated with, the DPMA.

- *Directory servers*: Directory servers are on-line repositories that provide certificates and certificate status information. In the Federal PKI, directory servers will provide information via the LDAP protocol or X.500 DSP chained operations, but they may also provide information in other ways. The FPMA will maintain an open BCA directory server for CA certificates and revocation information. Border directory servers contain certificates and CRLs for CAs and end-entities in their domain.

- *BCA Directory Server*: The BCA directory server will be open to Internet access by anyone, and will make the following available:

  – All certificates issued to or by the BCA;
  – All cross certificate pairs containing certificates issued to or by the BCA;
  – All CA and end-entity certificates required for interoperability within the overall Federal PKI;
  – The Federal BCA CRL.
  – Other certification information, as determined by the FPMA;

2.2 ARCHITECTURAL OVERVIEW

Figure 2 illustrates the bridge CA with its associated directory server. The Federal Bridge CA serves as a cross-certification entity for the principal CA associated with each trust domain. Each component PKI within the Federal PKI will be represented by a minimum of one border directory server, which may be provided by the trust domain itself, provided by another trust domain, or outsourced to an external service provider.

Border directory servers will connect via the bridge CA directory server to provide a government-wide certificate management repository. The border directory server will provide each trust domain with a publicly visible repository for certificates, certification revocation information, and certification practice statements. The border directory server need not replicate a trust domain's entire internal directory information base (DIB).

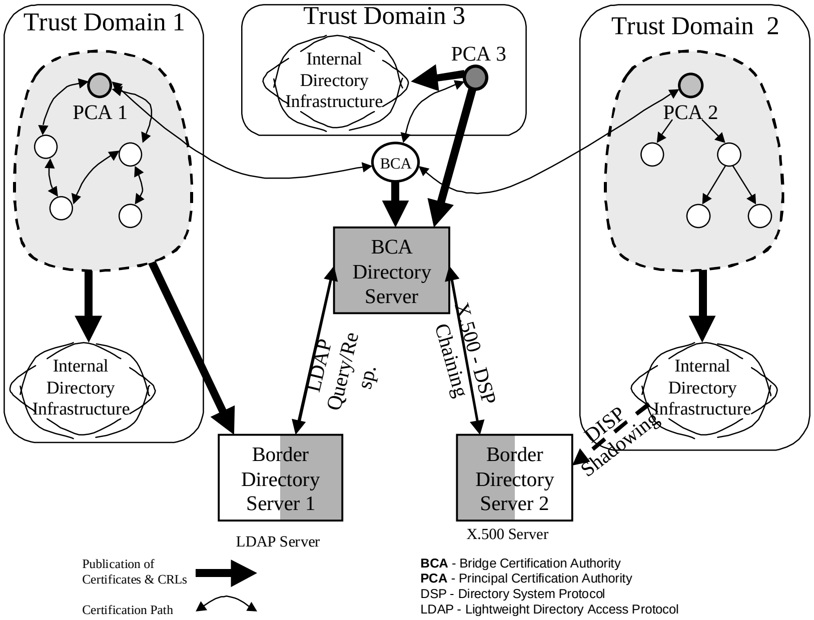
Trust Domain 1
Trust Domain 3
Trust Domain 2
Internal Directory Infrastructure
PCA 3
PCA 1
PCA 2
BCA
BCA Directory Server
LDAP Query/Resp.
X.500 - DSP Chaining
Internal Directory Infrastructure
Internal Directory Infrastructure
DISP Shadowing
Border Directory Server 1
Border Directory Server 2
LDAP Server
X.500 Server
Publication of Certificates & CRLs
Certification Path
BCA - Bridge Certification Authority
PCA - Principal Certification Authority
DSP - Directory System Protocol
LDAP - Lightweight Directory Access Protocol

Figure 2. Federal Border Directory Architecture

Figure 2 illustrates the interconnection of three trust domains, each having different internal directory structures, with the BCA directory server. (Note: External connections, such as the Internet, are not included in the figure.)

> Trust Domain 1 publishes certificate information to border directory server 1 through any protocol it chooses. The border directory server is provided by the trust domain and supports the Lightweight Directory Access Protocol (LDAP) for queries and responses from other trust domains.
>
> Trust Domain 2 publishes information from its internal directory to its border directory server using the X.500 Directory Information Shadowing Protocol (DISP). The border directory server supports the X.500 Directory System Protocol (DSP), using chaining to support queries and responses from other trust domains.

Trust Domain 3 does not have a separate border directory server. Instead, the PCA is responsible for posting certificate information directly to the BCA directory server.

2.3 CONCEPT OF OPERATION

User access to FPKI directory servers (i.e., the BCA directory server and border directory servers) should be limited to permit only interrogation services (i.e., read, compare, list, search, and abandon). We recommend that modification and administrative services be restricted to administrative users only. In this section, the term "access" implies interrogation services only.

The primary use of the FPKI directory is expected to be to provide certificates for relying parties in different trust domains to support digital signature validation. Figure 3 illustrates a typical scenario, wherein a relying party (the directory user) requests the signing certificate of a remote user from whom a signed message has been received:

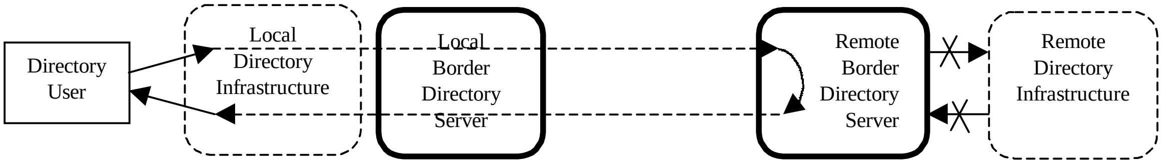
Directory User
Local Directory Infrastructure
Local Border Directory Server
Remote Border Directory Server
Remote Directory Infrastructure

Figure 3. Example FPKI Directory Usage Scenario

In this scenario, the user submits a read request for the sender's directory entry (which contains the sender's public digital signature key certificate) to a local directory server. The local server determines that the requested directory entry is not present and forwards the request to its local border directory server. (This is referred to as an "outgoing request.") Since the local border (in this case) includes only a subset of the information provided by the local directory infrastructure, the border directory server also determines that the requested entry is not present and forwards the request to a remote border directory server, possibly via the BCA directory server.

The remote border directory server receives this "incoming request" and retrieves the requested directory entry, sending it back to the relying party (i.e., the original directory user) as a response. Note that the remote border directory server need not retrieve information via its connection its local directory infrastructure, as indicated by the "x" on those information flows. This model allows trust domains to limit the risk of unauthorized disclosure or modification of information within their domain infrastructures.

Incoming directory requests can be originated by trust domain users, authorized external users, or members of the public. Each trust domain must evaluate the information provided within its directory infrastructure to determine the risk associated with access by these user categories. The remainder of this section presents four example directory configurations, presented in order of increased trust domain protection, and discusses some of the protection issues relating to them.

2.3.1 Free Access

In this example, illustrated in Figure 4, a separate border directory server is not provided. Instead, the internal directory infrastructure is directly connected to the "outside world." The directory infrastructure is implemented by one or more directory products that do not require user authentication for directory operations. Furthermore, no distinction is made between information that is available to domain users and information that is available to the public. This approach allows incoming requests from any originator, including members of the public and provides no confidentiality over any of the information stored in the trusted domain's internal infrastructure.

The free access approach could be viable in some instances where none of the information maintained in the infrastructure is of a sensitive nature. An example of where this approach might apply would be a professional organization, such as ACM or IEEE, where the directory might be maintained solely to provide public access to directory information (e.g., email address, public signature key certificate) provided by each member.

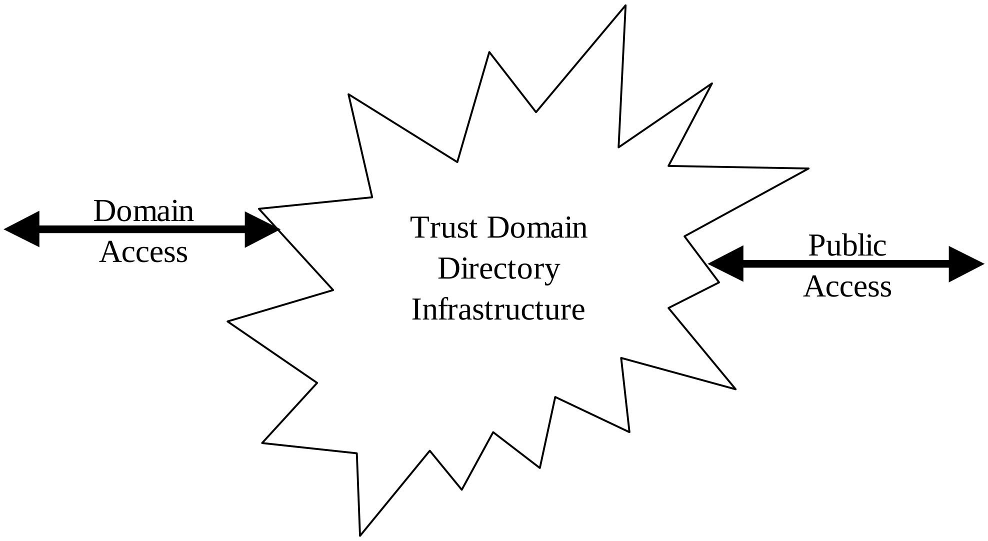

Figure 4. Free Access

2.3.2 Commercially Restricted Access

In this example, illustrated in Figure 5, the internal directory infrastructure is directly connected to the "outside world" as in the preceding example, but the trust domain's directory infrastructure is implemented using directory servers that implement commercial-quality protection mechanisms providing such features as end user authentication and identity-based access controls. The restricted access server provides the capability to control directory requests based on user identification or class. This provides a moderate level of protection against compromise for protected information stored in the trusted domain's internal infrastructure. The level of assurance provided by such a system is an important criterion that should be based on independent assessment against accepted security criteria. The validated level of protection should be used to determine whether the system provides adequate protection for the sensitive information.

This type of configuration can be used to provide a more granular approach to infrastructure protection. That is, data within the infrastructure could be divided into different "buckets" based on how access to the data should be controlled. Some data could be provided directly to the public, while other data could require authentication to limit dissemination to domain users only prior to permitting access.

An example of where this approach might apply would be a company where the directory might be maintained to provide public access to directory information (e.g., email address, public signature key certificate) for employees in public relations or sales, but restricted access to directory information for other employees.

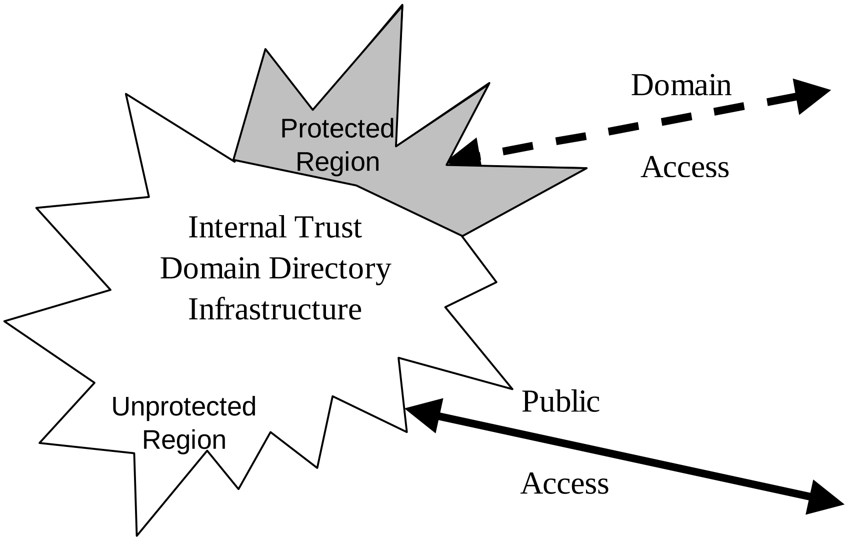

Figure 5. Commercially Restricted Access

2.3.3 Free Access Border Directory Server

In this example, illustrated in Figure 6, the trust domain's directory infrastructure is maintained within a protection boundary. Depending upon the its protection requirements, the trust domain's protection boundary may be enforced by one or more independently validated security mechanisms (e.g., packet-filtering firewalls, application-level firewalls, message guards, or directory guards). A separate border directory server is installed on the unprotected side of a trust domain's protected infrastructure. The free access border directory server allows incoming requests from any originator, including members of the public. This approach provides no confidentiality over the information stored in the border directory server, but provides protection over the information stored in the trusted domain's internal directory infrastructure to the degree provided by the protection domain mechanisms.

In this case, administrative users within the trust domain are responsible for populating information out to the border directory. This allows filtering of the information to limit the sensitive information made available via the border directory server. As with the first example, all information within the border directory server is available to the public. However, domain users may access the trust domain infrastructure via the protection boundary mechanisms. Such domain user access would not involve the border directory server.

This type of architecture can provide additional protection over information maintained within the trusted domain, since specific administrative activities are required to move information out to the border directory server. The information within the domain is protected using the protection boundary mechanisms (e.g., packet-filtering firewalls, application-level firewalls, message guards, or directory guards).

An example of where this approach might apply would be a governmental agency that needs to provide unrestricted public access to directory information for a subset of employees or office codes, but views such information for all other employees and office codes as confidential within the organization.

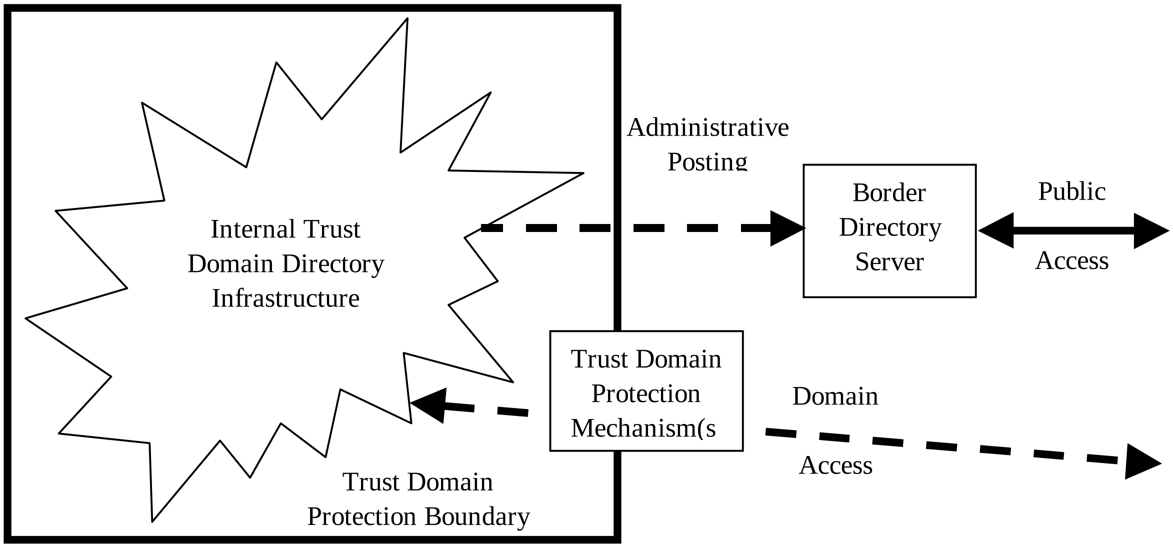

Figure 6. Free Access Border Directory Server

2.3.4 Restricted Access Border Directory Server

In this example, illustrated in Figure 7, the border directory server is implemented using a directory server that implements commercial-quality protection mechanisms that provide such features as end user authentication and identity-based access controls. The border directory server is installed on the unprotected side of the trust domain's protected infrastructure, which is protected by one or more independently validated security mechanisms (e.g., packet-filtering firewalls, application-level firewalls, message guards, or directory guards). The restricted access border directory server provides the capability to control directory requests based on user identification or class. This provides protection (based on the level of assurance) against compromise for the information stored in border directory server. In addition, this approach provides protection over the information stored in the trusted domain's internal directory infrastructure to the degree provided by the protection domain mechanisms.

In this case, administrative users within the trust domain are responsible for populating information out to the border directory. This allows the administrative users to segregate the posted information between information made available to the public and sensitive information made available only to authenticated users, including perhaps some group of individuals who are not in the domain, but have a need to access the restricted information. In this case, domain users may access the trust domain infrastructure via the protection boundary mechanisms. Such domain user access would not involve the border directory server.

An example of where this approach might apply would be a governmental law enforcement agency (e.g., FBI or DEA) that needs to provide unrestricted public access to directory information for a subset of employees or office codes. The border directory server would also provide restricted access to information that is necessary to support state and local law enforcement agencies. The trusted domain directory infrastructure would still need to be protected by one or more independently validated security mechanisms (e.g., packet-filtering firewalls, application-level firewalls, message guards, or directory guards). Agency users would access the trust domain infrastructure via the protection boundary mechanisms. Such domain user access would not involve the border directory server.

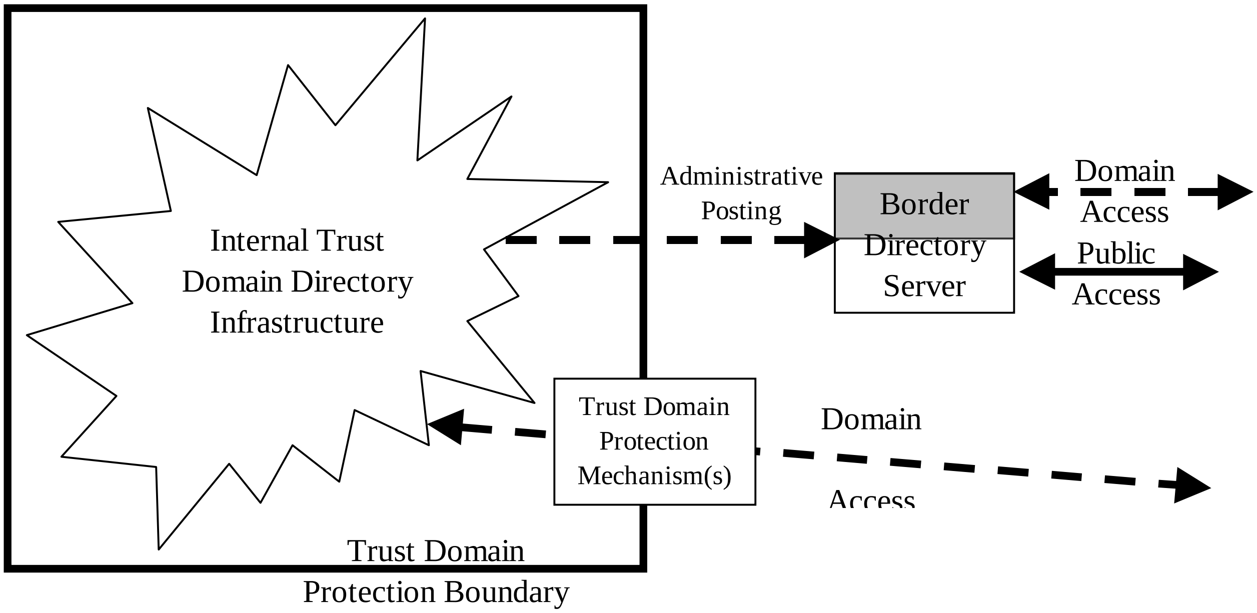

Figure 7. Restricted Access Border Directory Server

3 PROTECTION ISSUES

A significant area of concern is how to protect information within the organization's internal directory infrastructure from disclosure through the border directory server. The FPKI directory infrastructure is intended to provide public access to Federal PKI certificate information while at the same time preventing disclosure of sensitive information that may exist within the trust domains' directory infrastructure.

In this section, the issue of protecting both the FPKI directory and the agencies' internal infrastructures is examined from two perspectives: (1) posting information and (2) limiting incoming requests.

3.1 POSTING INFORMATION TO THE BORDER DIRECTORY SERVER

The only communications between a trust domain's internal directory infrastructure and its border directory server are expected to be: (1) restricted, administrative posting of appropriate information to the border directory server and (2) in support of certificate verification by relying parties inside the domain. The guidance for what to post could be stated as, "Post anything you wouldn't mind seeing on the front page of the New York Times."

By identifying who is authorized to post information and controlling exactly what type of information may be posted, a trust domain can gain some control over disclosure of information from within the trusted domain. The four example approaches in this section illustrate how posting information from the trusted domain infrastructure can be controlled and discuss security, performance, and operational impacts.

3.1.1 Separate PCA Posting to BCA Directory Server

In this approach, illustrated in Figure 2 by Trust Domain 3, the PCA posts information independently to the internal directory infrastructure and the BCA directory server.

- Security: Security is relatively good, since the PCA is responsible for posting the specific information from the trust domain to the BCA directory server, the posted information can be limited to only nonsensitive information.
- Performance: Query performance should be relatively good, since the first directory server encountered on outgoing requests is the BCA directory server. On incoming requests, the path to the BCA directory server is expected to be relatively short, since the BCA directory server is expected to be connected via a short path to each border directory server.
- Operational Impact: Since the PCA is responsible for sanitizing and posting information to the directory server, this is expected to be a major impact.

3.1.2 Separate CA Posting to Border Directory Server

In this approach, illustrated in Figure 8, the CA posts directory information to both the internal directory infrastructure and the border directory server, but communicates with only one of the two at any time. Administrator activities on both the internal directory system and border directory server should be audited and the audit logs periodically inspected.

- Security: Security is relatively good, since the PCA is responsible for posting the specific information from the trust domain to the border directory server, the posted information can be limited to only nonsensitive information.
- Performance: Query performance would generally not be as good as the preceding example. The first directory server encountered on outgoing requests could either be the local border directory server or the BCA directory server. On incoming requests, the path to the directory server would include at least one additional directory server (i.e., the BCA directory server). In the event communications between two trust domains is frequent, a direct link between the two directory servers could improve this performance.
- Operational Impact: Since the PCA is responsible for sanitizing and posting information to the directory server, this is expected to be a major impact.

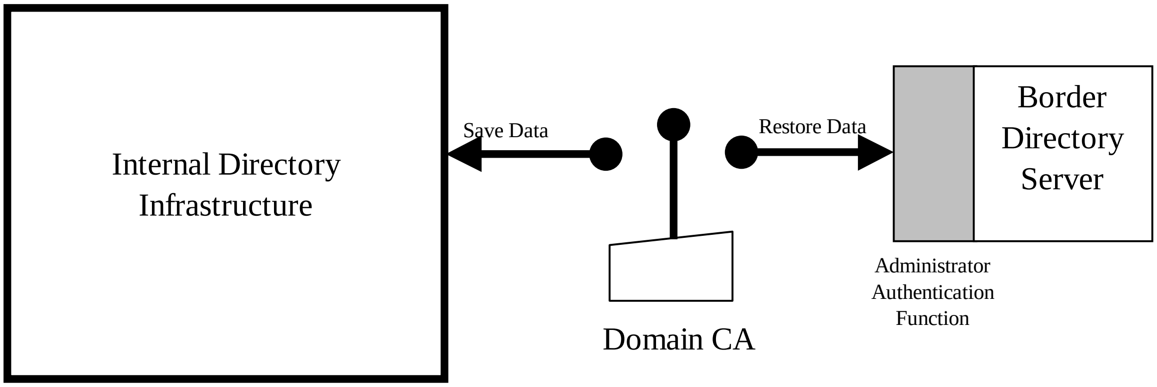
Internal Directory Infrastructure
Save Data
Restore Data
Border Directory Server
Administrator Authentication Function
Domain CA

Figure 8. Separate Administrative Posting

3.1.3 Administrative Posting from Domain Infrastructure

In this approach, illustrated in Figure 9, the administrator uses administrative and user interfaces (e.g., DAP or LDAP) to the internal directory infrastructure to extract information that is then posted to the border directory server using similar administrative and user interfaces. The border directory server would need to perform an authentication check against the administrator, who must be authorized to write or modify information.

Administrative posting to the border directory server should be audited and the audit logs periodically inspected.

- Security: Security is relatively good, since the PCA is responsible for posting the specific information from the trust domain to the border directory server. The posted information can be limited to only nonsensitive information.
- Performance: Query performance would be the same as the preceding example, since the border directory configuration is the same.
- Operational Impact: Since the PCA is responsible for sanitizing and posting information to the directory server, this continues to be a major impact.

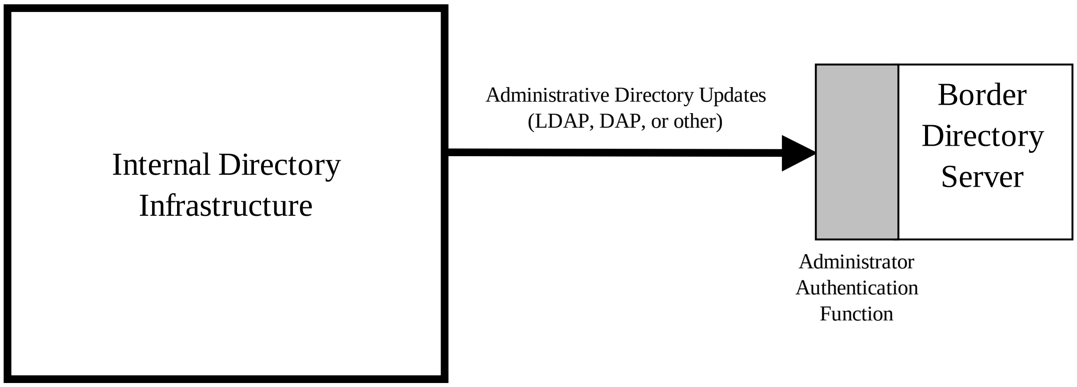
Internal Directory Infrastructure
Administrative Directory Updates (LDAP, DAP, or other)
Border Directory Server
Administrator Authentication Function

Figure 9. Networked Administrative Posting

3.1.4 Replication (Shadowing) from Domain Infrastructure

In this approach, illustrated in Figure 9, a directory server inside the trust domain's internal directory infrastructure is designated as a "Master" Directory System Agent (DSA) in a shadowing agreement with a "Shadow" border or BCA directory server. The shadowing agreement identifies specific directory subtrees to replicate from the master to the shadow directory server. The agreement describes the frequency (e.g., on change or periodic) and nature (e.g., supplier-initiated or consumer-initiated) of shadow updates. Construction and activation of shadowing agreements should be restricted to trusted administrators and access to the agreements (e.g., creation, initiation, deletion should be audited and the audit logs periodically inspected.

- Security: Security is relatively weaker than the preceding examples, since the specification of information to be shadowed may not allow sensitive and nonsensitive information to be adequately separated.
- Performance: Query performance would be the same as the preceding example, since the border directory configuration is the same.
- Operational Impact: Since information is replicated via one or a small number of existing shadow agreements, the operational impact should be substantially less than the preceding examples. Depending on the shadow agreement, replication can be performed automatically.

3.2 LIMITING INCOMING DIRECTORY REQUESTS

The only information flows from a border directory server to its local domain infrastructure should be to provide relying parties within the domain with appropriate responses. In particular, a trust domain should not permit its border directory server to chain into its own internal directory infrastructure (e.g., by using knowledge references) to satisfy incoming directory requests. In the event that outside access to the internal directory infrastructure is desired (e.g., for traveling members of the trust domain) this outside access should use a separate path, owned by the trust domain. The separate path should be suitably protected (e.g., network-level firewalls and proxies, application-level active firewalls, encrypted dial-up access, or guards), and suitable cryptographic protections (e.g., message authentication codes, digital signatures, or encryption) should be applied to the transmitted or received information.

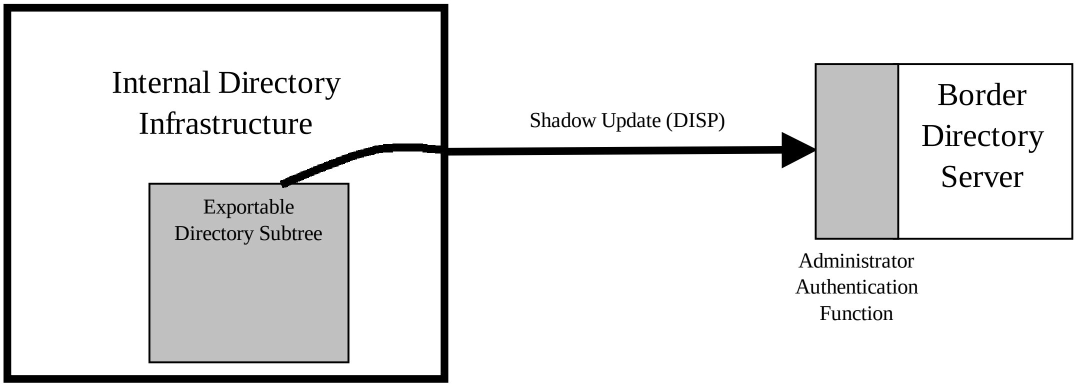
Internal Directory Infrastructure
Exportable Directory Subtree
Shadow Update (DISP)
Border Directory Server
Administrator Authentication Function

Figure 10. Replication (Shadowing)

3.3 LIMITING MALICIOUS INPUT TO THE BORDER DIRECTORY

External users (i.e., users from other trust domains or the public) should not be able to post anything to the border directory server. In the event that a trust domain wants to post external user information, it is recommended that (1) the information be submitted to an administrator using an out-of-band mechanism (e.g., digitally signed email or secure web transaction), (2) the administrator use suitable means to ensure the “goodness” of the information, and (3) that the administrator post the information using one of the foregoing mechanisms.

4 DESIGN ISSUES

This section discusses a number of design issues relative to the fielding of the FPKI directory. The issues are not presented in any order.

4.1 DIRECTORY INFORMATION BASE SCHEMA

The FPKI directory schema is intended to provide the minimal set of rules necessary to define the directory entry types, object classes, attributes, matching rules, name forms, and structure rules to support interoperability among the users of the FPKI directory. At present, the FPKI directory schema will be designed to conform to the Internet X.509 PKI LDAPv2 Schema [LDAP-S].

The National Technical Information Service is currently developing the U. S. Government On Line Directories (USGOLD) [USGold] as the Federal Government's directory for telephone and email access. Since the Federal PKI directory and the USGOLD may eventually coalesce, the USGOLD directory specifications will also be followed, where applicable.

4.2 TIME SYNCHRONIZATION FOR CHAINED QUERIES

Each directory service (e.g., read, list, or search) can include a timeLimit parameter, which indicates the maximum elapsed time, in seconds, within which the service must be provided [X.511]. Furthermore, "the timeLimit component does not imply the length of time spent processing the request during the elapsed time: any number of DSAs [i.e., directory servers] may be involved in processing the request during the elapsed time." In a distributed environment, lack of time synchronization among directory servers could raise a significant problem, especially in chained queries. A server with current time that differs from another by an amount greater than the timeout period could cause protocol servicing to timeout immediately once that server is reached.

This potential for such premature service timeouts can be remedied using either of two approaches:

1 The timeLimit parameter can be omitted, which implies no time limit. Since loops in protocol processing are detected as part of their processing, this approach would only be a problem if completion of a service actually takes a long time.

Excessively long interrogation service execution time can be addressed by issuing an abandon service request.

2 Directory servers could be required to periodically synchronize their clocks with an agreed-upon time source to maintain clock synchronization within some defined number of seconds. The timeLimit parameter would then be useable for times greater than the variance from the time source.

The first approach should require no hardware or software modifications (assuming directory user agent software provides the abandon service), but would require users to be aware of the abandon operation. The second approach would be transparent to the users, but would require some engineering effort to provide clock synchronization among all directory servers in the FPKI directory system.

4.3 DIRECTORY INTEGRITY

4.3.1 FPKI Directory Server Authentication

Each directory server in the Federal PKI directory should support strong authentication (and signed operations, if possible) for communications among servers. This provides corroboration that an entity in an instance of communication is the one claimed.

4.3.2 Data Integrity

Data integrity provides proof of the integrity of the information, either in storage or during transmission. The mechanism involves encryption of a compressed string (e.g., a message authentication code – MAC) of the relevant data to be stored or transmitted. The encryption is performed using the private digital signature key of the sender. Where integrity is required on directory information (e.g., digital signature public key certificates) the information must be signed by an appropriate authority (e.g., a CA or PCA)

4.4 DIRECTORY MANAGEMENT

4.4.1 Availability

Generally, the availability requirement for the FPKI directory is expected to be 24 by 7. Deviations from scheduled full time availability should be negotiated between the trust domain PCA and the FPMA.

4.4.2 Key Management

Common cryptographic algorithms and their intended usage need to be supported. The trust domain’s CPS, which should be published on its border directory server, should define acceptable cryptographic algorithms and required usages. External relying parties can use the CPS as guidance to facilitate the necessary interoperability. Furthermore, the use of certificate revocation information (e.g., certificate revocation lists – CRLs) should be documented as part of the domain’s CPS to support certificate path validation. This should describe any mechanisms used to assure the applicability (e.g., freshness).of the revocation information

Each trust domain is responsible for implementation and maintenance of its own certificate management infrastructure (CMI), including CAs, to ensure the ability for an external relying party to construct certificate paths from its trust domain to a cross-certified entity within the FPKI.

4.4.3 Unique User Identification

Within each trust domain, the PCA is responsible for assuring that each user has a unique identification (i.e., DNS name or distinguished name). The FPMA is responsible for assuring that each trust domain has a unique identification. The combination of unique user id and unique domain id should be combined to provide FPKI directory unique identifications for each user.

4.5 SHADOWING

Shadowing is the replication capability provided as part of the X.500 directory structure [X.525]. At present, although many directory products support shadowing using the Directory Information Shadowing Protocol (DISP), interoperability among different vendors’ products is rare. Nonetheless, use of replication, where possible, could improve directory query performance significantly.

In Section 3, directory shadowing was identified as a mechanism that a trust domain can use to populate its border directory. The domain’s border directory server must include a logically complete (or at least sufficiently complete) set of information relative to the domain, since external queries would otherwise need to chain into the trust domain’s internal directory infrastructure.

More generally, replication can be used to improve the performance of the FPKI directory by reducing the number of referrals or chains needed to satisfy a query. Replication of BCA directory information on each border directory would reduce the need to chain to the BCA directory merely to be redirected to another border directory. Also, the information in that directory is expected to be more static than other directories. Thus, the likelihood of inconsistencies between master and shadow directories would be a less critical issue.

Replication of information between border directories or from border directories to the BCA directory can also be expected to improve performance on directory queries to remote domains. However, information in border directories is expected to be considerably less static than that in the BCA directory. Thus, the likelihood of inconsistency between master and shadow (or shadow and secondary shadow) directory information is expected to be a more critical issue, requiring more frequent updates than would be expected if only BCA directory information were shadowed.

Shadow agreements among border directories requires coordination among the PCAs responsible for each party of the agreement. If the border directory is a shadow that is mastered within the trust domain's internal directory infrastructure, a secondary shadowing agreement would be necessary to provide information from one border directory to another. This would require coordination among the border directory PCAs and the authority responsible for the internal directory infrastructure that masters the originating border directory shadow.

5 FEDERAL PKI DIRECTORY EVOLUTION

The Federal PKI Directory will be an evolutionary enterprise. The initial BCA directory implementation will support the following three models of client access:

1. The client accesses its internal directory server as is done today (i.e., via whatever mechanism is currently in place). The internal server chains to its border directory server, which chains to the BCA directory server, which may continue the chain as necessary.
2. The client accesses its internal directory server as is done today. The internal server chains directly to the BCA directory server, which may continue the chains as necessary.
3. The client accesses its internal directory server using LDAP v3. If the server does not have the data, it returns a referral to the client. The referral may identify the BCA directory server or one or more border directory servers or both.

The initial BCA directory server will be a “proper” X.500 directory system agent (DSA) that implements DSP chaining and also implements LDAP v3 client access. The BCA DSA must be able to receive an LDAP query and resolve it by chaining to other DSAs.

In the second stage, the BCA DSA will support two new capabilities. First, it will provide referrals to LDAP v3 clients when the directory server holding the information only supports LDAP and does not provide chaining. Second, it will provide an “LDAP query gateway,” where the query arrives at the BCA from a DSA as a chained DSP request. The LDAP query gateway will process the request using LDAP operations for LDAP-only servers.

The goal is to provide a border directory server to host each agency’s externally accessible certificate information. This does not, however, mean that each trust domain is expected to “stand up” its own border directory server. Some trust domains have indicated a willingness to host the certificate information for “subscriber” trust domains on their border directory server. Such subscriber agreements could help in populating the Federal PKI Directory in the near term. Alternatively, a trust domain could engage an external contractor to provide the border directory server service on their behalf.

APPENDIX: DIRECTORY PROTOCOLS

This section provides a short overview of the directory protocol suites discussed in this paper.

X.500 PROTOCOL SUITE

DAP is the X.500 protocol that provides access between the X.500 User Agent (UA) and the X.500 Directory System Agent (DSA). The X.500 protocols are specified in Abstract Syntax Notation 1 (ASN.1) and encoded using various encoding mechanisms (e.g., Basic Encoding Rules - BER, Distinguished Encoding Rules - DER, Packed Encoding Rules - PER). In the event that the DSA cannot resolve the user's request, the DAP response can include referral information. DAP is just one of a suite of four directory protocols provided under the X.500 umbrella (X.519). Directory System Protocol (DSP) is used to provide a distributed directory environment and is used to support DAP "chaining," in which the DSA that received the DAP request forwards the request to one or more other DSAs [X.518]. Directory Information Shadowing Protocol (DISP) is provided to support replication, again in the distributed directory environment. Directory Operational Binding Management Protocol (DOP) was intended for use in authentication among DSAs, but is rarely implemented in X.500-compliant products.

A significant issue relative to X.500 directory products results from the standard itself. Since the standard admits many compliant variations, products that are X.500 compliant often cannot interoperate. This is particularly true with the DISP protocol, which works well within a single vendor's product line, but rarely works at all when implemented across multiple vendors' products.

Directory Access Protocol (DAP)

DAP supports connection, read, search, and modify packages. Each package supports a set of protocol service requests and their responses as follows (responses not shown):

- Connection Package: directoryBind and directoryUnbind
- Read Package: read, compare, and abandon
- Search Package: list and search
- Modify Package: addEntry, removeEntry, modifyEntry, and modifyDN

Directory System Protocol (DSP)

The DSP protocol includes the following packages and services:

- Connection Package: dSABind and dSAUnbind
- Chained Read Package: chainedRead, chainedCompare, and chainedAbandon
- Chained Search Package: chainedList and chainedSearch
- Chained Modify Package: chainedAddEntry, chainedRemoveEntry, chainedModifyEntry, and chainedModifyDN

Directory Information Shadowing Protocol (DISP)

The DISP protocol includes the following packages and services:

- Connection Package: dSAShadowBind and dSAShadowUnbind
- Shadow Consumer Package: requestShadowUpdate (by consumer) and updateShadow(by supplier)
- Shadow Supplier Package: coordinateShadowUpdate and updateShadow (both by supplier)

Directory Operational Binding Management Protocol (DOP)

The DOP protocol includes the following packages and services:

- DOP Connection Package: dSAOperationalBindingManagementBind and dSAOperationalBindingManagementUnbind
- DOP Package: establishOperationalBinding, modifyOperationalBinding, and terminateOperationalBinding

LIGHTWEIGHT DIRECTORY ACCESS PROTOCOL (LDAP)

LDAP was developed by the Internet Engineering Task Force (IETF) as an alternative to DAP. The source document for LDAP is IETF RFC 1777 [LDAP] (this is referred to as "LDAPv2" in RFC 2251, the specification for LDAPv3). The initial intention was to eliminate the complexity of the X.500 protocol suite, thereby making directory access more efficient.

LDAP Version 2 (IETF RFC 1777)

LDAP protocol elements are carried directly over TCP or other transport layer protocols, the intent being to bypass both session and presentation protocol layers. Many of the protocol data elements (e.g., DNs) are encoded as ordinary strings. A "lightweight" BER encoding is used for all protocol elements. The following protocol operations are provided by LDAP:

- bindRequest
- unbindRequest

- searchRequest
- modifyRequest
- addRequest
- delRequest
- modifyRDNRequest
- compareDNRequest
- abandonRequest
- bindResponse
- unbindResponse
- searchResponse
- modifyResponse
- addResponse
- delResponse
- modifyRDNResponse
- compareDNResponse

Kerberos version 4 authentication to the LDAP server is provided using the bindRequest operation. No provision is made for distributed servers, thus referral, chaining, and replication are not supported.

LDAP Version 3 (IETF RFC 2251)

LDAP Version 3 is documented in IETF RFC 2251 [LDAPv3]. It supports all protocol elements of LDAPv2 (i.e., RFC 1777), but includes the following new features:

- Referrals to other servers may be returned
- Simple Authentication and Security Layer (SASL) mechanisms may be used to provide association security services
- Attribute values and DNs have been internationalized using ISO 10646
- Extensibility is provided to support new operations or extend existing operations
- The schema is published in the directory

LDAP v3 protocol elements are carried directly over TCP or other transport layer protocols, as in LDAP v2. The protocol elements of LDAP v3 are encoded using a restricted version of BER. The protocol operations provided in LDAP v2 continue to be supported in LDAP v3 with the following exceptions:

- The unbindResponse operation has been removed from LDAP v3. An unsolicited notification LDAP message type, which includes a “notification of disconnection” is provided to notify the client that a connection has been closed due to an error condition. No response is provided to an unbindRequest operation, but the client may assume that the server will terminate the protocol session.
- The searchResponse operation has been replaced with searchResDone and searchResRef to provide the referral capability. The searchResRef operation includes a list of LDAP URLs that can be used by the client to complete the search request.
- Two new operations, extendedReq and extendedResp have been added to promote future protocol extensibility

The bindRequest and bindResponse operations have been modified to support SASL mechanisms to provide association security services. The searchResRef response was provided in support of distributed servers, however LDAP v3 provides no protocol-level support for either chaining or replication.

REFERENCES

[BURR] Proposed Federal Public Key Infrastructure Concept of Operations, 4 September, 1998

[CHOK] Certificate Policy and Certification Practices Framework, S. Chokhani and W. Ford, Informational RFC, IETF PKIX Part IV, July 1997.

[CONOPS] TWG-98-31, *Draft Federal PKI Concept of Operations*, 3 June 1998

[LDAP] IETF RFC 1777, “Lightweight Directory Access Protocol,” March 1995.

[LDAP-S] S. Boeyen, T. Howes, and P. Richard, “Internet X.509 Public Key Infrastructure LDAPv2 Schema,” Internet Draft <draft-ietf-pkix-LDAPv2-schema-02.txt>, IETF PKIX Working Group, September 1998.

[LDAPV3] IETF RFC 2251, “Lightweight Directory Access Protocol (v3),” December 1997.

[TWG-98-29] W. E. Burr, “Proposed Federal PKI Architecture,” 19 May 1998

[USGold] Electronic Mail Program Management Office, Office of Management and Budget, “Strategic Planning Guidance for the Government Electronic Directory”

[X.511] ITU-T Recommendation X.511: “Information Technology – Open Systems Interconnection – The Directory: Abstract Service Definition,” November 1993.

[X.518] ITU-T Recommendation X.518: “Information Technology – Open Systems Interconnection – The Directory: Procedures for Distributed Operation,” November 1993.

[X.519] ITU-T Recommendation X.519: “Information Technology – Open Systems Interconnection – The Directory: Protocol Specifications,” November 1993.

[X.525] ITU-T Recommendation X.525: “Information Technology – Open Systems Interconnection – Replication,” June 1997.

ACRONYMS

ASN.1	Abstract Syntax Notation One
BCA	Bridge Certification Authority
BER	Basic Encoding Rules
CA	Certification Authority
CRL	Certificate Revocation List
DAP	Directory Access Protocol (X.500)
DER	Distinguished Encoding Rules
DIB	Directory Information Base
DISP	Directory Information Shadowing Protocol (X.500)
DIT	Directory Information Tree
DOP	Directory Operational Binding Management Protocol (X.500)
DPMA	Domain Policy Management Authority
DSA	Directory System Agent (X.500)
DSP	Directory System Protocol (X.500)
FPKI	Federal Public Key Infrastructure
FPMA	Federal Policy Management Authority
IETF	Internet Engineering Task Force
LDAP	Lightweight Directory Access Protocol
PCA	Principal Certification Authority
PER	Packed Encoding Rules
PKI	Public Key Infrastructure
SASL	Simple Authentication and Security Layer
S/MIME	Secure Multipurpose Internet Mail Extensions
TWG	Technical Working Group
USGOLD	U. S. Government On Line Directories