

PKIX Working Group  
INTERNET DRAFT

A. Arsenault  
DOD  
S. Turner  
IECA

Expires April 22, 2000

October 22, 1999

Internet X.509 Public Key Infrastructure  
PKIX Roadmap  
<[draft-ietf-pkix-roadmap-04.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of 6 months and may be updated, replaced, or may become obsolete by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of current Internet-Drafts Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Abstract

This document provides an overview or "roadmap" of the work done by the IETF PKIX working group. It describes some of the terminology used in the working group's documents, and the theory behind an X.509-based Public Key Infrastructure and Privilege Management Infrastructure (PMI). It identifies each document developed by the PKIX working group, and describes the relationships among the various documents. It also provides advice to would-be PKIX implementors about some of the issues discussed at length during PKIX development, in hopes of making it easier to build implementations that will actually interoperate.

INTERNET DRAFT

PKIX Roadmap

22 October 1999

## [1](#) Introduction

### [1.1](#) This Document

This document is an informational Internet draft that provides a "roadmap" to the documents produced by the PKIX working group. It is intended to provide information; there are no requirements or specifications in this document.

[Section 2](#) of this document defines key terms used in this document. [Section 3](#) covers "PKIX theory;" it explains what the PKIX working group's basic assumptions were. [Section 4](#) provides an overview of the various PKIX documents. It identifies which documents address which areas, and describes the relationships among the various documents. [Section 5](#) contains "Advice to implementors." Its primary purpose is to capture some of the major issues discussed by the PKIX working group, as a way of explaining WHY some of the requirements and specifications say what they say. This should cut down on the number of misinterpretations of the documents, and help developers build interoperable implementations. [Section 6](#) contains a list of contributors we wish to thank. [Section 7](#) provides a list references. [Section 8](#) discusses security considerations, and [Section 9](#) provides contact information for the editors. Finally, [Section 10](#) provides a disclaimer.

### [1.2](#) Changes Since Last Version

Updated references to current drafts.

Added terminology to paragraph 2 for Attribute Authorities, Attribute Certificates, Certificates, Public Key Certificates, Public Key Infrastructure, and Privilege Management Infrastructure.

Split paragraph 3.1 and 3.3 in two to allow separate descriptions for PKI (i.e., PKC discussions) and PMI (i.e., AC discussions).

Updated PKIX History (paragraph 3.2).

Split paragraph 3.4 in two to accomodate discussions for both X.509 Certificates and Attribute Certificates.

Updated Profiles, Operational Protocols, and Management Protocols paragraphs 3.6.1, 3.6.2, and 3.6.3, respectively.

Updated Revocation paragraph 3.5.8 to indicate why a certificate is retained on a CRL for one additional period.

Added descriptions for new drafts in 4.1-4.5: Operation Protocols -

LDAPv3, Simple Certificate Validation Protocol (SCVP), Using HTTP as Transport Protocol for CMP, Using TCP as Transport Protocol for CMP, Limited Attribute Certificate Acquisition Protocol (LAAP), OCSP Extensions

Added paragraph 4.6 to talk about drafts that didn't make it through working group review.

Removed references to [BERT1], [CACHE], and [WEB] from paragraph 7.

### [1.3](#) To Do

Add text in paragraph 5.3 to talk about extended key usages.

Add in paragraph to talk about PMI functions.

Add in paragraph to talk about delta between [RFC2459](#) and the updated [RFC2459](#).

## [2](#) Terminology

There are a number of terms used and misused throughout PKI-related and PMI-related literature. To limit confusion caused by some of those terms, throughout this document, we will use the following terms in the following ways:

- Attribute Authority (AA) - An authority trusted by one or more users to create and sign attribute certificates. It is important to note that the AA is responsible for the attribute certificates during their whole lifetime, not just for issuing them.
- Attribute Certificate (AC) - A data structure containing a set of attributes for an end-entity and some other information, which is digitally signed with the private key of the AA which issued it.

- Certificate - Can refer to either an AC or a public key certificate. Where there is no distinction made the context should be assumed to apply to both an AC and a public key certificate.
- Certification Authority (CA) - An authority trusted by one or more users to create and assign public key certificates. Optionally the CA may create the user's keys. It is important to note that the CA is responsible for the public key certificates during their whole lifetime, not just for issuing them.

- Certificate Policy (CP) - A named set of rules that indicates the applicability of a public key certificate to a particular community and/or class of application with common security requirements. For example, a particular certificate policy might indicate applicability of a type of public key certificate to the authentication of electronic data interchange transactions for the trading of goods within a given price range.
- Certification Practice Statement (CPS) - A statement of the practices which a CA employs in issuing public key certificates.
- End-entity - A subject of a certificate who is not a CA in the PKIC or an AA in the PMI. (An EE from the PKI can be an AA in the PMI.)
- Public Key Certificate (PKC) - A data structure containing the public key of an end-entity and some other information, which is digitally signed with the private key of the CA which issued it.
- Public Key Infrastructure (PKI) - The set of hardware, software, people, policies and procedures needed to create, manage, store, distribute, and revoke PKCs based on public-key cryptography.
- Privilege Management Infrastructure (PMI) - A collection of ACs, with their issuing AA's, subjects, relying parties, and repositories, is referred to as a Privilege Management Infrastructure.
- Registration Authority (RA) - An optional entity given

responsibility for performing some of the administrative tasks necessary in the registration of subjects, such as: confirming the subject's identity; validating that the subject is entitled to have the values requested in a PKC; and verifying that the subject has possession of the private key associated with the public key requested for a PKC.

- Relying party - A user or agent (e.g., a client or server) who relies on the data in a certificate in making decisions.
- Root CA - A CA that is directly trusted by an EE; that is, securely acquiring the value of a Root CA public key requires some out-of-band step(s). This term is not meant to imply that a Root CA is necessarily at the top of any hierarchy, simply that the CA in question is trusted directly.
- Subordinate CA - A "subordinate CA" is one that is not a Root CA for the EE in question. Often, a subordinate CA will not be a Root CA for any entity but this is not mandatory.

- Subject - A subject is the entity (AA, CA, or EE) named in a certificate. Subjects can be human users, computers (as represented by Domain Name Service (DNS) names or Internet Protocol (IP) addresses), or even software agents.
- Top CA - A CA that is at the top of a PKI hierarchy.

Note: This is often also called a "Root CA," from since in data structures terms and in graph theory, the node at the top of a tree is the "root." However, to minimize confusion in this document, we elect to call this node a "Top CA," and reserve "Root CA" for the CA directly trusted by the EE. Readers new to PKIX should be aware that these terms are not used consistently throughout the PKIX documents, as [[FORMAT](#)] uses "Root CA" to refer to what this and other documents call a "Top CA," and "most-trusted CA" to refer to what this and other documents call a "Root CA."

## [3](#) PKIX Theory

### [3.1](#) Certificate-using Systems

### 3.1.1 Certificate-using Systems and PKIs

At the heart of recent efforts to improve Internet security are a group of security protocols such as Secure Multipurpose Internet Mail Extensions (S/MIME), Transport Layer Security (TLS), and Internet Protocol Security (IPSec). All of these protocols rely on public-key cryptography to provide services such as confidentiality, data integrity, data origin authentication, and non-repudiation. The purpose of a PKI is to provide trusted and efficient key and public key certificate management, thus enabling the use of authentication, non-repudiation, and confidentiality.

Users of public key-based systems must be confident that, any time they rely on a public key, the associated private key is owned by the subject with which they are communicating. (This applies whether an encryption or digital signature mechanism is used.) This confidence is obtained through the use of PKCs, which are data structures that bind public key values to subjects. The binding is achieved by having a trusted CA verify the subject's identity and digitally sign each PKC.

A PKC has a limited valid lifetime which is indicated in its signed contents. Because a PKC's signature and timeliness can be independently checked by a certificate-using client, PKCs can be distributed via untrusted communications and server systems, and can be cached in unsecured storage in certificate-using systems.

PKCs are used in the process of validating signed data. Specifics vary according to which algorithm is used, but the general process works as follows:

Note: there is no specific order in which the checks listed below must be made; implementors are free to implement them in the most efficient way for their systems.

- The recipient of signed data verifies that the claimed identity of the user is in accordance with the identity contained in the PKC;
- The recipient validates that no PKC in the path is revoked (e.g., by retrieving a suitably-current Certificate Revocation List

(CRL) or querying an on-line certificate status responder), and that all PKCs are within their validity periods at the time the data was signed;

- The recipient verifies that the data are not claimed to have any values for which the PKC indicates that the signer is not authorized;
- The recipient verifies that the data have not been altered since signing, by using the public key in the PKC.

If all of these checks pass, the recipient can accept that the data was signed by the purported signer. The process for keys used for encryption is similar.

Note: It is of course possible that the data was signed by someone very different from the signer, if for example the purported signer's private key was compromised. Security depends on all parts of the certificate-using system, including but not limited to: physical security of the place the computer resides; personnel security (i.e., the trustworthiness of the people who actually develop, install, run, and maintain the system); the security provided by the operating system on which the private key is used; and the security provided the CA. A failure in any one of these areas can cause the entire system security to fail. PKIX is limited in scope, however, and only directly addresses issues related to the operation of the PKI subsystem. For guidance in many of the other areas, see [POLPROC].

### [3.1.2](#) Certificate-using Systems and PMIs

Many systems use the the PKC to perform identity based access control decisions (i.e., the identity may be used to support identity-based

access control decisions after the client proves that it has access to the private key that corresponds to the public key contained in the PKC). For many systems this is sufficient, but increasingly systems are beginning to find that rule-based, role-based, and rank-based access control is required. These forms of access control decisions require additional information that is normally not included in a PKC, because the lifetime of the information is much shorter than the lifetime of the public-private key pair. To support

binding this information to a PKC the Attribute Certificate (AC) was defined in ANSI and later incorporated into ITU-T Recommendation X.509. The AC format allows any additional information to be bound to a PKC by including, in a digitally signed data structure, a reference back to one specific PKC or to multiple PKCs, useful when the subject has the same identity in multiple PKCs. Additionally, the AC can be constructed in such a way that it is only useful at one or more particular targets (e.g., web server, mail host).

Users of a PMI must be confident that the identity purporting to possess an attribute has the right to possess that attribute. This confidence may be obtained through the use of PKCs or it may be configured in the AC-using system. If PKCs are used the party making the access control decision can determine "if the AC issuer is trusted to issue ACs containing this attribute."

### [3.2](#) PKIX History

In the beginning there was ITU-T Recommendation X.509. It defines a widely accepted basis for a PKI, including data formats and procedures related to distribution of public keys via PKCs digitally signed by CAs. X.509 does not however include a profile to specify the support requirements for many of the PKC data structure's sub-fields, for any of the extensions, nor for certain data values. PKIX was formed in October 1995 to deliver a profile for the Internet PKI of X.509 version 3 PKCs and version 2 CRLs. The Internet PKI profile [[FORMAT](#)] went through eleven draft versions before becoming an RFC. Other profiles have been developed in PKIX for particular algorithms to make use of [[FORMAT](#)]. There has been no sense of conflict between the groups that developed these profiles as they are seen as complimentary.

The development of the management protocols has not been so straightforward. The Certificate Management Protocol (CMP) was developed to define a message protocol that is used between entities in a PKI [[CMP](#)]. The demand for an enrollment protocol and the desire to use PKCS-10 message format as the certificate request syntax lead to the development of two different documents in two different groups. The Certificate Request Syntax (CRS) draft was developed in the SMIME WG which used PKCS-10 [[PKCS10](#)] as the certification request

message format. Certificate Request Message Format [[CRMF](#)] draft was

also developed but in the PKIX WG. It was to define a simple enrollment protocol that would subsume both the CMP and CRS enrollment protocols, but it did not use PKCS-10 as the certificate request message format. Then the certificate management message format document, was developed to define an extended set of management messages that flow between the components of the Internet PKI. Certificate Management Messages over CMS (CMC) was developed to allow the use of an existing protocol (S/MIME) as a PKI management protocol, without requiring the development of an entirely new protocol such as CMP [[CMP](#)]. It also included [[PKCS10](#)] as the certificate request syntax, which caused work on the CRS draft to stop. Information from the certificate management message format document was moved into [[CMP](#)] and [[CMC](#)] so work on the certificate management message format document was discontinued.

Another long debated topic in the WG dealt with certificate revocation. Numerous drafts have been developed to address different issues related certificate revocations. CMP supports revocation request, response, revocation announcement, and requests for CRL messages. CMC defines revocation request, revocation response, and requests for CRL messages messages, but uses CMS as the encapsulating protocol. [[OCSP](#)] was developed to address concerns that not all relying parties want to go through the process checking CRLs from every CA in the certification path. It defines an on-line mechanism to determine the status of a given certificate, which may provide more timely revocation information than is possible with CRLs. The Simple Certification Verification Protocol (SCVP) was produced to allow relying parties to off-load all of their certification verification to another entity [[SCVP](#)]. The WG was arguable split over whether such a function should be supported and whether it should be its own protocol or included in OCSP. In response, a draft defining OCSP Extensions [[OCSPX](#)] was produced to include the functions of SCVP. One other draft called Open CRL Distribution Point (OCDP) was produced which documented two extensions: one to support an alternative CRL partitioning mechanism to the CRL Distribution Point mechanism documented in [[FORMAT](#)] and one to support identifying other revocation sources available to certificate-users. The work from this draft was subsumed by an ITU-T | ISO/IEC Amendment to X.509, hence work on this draft was halted.

Development of the operational protocols has been slightly more straightforward. Three documents for the Light Weight Directory Access Protocol (LDAP) have been developed one for defining LDAPv2 as an access protocol to repositories [[PKI-LDAPv2](#)]; one for storing PKI information in an LDAP directory [[SCHEMA](#)]; and one for LDAPv3 requirements for PKI [[PKI-LDAPv3](#)]. Using the File Transfer Protocol (FTP) and the Hyper Text Transmission Protocol (HTTP) to retrieve

PKCs and CRLs from PKI repositories was documented in [[FTPHTTP](#)].

In late 1998 the PKIX charter was revised to include protocols for time stamping and data certification services. [[TSP](#)] was developed to define protocols required to interact with a Time Stamp Authority (TSA) who asserts that a datum existed at a given time. Of course, if a true non-repudiation service is to be provided additional services that prove the data was actually in the possession of the subject requesting the time stamp. So, the [[DVCS](#)] draft was developed to provide two mechanisms to prove the subject actually possessed the data. In addition, [[DVCS](#)] provides two additional services: one to verify all signatures attached to the signed document using all appropriate status information and PKCs and one to verify and assert the validity of one or more PKCs at the specified time. Thoughtfully, [[DVCS](#)] permits the [[TSP](#)] protocol to be used as one of the time stamp tokens. Both [[DVCS](#)] and [[TSP](#)] use [[CMS](#)] as an encapsulating (though in [[TSP](#)] request for a time stamp are not required to use [[CMS](#)]). [[ETNPT](#)] was developed to use [[DVCS](#)] to maintain the trust in a token issued by a non-repudiation Trusted Third Party (NR TTP) after the key initially used to establish trust in the token expires.

Around the same time, a work item for ACs, defined in [[X.509](#)], was added. ACs are similar to PKCs, but they do not bind public keys to identities rather they bind attributes to identities. The attributes bound to the identity can represent anything, but are mostly used to support rule-based, role-based, and rank-based access control decisions. Two drafts have since been developed: the Internet Attribute Certificates Profile for Authorizations [[AC](#)] and the Limited AttributeCertificate Acquisition Protocol [[LAAP](#)]. The first profiles the fields and extensions of the AC and the second provides a deliberately limited protocol to access a repository when LDAP is not appropriate.

Other drafts have been produced to address specific issues. [[DHPOP](#)] was developed to define two mechanisms by which a signature can be produced using a Diffie-Hellman pair. This draft provides a mechanism to Diffie-Hellman key pairs to issue a PKCS-10 certification request. After some operational experience with [[CMP](#)], two drafts, one for using HTTP as the transport protocol [[CMP-HTTP](#)] and one for Transmission Control Protocol (TCP) [[CMP-TCP](#)], were written to solve problems encountered by implementors.

From the alphabet soup above, it is clear why this roadmap is required.

### [3.3](#) Overview of the PKIX Approach

#### [3.3.1](#) PKI

PKIX is an effort to develop specifications for a PKI for the Internet using PKCs. A PKI is defined as:

The set of hardware, software, people, policies and procedures needed to create, manage, store, distribute, and revoke PKCs based on public-key cryptography.

A PKI consists of five types of components [[MISPC](#)]:

- Certification Authorities (CAs) that issue and revoke PKCs;
- Organizational Registration Authorities (ORAs) that vouch for the binding between public keys and certificate holder identities and other attributes;
- Certificate holders that are issued certificates and can sign digital documents and encrypt documents;
- Clients that validate digital signatures and their certification paths from a known public key of a trusted CA;
- Repositories that store and make available certificates and Certificate Revocation Lists (CRLs).

Figure 1 is a simplified view of the architectural model assumed by the PKIX Working Group.

INTERNET DRAFT

PKIX Roadmap

22 October 1999

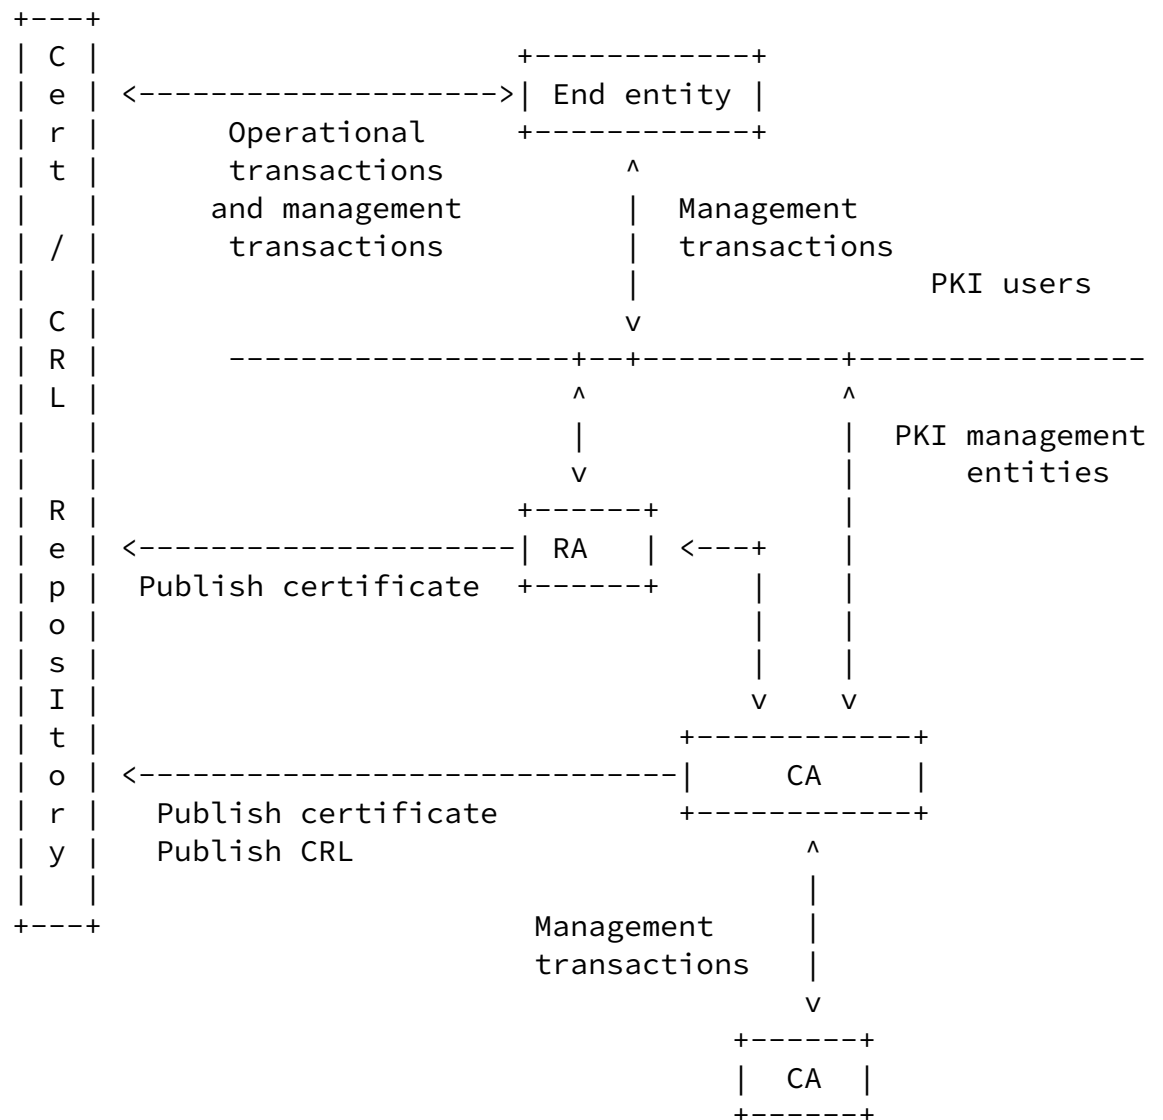


Figure 1 - PKI Entities

### 3.3.2 PMI

PKIX is also an effort to develop specifications for a Privilege Management Infrastructure for the Internet using ACs. A Privilege Management Infrastructure, or PMI, is defined as:

The set of hardware, software, people, policies and procedures needed to create, manage, store, distribute, and revoke ACs.

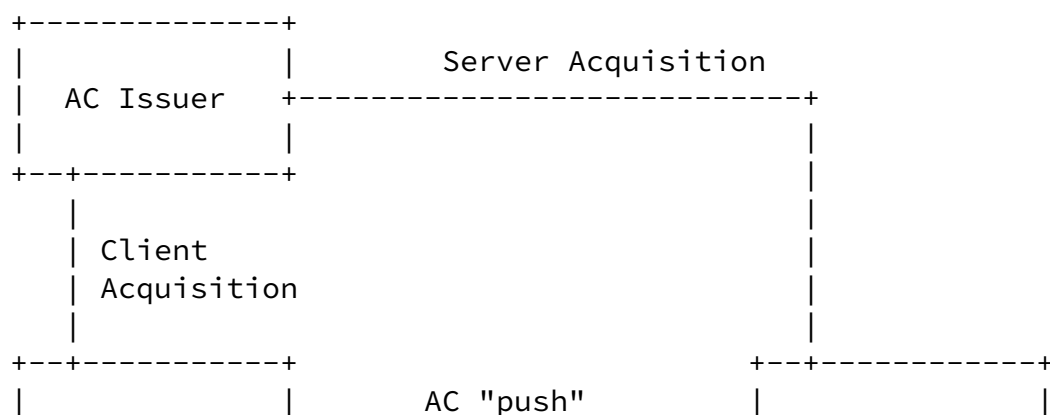
A PMI consists of five types of components [AC]:

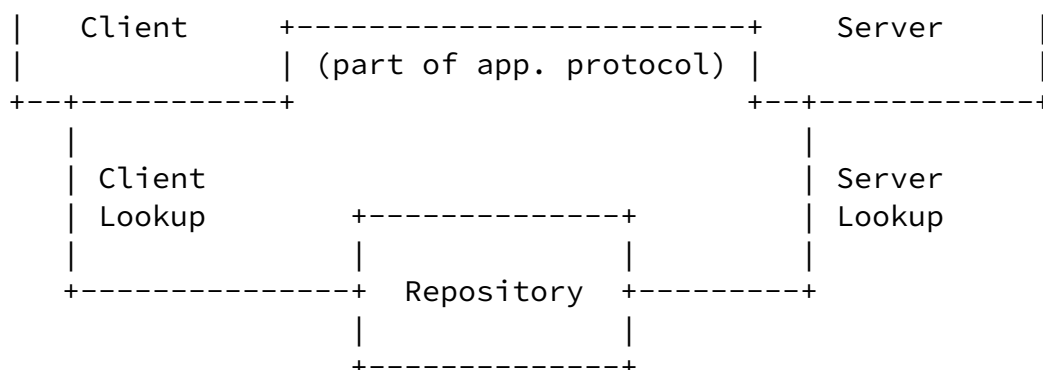
- Attribute Authorities (AAs), or Attribute Certificate Issuer, that issue and revoke ACs;

Note: AAs may implicitly revoke ACs by using very short validity periods.

- Attribute Certificate Users that parses or processes an AC;
- Attribute Certificate Verifiers that check the validity of an AC and then makes use of the result;
- Clients that request an action for which authorization checks are to be made;
- Repositories that store and make available certificates and Certificate Revocation Lists (CRLs).

Figure 2 is an example of the exchanges that may involve ACs.





### 3.4.1 Public Key Certificates

standards or may be defined and registered by any organization or community. In June 1996, standardization of the basic v3 format was completed [[X.509](#)].

ISO/IEC/ITU and ANSI X9 have also developed standard extensions for use in the v3 extensions field [[X.509](#)][X9.55]. These extensions can convey such data as additional subject identification information, key attribute information, policy information, and certification path constraints. However, the ISO/IEC/ITU and ANSI X9 standard extensions are very broad in their applicability. In order to develop interoperable implementations of X.509 v3 systems for Internet use, it is necessary to specify a profile for use of the X.509 v3 extensions tailored for the Internet. It is one goal of PKIX to specify a profile for Internet, electronic mail, and IPsec applications, etc. Environments with additional requirements may build on this profile or may replace it.

### [3.4.2](#) Attribute Certificates

ANSI X.9 first published the Attribute Certificate format in [put date in here] as part of [put reference in here]. It defined the standard version 1 (v1) AC format. They later created a version 2 (v2) AC by modifying the owner field to point to either an identity or a specific PKC and including an extension mechanism. In 1997 ITU-T included it in [[X.509](#)].

ANSI, ITU-T, and IETF have developed standard extensions and attributes for use in the v2 ACs. Extensions can convey such information as an audit identity that can be used to create an audit trail, identity specific servers/services where the AC owner can use

their AC, point to a specific issuer's key, and indicate where to get revocation information. The AC is generic enough to allow any attribute to be conveyed in the data structure. Without limiting the attributes and extensions that can be included in an AC it is very difficult to develop interoperable implementations for Internet use. It is the goal of PKIX to specify a profile for the Internet, electronic mail, IPsec applications, etc. Environments with additional requirements may build on this profile or replace it.

### [3.5](#) Functions of a PKI

This section describes the major functions of a PKI. In some cases, PKIs may provide extra functions.

### [3.5.1](#) Registration

This is the process whereby a subject first makes itself known to a CA (directly, or through an RA), prior to that CA issuing a PKC or PKCs for that subject. Registration involves the subject providing its name (e.g., common name, fully-qualified domain name, IP address), and other attributes to be put in the PKC, followed by the CA (possibly with help from the RA) verifying in accordance with its Certification Practice Statement (CPS) that the name and other attributes are correct.

### [3.5.2](#) Initialization

Initialization is when the subject (e.g., the user or client system) gets the values needed to begin communicating with the PKI. For example, initialization can involve providing the client system with the public key and/or PKC of a CA, or generating the client system's own public/private key pair.

### [3.5.3](#) Certification

This is the process in which a CA issues a PKC for a subject's public key, and returns that PKC to the subject and/or posts that PKC in a repository.

### [3.5.4](#) Key Pair Recovery

In some implementations, key exchange or encryption keys will be required by local policy to be "backed up," or recoverable in case the key is lost and access to previously-encrypted information is needed. Such implementations can include those where the private key exchange key is stored on a hardware token which can be lost or broken, or when a private key file is protected by a password which can be forgotten. Often, a company is concerned about being able to

read mail encrypted by or for a particular employee when that employee is no longer available because she is ill or no longer works for the company.

In these cases, the user's private key can be backed up by a CA or by a separate key backup system. If a user or her employer needs to recover these backed up key materials, the PKI must provide a system that permits the recovery without providing an unacceptable risk of compromise of the private key.

#### [3.5.5](#) Key Generation

Depending on the CA's policy, the private/public key pair can either be generated by the user in his local environment, or generated by the CA. In the latter case, the key material may be distributed to the user in an encrypted file or on a physical token (e.g., a smart card or PCMCIA card).

#### [3.5.6](#) Key Update

All key pairs need to be updated regularly (i.e., replaced with a new key pair) and new PKCs issued. This will happen in two cases: normally, when a key has passed its maximum usable lifetime; and exceptionally, when a key has been compromised and must be replaced.

##### [3.5.6.1](#) Key Expiry

In the normal case, a PKI needs to provide a facility to gracefully transition from a PKC with an existing key to a new PKC with a new key. This is particularly true when the key to be updated is that of a CA. Users will know in advance that the key will expire on a certain date; the PKI, working together with PKC-using applications, should allow for appropriate keys to work before and after the transition. There are a number of ways to do this; see [[CMP](#)] for an example of one.

##### [3.5.6.2](#) Key Compromise

In the case of a key compromise, the transition will not be "graceful" in that there will be an unplanned switch of PKCs and keys; users will not have known in advance what was about to happen. Still, the PKI must support the ability to declare that the previous PKC is now invalid and shall not be used, and to announce the validity and availability of the new PKC.

Note: compromise of a private key associated with a Root CA is catastrophic for users relying on that Root CA. If a Root CA's private key is compromised, that CA's PKC must be revoked and all

PKCs subordinate to it must also be revoked. Until such time as the Root CA has been issued a new PKC and the Root CA issues PKCs to users relying upon it, users relying on that Root CA are cut off from the rest of the system, as there is no way to develop a valid certification path back to a trusted node.

Further, users will likely have to be notified by out-of-band mechanisms about the change in CA keys. If the old key is compromised, any "update" message telling subordinates to switch to a new key could have come from an attacker in possession of the old key, and could point to a new public key for which the attacker already has the private key. It is possible to have anticipated this event, and "pre-placed" replacement Root CA keys with all relying parties, but some secure, out-of-band mechanism will have to be used to tell users to make the switch, and this will only help if the replacement key has not been compromised.

Additionally, once the Root CA is brought back up with a new key, it will likely be necessary to re-issue PKCs, signed with the new key, to all subordinate users, since their current PKC would be signed with a now-revoked key.

#### [3.5.7](#) Cross-certification

A cross-certificate is a PKC issued by one CA to another CA which contains a public CA key associated with the private CA signature key used for issuing PKCs. Typically, a cross-certificate is used to allow client systems/end entities in one administrative domain to communicate securely with client systems/end users in another administrative domain. Use of a cross-certificate issued from CA\_1 to CA\_2 allows user Alice, who trusts CA\_1, to accept a PKC used by Bob, which was issued by CA\_2. Cross-certificates can also be issued from one CA to another CA in the same administrative domain, if required.

Cross-certificates can be issued in only one direction, or in both directions, between two CA's. That is, just because CA\_1 issues a cross-certificate for CA\_2, CA\_2 does not have to issue a cross-certificate for CA\_1.

#### [3.5.8](#) Revocation

When a PKC is issued, it is expected to be in use for its entire validity period. However, various circumstances may cause a PKC to become invalid prior to the expiration of the validity period. Such circumstances include change of name, change of association between subject and CA (e.g., an employee terminates employment with an organization), and compromise or suspected compromise of the

corresponding private key. Under such circumstances, the CA needs to

revoke the PKC.

X.509 defines one method of PKC revocation. This method involves each CA periodically issuing a signed data structure called a certificate revocation list (CRL). A CRL is a time stamped list identifying revoked PKCs which is signed by a CA and made freely available in a public repository. Each revoked PKC is identified in a CRL by its PKC serial number. When a certificate-using system uses a PKC, that system not only checks the PKC signature and validity but also acquires a suitably-recent CRL and checks that the PKC serial number is not on that CRL. The meaning of "suitably-recent" may vary with local policy, but it usually means the most recently-issued CRL. A CA issues a new CRL on a regular periodic basis (e.g., hourly, daily, or weekly). CA's may also issue CRLs aperiodically. For example, if an important key is deemed compromised, the CA may issue a new CRL to expedite notification of that fact, even if the next CRL does not have to be issued for some time. (A problem of aperiodic CRL issuance is that end-entities may not know that a new CRL has been issued, and thus may not retrieve it from a repository.)

An entry is added to the CRL as part of the next update following notification of revocation. An entry may be removed from the CRL after appearing on one regularly scheduled CRL issued beyond the revoked PKC's validity period. Leaving the revoked PKC on the CRL for this extra period allows for PKCs that are revoked prior to issuing a new CRL and whose invalidity date falls before the CRL issuing time to be accounted for. If the revoked PKC is not retained on the CRL for this extra period then the possibility arises that a revoked PKC may never appear on a CRL.

An advantage of the CRL revocation method is that CRLs may be distributed by exactly the same means as PKCs themselves, namely, via untrusted communications and server systems.

One limitation of the CRL revocation method, using untrusted communications and servers, is that the time granularity of revocation is limited to the CRL issue period. For example, if a revocation is reported now, that revocation will not be reliably notified to certificate-using systems until the next CRL is issued - this may be up to one hour, one day, or one week depending on the

frequency that the CA issues CRLs.

As with the X.509 v3 PKC format, in order to facilitate interoperable implementations from multiple vendors, the X.509 v2 CRL format needed to be profiled for Internet use. This was done as part of [\[FORMAT\]](#). However, PKIX does not require CAs to issue CRLs. On-line methods of revocation notification may be applicable in some environments as an alternative to the X.509 CRL. PKIX defines a protocol known as OCSP

to facilitate on-line checking of the status of PKCs [\[OCSP\]](#).

On-line revocation checking may significantly reduce the latency between a revocation report and the distribution of the information to relying parties. Once the CA accepts the report as authentic and valid, any query to the on-line service will correctly reflect the PKC validation impacts of the revocation. However, these methods impose new security requirements; the PKC validator must trust the on-line validation service while the repository does not need to be trusted.

#### [3.5.9](#) Certificate and Revocation Notice Distribution/Publication

As alluded to in sections [3.1](#) and [3.5.8](#) above, the PKI is responsible for the distribution of PKCs and PKC revocation notices (whether in CRL form or in some other form) in the system. "Distribution" of PKCs includes transmission of the PKC to its owner, and may also include publication of the PKC in a repository. "Distribution" of revocation notices may involve posting CRLs in a repository, transmitting them to end-entities, and/or forwarding them to on-line responders.

#### [3.6](#) Parts of PKIX

This section identifies the six different areas in which the PKIX working group has developed documents. The first area involves profiles of the X.509 v3 PKC standards and the X.509 v2 CRL standards for the Internet. The second area involves operational protocols, in which relying parties can obtain information such as PKCs or PKC status. The third area covers management protocols, in which different entities in the system exchange information needed for proper management of the PKI. The fourth area provides information about certificate policies and certificate practice statements, covering the areas of PKI security not directly addressed in the rest

of PKIX. The fifth area deals with providing time stamping and data certification services, which can be used to build such services as non-repudiation.

### [3.6.1](#) Profiles

An X.509 v3 PKC is a very complex data structure. It consists of basic information fields, plus a number of optional extensions. Many of the fields and numerous extensions can take on a wide range of options. This provides an enormous degree of flexibility, which allows the X.509 v3 PKC format to be used with a wide range of applications in a wide range of environments. Unfortunately, this same flexibility makes it extremely difficult to produce independent implementations that will actually interoperate with one another. In order to build an Internet PKI based on X.509 v3 PKCs, the PKIX

working group had to develop a profile of the X.509 v3 PKC specification.

A profile of the X.509 v3 PKC specification is a description of the contents of the PKC and which extensions must be supported, which extensions may be supported, and which extensions may not be supported. [\[FORMAT\]](#) provides such a profile of X.509 v3 PKC for the Internet PKI. In addition, [\[FORMAT\]](#) suggests ranges of values for many of the extensions.

[\[FORMAT\]](#) also provides a profile for Version 2 CRLs for use in the Internet PKI. CRLs, like PKCs, have a number of optional extensions. In order to promote interoperability, it is necessary to constrain the choices an implementor supports.

In addition to profiling the PKC and CRL formats, it is necessary to define particular Object Identifiers (OIDs) for certain encryption algorithms, because there are a variety of OIDs registered for some algorithm suites. Many of the OIDs are defined in [\[FORMAT\]](#) to promote interoperability. Also, PKIX has produced two documents ([\[ECDSA\]](#) and [\[KEA\]](#)) which provide guidance on the proper implementation of specific algorithms.

Some countries are in a process of updating their legal frameworks in order to regulate and incorporate recognition of signatures in electronic form. Many of these frameworks introduce certain basic

requirements on PKCs, often termed Qualified Certificates, supporting these types of "legal" signatures. Partly as a result of this there is a need for a specific PKC profile providing standardized support for certain related issues such as a common structure for expressing unambiguous identities of certified subjects (unmistakable identity). In December 1998, PKIX adopted as a work item the development of a refinement of [\[RFC2459\]](#) that further profiles PKIX PKC into qualified certificates. This work is reflected in [\[QC\]](#).

Like the X.509 v3 PKC the AC also a very complex data structure consisting of basic information fields, a number of optional extensions, and a virtually unlimited number of attributes. Again, many of the fields, extensions, and attributes can take on a wide range of options allowing an enomorous degree of flexibility. In order to build an Internet PMI based on ACs, the PKIC working group had to develop a profile of the AC.

The AC profile is description of the contents of the AC, the allowed and required extensions, and applicable attributes. [\[AC\]](#) provides such a profile of the X.509 v2 AC.

### [3.6.2](#) Operational Protocols

Operational protocols are required to deliver certificates and CRLs (or other certificate status information) to certificate using systems. Provision is needed for a variety of different means of certificate and CRL delivery, including distribution procedures based on LDAP, HTTP, FTP, and X.500. Operational protocols supporting these functions are defined in [\[FTPHTTP\]](#), [\[OCSP\]](#), [\[PKI-LDAPv2\]](#), and [\[PKI-LDAPv3\]](#). A limited protocol to support AC retrieval has also been document in [\[LAAP\]](#).

[\[DHPOP\]](#) defines a procedure for producing signatures withg the Diffie-Hellman key agreement algorithm. This signature mechanism was developed to support PKCS-10 certificate requests.

### [3.6.3](#) Management Protocols

Management protocols are required to support on-line interactions between PKI user and management entities. For example, a management

protocol might be used between a CA and a client system with which a key pair is associated, or between two CAs which cross-certify each other. A management protocol can be used to carry user or client system registration information, or a request for revocation of a certificate.

There are two parts to a "management protocol." The first is the format of the messages that will be sent, and the second is the actual protocol that governs the transmission of those messages. Originally, the PKIX working group developed two documents, [\[CRMF\]](#) and certificate management message format (CMMF), that together described the necessary set of message formats, and two other documents, [\[CMP\]](#) and [\[CMC\]](#), that described protocols for exchanging those messages. However, the message formats defined in the CMMF draft were inserted into both [\[CMP\]](#) and [\[CMC\]](#), and thus the (CMMF) draft has been dropped as a PKIX document.

[\[CMP-HTTP\]](#) and [\[CMP-TCP\]](#) were developed, after some implementation experience, to update the procedure documented in [\[CMP\]](#) for using CMP with HTTP and TCP and the transport protocols [\[CMP\]](#).

#### [3.6.4](#) Policy Outline

As mentioned before, profiling certificates and specifying operational and management protocols only addresses a part of the problem of actually developing and implementing a secure PKI. What is also needed is the development of a certificate policy (CP) and certification practice statement (CPS), and then following those documents. The CP and CPS should address physical and personnel

security, subject identification requirements, revocation policy, and a number of other topics. [\[POLPROC\]](#) provides a framework for certification practice statements.

#### [3.6.5](#) Time-Stamp and Data Certification Services

In late 1998, the PKIX working group began two efforts that were not in the original working group charter, but were deemed to be appropriate because they described infrastructure services that could be used to provide desired security services. The first of these is time stamping, described in [\[TSP\]](#). Time stamping is a service in which a trusted third party - a Time Stamp Authority, or TSA - signs

a message, in order to provide evidence that it existed prior to a given time. Time stamping provides some support for non-repudiation, in that a user cannot claim that a transaction was later forged after compromise of a private key, because the existence of the signed time stamp indicates that the transaction in question could not have been created after the indicated time.

[TSP] also defines the role of a Temporal Data Authority, or TDA. A TDA is a Trusted Third Party (TTP) that creates a temporal data token. This temporal data token associates a message with a particular event and provides supplementary evidence for the time included in the time stamp token. For example, a TDA could associate the message with the most recent closing value of the Dow Jones Average. The temporal data with which the message is associated should be unpredictable in order to prevent forward dating of tokens. The third iteration of the draft removed support for TDAs as no one in the WG expressed a requirement for the role.

At the Minneapolis IETF meeting, it was disclosed that the materials covered in [TSP] draft may be covered by patent(s). Use of the material covered by the patent(s) in question has not been granted by the patentholder. Thus, anyone interested in implementing the PKIX [TSP] draft must be aware of this intellectual property issue.

The second new effort is the definition of a Data Validation and Certification Server, or DVCS, protocol [DVCS]. A DVCS is a Trusted Third Party that verifies the correctness of specific data submitted to it.

This is different from the TSP service in that a TSA will not attempt to parse and/or verify a message sent to it for certification; instead, it will merely append a reliable indication of the current time, and sign the resulting string-of-bits. This offers an indication that the given string-of-bits existed at a specified time; it does not offer any indication of the correctness or relevance of that string of bits. By contrast, the DVCS certifies possession of

data or the validity of another entity's signature. As part of this, the DVCS verifies the mathematical correctness of the actual signature value contained in the request and also checks the full certification path from the signing entity to a trusted point (e.g., the DVCS's CA, or the Root CA in a hierarchy).

The DVCS supports non-repudiation in two ways. First, it provides evidence that a signature or PKC was valid at the time indicated in the token. The token can be used even after the corresponding PKC expires and its revocation information is no longer available on CRLs (for example). Second, the production of a data certification token in response to a signed request for certification of another signature or PKC also provides evidence that due diligence was performed by the requester in validating the signature or PKC.

## [4](#) PKIX Documents

This section describes each of the documents written by the PKIX working group. As PKIX progresses, this section will need to be continually updated to reflect the status of each document (e.g., Proposed Standard, Draft Standard, Standard, Informational Draft, Informational RFC, something-that-was-just-thrown-out-for-consideration, etc.).

### [4.1](#) Profile

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Certificate and CRL Profile ([RFC 2459](#))

DESCRIPTION: This document describes the profiles to be used for X.509 v3 PKCs and version 2 CRLs by Internet PKI participants. The profiles include the identification of ISO/IEC/ITU and ANSI extensions which may be useful in the Internet PKI. The profiles are presented in the 1988 Abstract Syntax Notation One (ASN.1) rather than the 1994 syntax used in the ISO/IEC/ITU standards. Would-be PKIX implementors and developers of certificate-using applications should start with [\[FORMAT\]](#) to ensure that their systems will be able to interoperate with other users of the PKI.

[\[FORMAT\]](#) also includes path validation procedures. The procedures presented are based upon the ISO/IEC/ITU definition, but the presentation assumes one or more self-signed trusted CA PKCs. The procedures are provided as examples only. Implementations are not required to use the procedures provided; they may implement whichever procedures are efficient for their situation. However, implementations are required to derive the same results as the example procedures.

STATUS: Proposed Standard.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Representation of Elliptic Curve Digital Signature Algorithm (ECDSA) Keys and Signatures in Internet X.509 Public Key Infrastructure Certificates <[draft-ietf-pkix-ipki-ecdsa-01.txt](#)>

DESCRIPTION: This document provides Object Identifiers (OIDs) and other guidance for IPKI users who use the Elliptic Curve Digital Signature Algorithm (ECDSA). It profiles the format and semantics of the subjectPublicKeyInfo field and the keyUsage extension in X.509 v3 PKCs containing ECDSA keys. This document should be used by anyone wishing to support ECDSA; others who do not support ECDSA are not required to comply with it.

STATUS: WG Last Call.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Representation of Key Exchange Algorithm (KEA) Keys in Internet X.509 Public Key Infrastructure Certificates ([RFC 2528](#))

DESCRIPTION: This document provides Object Identifiers (OIDs) and other guidance for IPKI users who use the Key Exchange Algorithm (KEA). It profiles the format and semantics of the subjectPublicKeyInfo field and the keyUsage extension in X.509 v3 PKCs containing KEA keys. This document should be used by anyone wishing to support KEA; others who do not support ECDSA are not required to comply with it.

STATUS: Informational RFC.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Qualified Certificates <[draft-ietf-pkix-qc-01.txt](#)>

DESCRIPTION: This document profiles the format for and defines requirements on information content in a specific type of PKCs called Qualified Certificates. A "Qualified Certificate" is a PKC that is issued to a natural person (i.e., a living human being); contains an unmistakable identity based on a real name or a pseudonym of the subject; exclusively indicates non-repudiation as the key usage for the certificate's public key; and meets a number of requirements.

STATUS: Under WG review.

DOCUMENT TITLE: An Internet AttributeCertificate Profile for Authorizations <[draft-ietf-pkix-acx509prof-01.txt](#)>

DESCRIPTION: This document profiles the format for an defines

INTERNET DRAFT

PKIX Roadmap

22 October 1999

requirements on X.509 v2 ACs to support authorization services required by various Internet protocols (TLS, CMS, and the consumers of CMS, etc.). Two profiles are defined on that supports basic authorizations and on the supports proxiable services.

STATUS: Under WG review.

#### [4.2](#) Operational Protocols

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Operational Protocols - LDAPv2 ([RFC 2559](#))

DESCRIPTION: This document describes the use of LDAPv2 as a protocol for PKI elements to publish and retrieve certificates and CRLs from a repository. [[LDAPv2](#)] is a protocol that allows publishing and retrieving of information.

STATUS: Proposed Standard.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure LDAPv2 Schema ([RFC 2587](#))

DESCRIPTION: This document defines a minimal schema necessary to support the use of LDAPv2 for PKC and CRL retrieval and related functions for PKIX. This document supplements [[LDAPv2](#)] by identifying the PKIX-related attributes that must be present.

STATUS: Proposed Standard.

DOCUMENT TITLE: X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP ([RFC 2560](#))

DESCRIPTION: This document specifies a protocol useful in determining the current status of a certificate without the use of CRLs. A major complaint about certificate-based systems is the need for a relying party to retrieve a current CRL as part of the certificate validation process. Depending on the size of the CRL, this can cause severe problems for bandwidth-challenged devices. Depending on the frequency of CRL issuance, this can also cause timeliness problems. (E.g., if CRLs are only published weekly, with no interim releases, a certificate could actually have been revoked for just short of one week without it being on the current CRL, and thus improper use of

that certificate could still be occurring.)

OCSP attempts to address those problems. It provides a mechanism whereby a relying party identifies one or more certificates to an approved OCSP "responder", and the responder sends back the current status of the certificate(s) - e.g., "revoked", "notRevoked",

"unknown". This can dramatically reduce the bandwidth required to transmit revocation status - a relying party does not have to retrieve a CRL of many entries to check the status of one certificate. It can (although it is not guaranteed to) improve the timeliness of revocation notification, and thus reduce the window of opportunity for someone trying to use a revoked certificate.

STATUS: Proposed Standard.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Operational Protocols: FTP and HTTP ([RFC 2585](#))

DESCRIPTION: This document describes the use of the File Transfer Protocol (FTP) and the Hyper-text Transfer Protocol (HTTP) to obtain certificates and CRLs from PKI repositories.

STATUS: Proposed Standard.

DOCUMENT TITLE: Diffie-Hellman Proof-of-Possession Algorithms <[draft-ietf-pkix-dhpop-02.txt](#)>

DESCRIPTION: This documents describes two signing algorithms using the Diffie-Hellman key agreement process to provide a shared secret as the basis of the signature. It allows Diffie-Hellman a key agreement algorithm to be used instead of requiring that the public key being requested for certification correspond to an algorithm that is capable of signing and/or encrypting. The first algorithm generates a signature for a specific verifier where the signer and recipient have the same public key parameters. The second algorithm generates a signature for arbitrary verifiers where the signer and recipient do not have the same public key parameters.

STATUS: Under WG review.

DOCUMENT TITLE: Limited Attribute Certificate Aquisition Protocol

[<draft-ietf-pkix-laap-00.txt>](#)

DESCRIPTION: This document specifies a deliberately limited protocol for requesting ACs from a server. It is intended to be complementary to the use of LDAP for AC retrieval, covering those cases where use of an LDAP server is not suitable due to the type of authorization model being employed.

STATUS: Under WG review.

### [4.3](#) Management Protocols

DOCUMENT TITLE: Certificate Management Messages over CMS [<draft-ietf-pkix-cmc-05.txt>](#)

DESCRIPTION: This document defines the means by which PKI clients and servers may exchange PKI messages when using S/MIME's Cryptographic Message Syntax [[CMS](#)] as a transaction envelope. CMC supports the certificate request message body specified in the Certificate Request Message Format [[CRMF](#)] documents, as well as a variety of other certificate management messages. The primary purpose of this specification is to allow the use of an existing protocol (S/MIME) as a PKI management protocol, without requiring the development of an entirely new protocol such as CMP. A secondary purpose is to codify in IETF standards the current industry practice of using PKCS-10 messages [[PKCS10](#)] for certificate requests.

STATUS: Under WG review.

DOCUMENT TITLE: Internet X.509 Certificate Request Message Format ([RFC 2511](#))

DESCRIPTION: CRMF specifies a format recommended for use whenever a relying party is requesting a certificate from a CA or requesting that an RA help it get a certificate. The request message format was needed before many of the other message formats had to be finalized, and so it was put into a separate document. This document only specifies the format of a message. Specification of a protocol to

transport that message is beyond the scope of CRMF.

STATUS: Proposed Standard.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Certificate Management Protocols ([RFC 2510](#))

DESCRIPTION: This document specifies a new protocol specifically developed for the purpose of transporting messages like those specified in CRMF among PKI elements. In general, CMP will be used in conjunction with CRMF, and will then be run over a transfer service (e.g., S/MIME, HTTP) to provide a complete PKI management service.

STATUS: Proposed Standard.

DOCUMENT TITLE: Simple Certificate Validation Protocol (SCVP) <[draft-ietf-pkix-scvp-01.txt](#)>

DESCRIPTION: The SCVP protocol allows a client to offload certificate handling to a server. The server can give a variety of valuable

information about the certificate, such as whether or not the certificate is valid, a chain to a trusted root, and so on.

STATUS: Under WG review.

DOCUMENT TITLE: Using HTTP as a Transport Protocol for CMP <[draft-ietf-pkix-cmp-http-00.txt](#)>

DESCRIPTION: This document describes how to layer [[CMP](#)] over [HTTP]. A simple method for doing so is described in [[CMP](#)], but that method does not accommodate a polling mechanism, which may be required in some environments. This document specifies an alternative method which uses the polling protocol defined in [[CMP](#)]. A new Content-Type for messages is also defined.

STATUS: Under WG review.

DOCUMENT TITLE: Using TCP as a Transport Protocol for CMP <[draft-ietf-pkix-cmp-tcp-00.txt](#)>

DESCRIPTION: This document describes how to layer Certificate

Management Protocols [[CMP](#)] over [TCP]. A method for doing so is described in [[CMP](#)], but that method does not solve problems encountered by implementors. This document specifies an enhanced method which extends the protocol.

STATUS: Under WG review.

DOCUMENT TITLE: OCSP Extensions <[draft-ietf-pkix-ocsp-x-00.txt](#)>

DESCRIPTION: This document defines Internet-standard extensions to OCSP that enable a client to delegate processing of certificate acceptance functions to a trusted server. The client may control the degree to which delegation takes place. In addition limited support is provided for delegating authorization decisions.

STATUS: Under WG review.

#### [4.4](#) Policy Outline

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework ([RFC 2527](#))

DESCRIPTION: As noted before, the specification and implementation of certificate profiles, operational protocols, and management protocols is only part of building a PKI. Equally as important - if not more important - is the development and enforcement of a certificate security policy, and a Certification Practice Statement (CPS). The

purpose of this document (PKIX-4) is to establish a clear relationship between certificate policies and CPSs, and to present a framework to assist the writers of certificate policies or CPSs with their tasks. In particular, the framework identifies the elements that may need to be considered in formulating a certificate policy or a CPS. The purpose is not to define particular certificate policies or CPSs, per se.

STATUS: Informational RFC.

#### [4.5](#) Time-Stamp and Data Certification Services

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Time Stamp Protocols <[draft-ietf-pkix-time-stamp-03.txt](#)>

DESCRIPTION: This document defines the specification for a time stamp service. It defines a Time Stamp Authority, or TSA, a trusted third party who maintains a reliable time service. When the TSA receives a time stamp request, it appends the current time to the request and signs it into a token to certify that the original request existed prior to the indicated time. This helps provide non-repudiation by preventing someone (either a legitimate user or an attacker who has successfully compromised a key) from "back-dating" a transaction. It also makes it more difficult to challenge a transaction by asserting that it has been back-dated. Note that the TSA does not provide any data parsing service; that is, the TSA does not attempt to validate that which it signs; it merely regards it as a string of bits whose meaning is unimportant, but existence is vital.

STATUS: Under WG review.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Data Certification Server Protocols <[draft-ietf-pkix-dcs-03.txt](#)>

DESCRIPTION: This document defines a data validation and certification service, or DVCS, which can be used to certify both the existence and correctness of a message or signature. In contrast to the time stamp service described above, the DVCS certifies possession of data or the validity of another entity's signature. As part of this, the DVCS verifies the mathematical correctness of the actual signature value contained in the request and also checks the full certification path from the signing entity to a trusted point (e.g., the DVCS's CA, or the Root CA in a hierarchy).

The DVCS supports non-repudiation in two ways. First, it provides evidence that a signature or public key certificate was valid at the time indicated in the token. The token can be used even after the corresponding public key certificate expires and its revocation

information is no longer available on CRLs (for example). Second, the production of a data certification token in response to a signed request for certification of another signature or public key certificate also provides evidence that due diligence was performed by the requester in validating the signature or public key certificate.

STATUS: Under WG review.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Extending trust in non repudiation tokens in time <[draft-ietf-pkix-extend-trust-non-repudiation-token-00.txt](#)>

DESCRIPTION: This document describes a method to maintain the trust in a token issued by a non-repudiation Trusted Third Party (NR TTP) (DVCS/TSA/TDA) after the key initially used to establish trust in the token expires. The document describes a general format for storage of DVCS/TS/TDA tokens for this purpose, which establishes a chain of custody for the data.

STATUS: Under WG review.

#### [4.6](#) Documents that never made it out of the working group

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Certificate Management Message Formats

DESCRIPTION: This document contains the formats for a series of messages important in certificate/PKI management. These messages let CA's, RA's, and relying parties communicate with each other. Note that this document only specifies message formats; it does not specify a protocol for transferring messages. That protocol can be either CMP or CMC, or perhaps another custom protocol.

STATUS: Work has been discontinued, as all useful information from it has been moved into [[CMP](#)] and [[CMC](#)].

DOCUMENT TITLE: WEB based Certificate Access Protocol-- WebCAP/1.0

DESCRIPTION: This document specifies a set of methods, headers, and content-types ancillary to HTTP/1.1 to publish, retrieve X.509 certificates and Certificate Revocation Lists. This protocol also facilitates determining current status of a digital certificate without the use of CRLs. This protocol defines new methods, request and response bodies, error codes to HTTP/1.1 protocol for securely publishing, retrieving, and validating certificates across a firewall.

STATUS: Work has been discontinued due to lack of interest.

DOCUMENT TITLE: Internet X.509 Public Key Infrastructure Enhanced CRL Distribution Options (OpenCDP)

DESCRIPTION: This document proposes an alternative to the CRL Distribution Point (CDP) approach documented in [[FORMAT](#)]. OCDP separates the CRL location function from the process of certificate and CRL validation, and thus claims some benefits over the CDP approach.

STATUS: Work has been discontinued, as all useful information has been incorporated into [[X.509](#)]. An updated [[FORMAT](#)] RFC should profile the use of the CDP approach.

DOCUMENT TITLE: Internet Public Key Infrastructure: Caching the Online Certificate Status Protocol

DESCRIPTION: To improve the degree to which it can scale, OCSP allows caching of responses - e.g., at intermediary servers, or even at the relying party's end system. This document describes how to support OCSP caching at intermediary servers.

STATUS: Work has been discontinued due to lack of interest.

DOCUMENT TITLE: Basic Event Representation Token <[draft-ietf-pkix-bert1-01.txt](#)>

DESCRIPTION: This document defines a finite method of representing a discrete instant in time as a referable event. The Basic Event Representation Token (BERT) is a light-weight binary token designed for use in large numbers over short periods of time. The tokens contain only a single instance of an event stamp and the trust process is referenced against an external reference.

STATUS: Work has been discontinued.

## [5](#) Advice to Implementors

This section provides guidance to those who would implement various parts of the PKIX suite of documents. The topics discussed in this section engendered significant discussion in the working group, and there was at times either widespread disagreement or widespread misunderstanding of them. Thus, this discussion is provided to help readers of the PKIX document set understand these issues, in the hope of fostering greater interoperability among eventual PKIX implementations.

INTERNET DRAFT

PKIX Roadmap

22 October 1999

## [5.1](#) Names

PKIX has been referred to as a "name-centric" PKI because the PKCs associate public keys with names of entities. Each PKC contains at least one name for the owner of a particular public key. The name can be an X.500 distinguished name, contained in the subjectDN field of the PKC. There can also be names such as [RFC822](#) e-mail addresses, DNS domain names, and uniform resource identifiers (URIs) associated with the key; these attributes are kept in the subjectAltName extension of the PKC. A PKC must contain at least one of these name forms, it may contain multiple forms if deemed appropriate by the CA based on the intended usage of the PKC.

### [5.1.1](#) Name Forms

There are two possible places to put a name in an X.509 v3 PKC. One is the subject field in the base PKC (often called the "Distinguished Name" or "DN" field), and the other is in the subjectAltName extension.

#### [5.1.1.1](#) Distinguished Names

According to [\[FORMAT\]](#), a PKIX PKC must have a non-null value in the subject field, except for an EE PKC, which is permitted to have an empty subject field. Furthermore, if a PKC has a non-null subject field, it MUST contain an X.500 Distinguished Name.

#### [5.1.1.2](#) SubjectAltName Forms

In addition to the DN, a PKIX PKC may have one or more values in the subjectAltName extension.

The subjectAltName extension allows additional identities to be bound to the subject of the PKC (e.g., it allows "umbc.edu" and "130.85.1.3" to be associated with a particular subject, as well as "C=US, O=University of Maryland, L=Baltimore, CN=UMBC"). X.509-defined options for this extension include: Internet electronic mail addresses; DNS names; IP addresses; and URIs. Other options can exist, including locally-defined name forms.

A single subjectAltName extension can include multiple name forms, and multiple instances of each name form.

Whenever such alternate name forms are to be bound into a PKC, the `subjectAltName` (or `issuerAltName`) extension must be used. It is technically possible to embed an alternate name form in the subject field. For example, one could make a DN contain an IP address via a kludge such as "C=US, L=Baltimore, CN=130.85.1.3". However, this

usage is deprecated; the alternative name extension is the preferred location for finding such information. As another example, some legacy implementations exist where an [RFC822](#) name is embedded in the subject distinguished name as an `EmailAddress` attribute. Per [\[FORMAT\]](#), PKIX-compliant implementations generating new PKCs with electronic mail addresses MUST use the `rfc822Name` in the `subjectAltName` extension to describe such EEs. Simultaneous inclusion of the `EmailAddress` attribute in the subject distinguished name to support legacy implementation is deprecated but permitted.

In line with this, if the only subject identity included in a PKC is an alternative name form, then the subject distinguished name must be empty (technically, an empty sequence), and the `subjectAltName` extension must be present. If the subject field contains an empty sequence, the `subjectAltName` extension must be marked critical.

If the `subjectAltName` extension is present, the sequence must contain at least one entry. Unlike the subject field, conforming CAs shall not issue PKCs with `subjectAltNames` containing empty `GeneralName` fields. For example, an `rfc822Name` is represented as an `IA5String`. While an empty string is a valid `IA5String`, such an `rfc822Name` is not permitted by PKIX. The behavior of clients that encounter such a PKC when processing a certification path is not defined by this working group.

Because the subject's alternative name is considered to be definitively bound to the public key, all parts of the subject's alternative name must be verified by the CA.

#### [5.1.1.2.1](#) Internet e-mail addresses

When the `subjectAltName` extension contains an Internet mail address, the address is included as an `rfc822Name`. The format of an `rfc822Name` is an "addr-spec" as defined in [\[RFC-822\]](#). An addr-spec has the form `local-part@domain`; it does not have a phrase (such as a common name) before it, or a comment (text surrounded in parentheses) after it,

and it is not surrounded by "<" and ">".

#### [5.1.1.2.2](#) DNS Names

When the subjectAltName extension contains a domain name service label, the domain name is stored in the dNSName attribute (an IA5String). The string shall be in the "preferred name syntax," as specified by [\[DNS\]](#). Note that while upper and lower case letters are allowed in domain names, no significance is attached to the case. In addition, while the string " " is a legal domain name, subjectAltName extensions with a dNSName " " are not permitted. Finally, the use of the DNS representation for Internet mail addresses (wpolk.nist.gov

instead of wpolk@nist.gov) is not permitted; such identities are to be encoded as rfc822Name.

#### [5.1.1.2.3](#) IP addresses

When the subjectAltName extension contains an iPAddress, the address shall be stored in the octet string in "network byte order," as specified in [\[IP\]](#). The least significant bit (LSB) of each octet is the LSB of the corresponding byte in the network address. For IP Version 4, as specified in [\[IP\]](#), the octet string must contain exactly four octets. For IP Version 6, as specified in [\[IPv6\]](#), the octet string must contain exactly sixteen octets.

#### [5.1.1.2.4](#) URIs

[FORMAT] notes "When the subjectAltName extension contains a URI, the name MUST be stored in the uniformResourceIdentifier (an IA5String). The name MUST be a non-relative URL, and MUST follow the URL syntax and encoding rules specified in [\[RFC 1738\]](#). The name must include both a scheme (e.g., "http" or "ftp") and a scheme-specific-part. The scheme-specific-part must include a fully qualified domain name or IP address as the host. As specified in [\[RFC 1738\]](#), the scheme name is not case-sensitive (e.g., "http" is equivalent to "HTTP"). The host part is also not case-sensitive, but other components of the scheme-specific-part may be case-sensitive. When comparing URIs, conforming implementations MUST compare the scheme and host without regard to case, but assume the remainder of the scheme-specific-part is case sensitive."

### [5.1.2](#) Scope of Names

The original X.500 work presumed that every subject in the world would have a globally-unique distinguished name. Thus, every subject could be easily located, and there would never be a conflict. All that would be needed is a sufficiently-large name space, and rules for constructing names based on subordination and location.

Obviously, that is not practical in the real world, for a variety of reasons. There is no single entity in the world trusted by everybody to reside at the top of the name space, and there is no way to enforce uniqueness on names for all entities. (These were among the reasons for the failure of PEM to be widely implemented.)

This does not mean, however, that a name-based PKI cannot work. It is important to recognize that the scope of names in PKIX is local; they need to be defined and unique only within their own domain.

Suppose for example that a Top CA is established with DN "O=IETF,

OU=PKIX, CN=PKIX\_CA". That CA will then issue PKCs for subjects subordinate to it. The only requirement, which can be enforced procedurally, is that no two distinct entities beneath this Top CA have the same name. We can't prevent somebody else in the world from creating her own CA, called "O=IETF, OU=PKIX, CN=PKIX\_CA", and it is not necessary to do so. Within the domain of the original Top CA, there will be no conflict, and the fact that there is another CA of the same name in some other domain is irrelevant.

This is analogous to the current DNS or IP address situations. On the Internet, there is only one node called `www.ietf.org`. The fact that there might be 10 different intranets, each with a host given the DNS name `www.ietf.org`, is irrelevant and causes no interoperability problems - those are different domains. However, if there were to be another node on the Internet with domain name `www.ietf.org`, then there would be a problem due to name duplication.

The same applies for IP addresses. As long as only one node on the Internet responds to the IP address 130.85.1.3, there is no problem, despite the fact that there are 100 different corporate Intranets, each using that same IP address. However, if two different nodes on the Internet each responded to the IP address 130.85.1.3, there would

be a problem.

### [5.1.3](#) Certificate Path Construction

Certificate path construction has been the topic of many discussions in the WG. The issue centered around how best to get a certificate when the connection between the subject's name and the name of their CA, as well as the CA's repository (or directory), may not be obvious. Many proposals were put forth, including implementing a global X.500 Directory Service, using DNS SRV records, and using an extension to point to the directory provider. At the end of the discussion the group decided to use the authority information access extension defined in [\[FORMAT\]](#), which allows the CA to indicate the access method and location of CA information and services. The extension would be included in subject's certificates and could be used to associate an Internet style identity for the location of repository to retrieve the issuer's certificate in cases where such a location is not related to the issuer's name.

Another discussion related to certificate path construction was where to store the CA and EE PKCs in the directory (specifically LDAPv2 directories). Two camps emerged with different views on where to store CA and cross-certificates. In the CA's directory entry, one camp wanted self-issued PKCs stored in the cACertificate attribute, PKCs issued to this CA stored in the forward element of the crossCertificatePair, and PKCs issued from this CA for other CAs in

the reverse element of the crossCertificatePair attribute. The other camp wanted all CA PKCs stored in the cACertificate attribute, and PKCs issued to/from another domain stored in the crossCertificatePair attribute. There was a short discussion that the second was more efficient than first, and that one solution or the other was more widely deployed. The end result was to indicate that self-issued PKCs and PKCs issued to the CA by CAs in the same domain as the CA are stored in the cACertificate attribute. The crossCertificatePair attribute's forward element will include all but self-issued PKCs issued to the CA. The reverse element may include a subset of PKCs issued by the CA to other CAs. With this resolution both camp's implementations are supported and are free to choose the location of CA PKCs to best support their implementation.

### [5.1.4](#) Name Constraints

A question that has arisen a number of times is "how does one enforce Name constraints when there is more than one name form in a PKC?" That is, [\[FORMAT\]](#) states that:

subject's alternative names may be constrained in the same manner as subject distinguished names using the name constraints extension as described in [section 4.2.1.11](#).

What does this mean? Suppose that there is a CA with DN "O=IETF, OU=PKIX, CN=PKIX\_CA", with the subjectAltName extension having dNSName "PKIX\_CA.IETF.ORG". Suppose that that CA has issued a PKC with an empty DN, with subjectAltName extension having dNSName set to "PKIX\_CA.IETF.ORG", and rfc822Name set to Steve@PKIX\_CA.IETF.ORG. The question is, are name constraints enforced on these two PKCs - is the name of the EE PKC considered to be properly subordinate to the name of the CA?

The answer is "yes". The general rules for deciding whether a PKC meets name constraints are:

- If a PKC complies with name constraints in any one of its name forms, then the PKC is deemed to comply with name constraints.
- If a PKC contains a name form that its issuer does not, the PKC is deemed to comply with name constraints for that name form.

In deciding whether a name form meets name constraints, the following rules apply (in all cases Name B is the name in the name constraints extension):

#### [5.1.4.1](#) rfc822Names

Three variations are allowed:

- If the name constraint is specified as "larry@mail.mit.edu", then Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name (specifies a particular mailbox). For example, larry@mail.mit.edu is subordinate, but

`larry_sanders@mail.mit.edu` is not.

- If the name constraint is specified as `"mail.mit.edu"`, then Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name (specified all mailboxes on host `mail.mit.edu`). For example, `curly@mail.mit.edu` and `mo@mail.mit.edu` are subordinate, but `mo@mail6.mit.edu` and `curly@smtp.mail.mit.edu` are not.
- If the name constraint is specified as `".mit.edu"`, then Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name, with the addition of zero or more additional host or domain names to the left of Name B's name. That is, `smtp.mit.edu` is subordinate to `.mit.edu`, as is `pop.mit.edu`. However, `mit.edu` is not subordinate to `.mit.edu`. When the constraint begins with a `"."` it specifies any address within a domain. In the previous example, `"mit.edu"` is not in the domain of `".mit.edu"`.

Note: If [rfc822](#) names are constrained, but the PKC does not contain a `subjectAltName` extension, the `EmailAddress` attribute will be used to constrain the name in the subject distinguished name. For example (c is country, o is organization, ou is organizational unit, and em is `EmailAddress`), Name A (`"c=US, o=MIT, ou=CS, em=curly@mail.mit.edu"`) is subordinate to Name B (`"c=US, o=MIT, ou=CS"`) (in B's subtree) if Name A contains all of Name B's names.

#### [5.1.4.2](#) dNSNames

Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name, with the addition of zero or more additional domain names to the left of Name B's name. That is, `www.mit.edu` is subordinate to `mit.edu`, as is `larry.cs.mit.edu`. However, `www.mit.edu` is not subordinate to `web.mit.edu`.

#### [5.1.4.3](#) x.400 addresses

Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name. For example (c is country-name, admd is administrative-domain-name, and o is organization-name, ou is

given-name in both Name A and B), the mnemonic X.400 address (using PrintableString choices for c and admd) "c=US, admd=AT&T, o=MIT, ou=cs, pn='sn=Doe,gn=John'" is subordinate to "c=US, admd=AT&T, o=MIT, ou=cs" and "c=US, admd=AT&T, o=MIT, pn='sn=DOE,gn=JOHN'" (pn is a SET that includes, among other things, sn and gn).

#### [5.1.4.5](#) DNS

Name A is subordinate to Name B (in B's subtree), if Name A contains all of Name B's name, treating attribute values encoded in different types as different strings, ignoring case in PrintableString (in all other types case is not ignored), removing leading and trailing white space in PrintableString, and converting internal substrings of one or more consecutive white space characters to a single space. For example, (c is country, o is organization, ou is organizational unit, and cn is common name):

- Assuming PrintableString types for all attribute values in Name A and B, "c=US, o=MIT, ou=CS" is subordinate to "c=us, o=MIT, ou=cs", as is "c=US, o=MIT, ou=CS, ou=administration". Another example, "c=US, o=MIT, ou=CS, cn= John Doe" (note the white spaces) is subordinate to "c=US, o=MIT, ou=CS, cn=John Doe".
- Assuming UTF8String types for all attribute values in Name A and B, "c=US, o=MIT, ou=CS, ou=administration" is subordinate to "c=US, o=MIT, ou=CS", but "c=US, o=MIT, ou=cs, ou=Administration". "c=US, o=MIT, ou=CS, cn= John Smith" (note the white spaces) is not subordinate to "c=US, o=MIT, ou=CS, cn= John Smith".
- Assuming UTF8String types for all attribute values in Name A and PrintableString types for all attribute values in Name B, "c=US, o=MIT, ou=cs" is subordinate to "c=US, o=MIT, ou=CS" if the verification software supports the full comparison algorithm in the X.500 series. "c=US, o=MIT, ou=cs" is NOT subordinate to "c=US, o=MIT, ou=CS" if the verification software supports the comparison algorithm in [\[FORMAT\]](#).

#### [5.1.4.6](#) URIs

The constraints apply only to the host part of the name. Two variations are allowed:

- If the name constraint is specified as ".mit.edu", then Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name, with the addition of zero or more additional host or domain names to the left of Name B's name. That is, www.mit.edu is subordinate to .mit.edu, as is curly.cs.mit.edu.

However, mit.edu is not subordinate to .mit.edu. When the constraint begins with a "." it specifies a host.

- If the name constraint is specified as "abc.mit.edu", then Name A is subordinate to Name B (in B's subtree) if Name A contains all of Name B's name. No leftward expansion of the host or domain name is allowed.

#### [5.1.4.7](#) iPAddresses

Two variations are allowed depending on the IP version:

- For IPv4 addresses: Name A (128.32.1.1 encoded as 80 20 01 01) is subordinate to Name B (128.32.1.0/255.255.255.0 encoded as 80 20 00 00 FF FF FF 00) (in B's subtree) if Name A falls within the net denoted in Name B's CIDR notation.
- For IPv6 addresses: Name A (4224.0.0.0.8.2048.8204.16762 encoded as 10 80 00 00 00 00 00 00 00 00 08 08 00 20 0C 41 7A) is subordinate to Name B (4224.0.0.0.8.2048.8204.0/65535.65535.65535.65535.65535.65535.65535.0 encoded as 10 80 00 00 00 00 00 00 00 00 08 08 00 20 0C 00 00 FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF 00 00) (in B's subtree) if Name A falls within the net denoted in Name B's CIDR notation.

#### [5.1.4.8](#) Others

As [FORMAT] does not define requirements for the name forms otherName, ediPartyName, or registeredID there are no corresponding name constraints requirements.

#### [5.1.5](#) Wildcards in Name Forms

A "wildcard" in a name form is a placeholder for a set of names (e.g., "\*.mit.edu" meaning "any domain name ending in .mit.edu", and \*@aol.com meaning "email address that uses aol.com"). There are many people who believe that allowing wildcards in name forms in PKIX PKCs would be a useful thing to do, because it would allow a single PKC to be used by all members of a group. These members would presumably have attributes in common; e.g., access rights to some set of resources, and so issuance of a PKC with a wildcard for the group would simplify management.

After much discussion, the PKIX working group decided to permit the use of wildcards in PKCs. That is, it is permissible for a PKIX-conformant CA to issue a PKC with a wildcard. However, the semantics

of subjectAltName extension that include wildcard characters are not addressed by PKIX. Applications with specific requirements may use

such names but must define the semantics.

#### [5.1.6](#) Name Encoding

A very important topic that consumed much of the WG's time was the implementation of the directory string choices. While the long term goal of the IETF was clear, use UTF8String, the short term goals were not so clear. Many implementations only use PrintableString, others use BMPString, and still others use Latin1String (ISO 8859-1) and tag it as TeletexString (there are others still). To ensure that there is consistency with encodings [\[FORMAT\]](#) defines a set of rules for the string choices. PrintableString was kept as the first choice because of it's widespread support by vendors. BMPString was the second choice, also for it's widespread vendor support. Failing support by PrintableString and BMPString, UTF8String must be used. In keeping with the IETF mandate, UTF8String can be used at any time if the CA supports it. Also, processing implementations that wish to support TeletexString should handle the entire ISO 8859-1 character set and not just the Latin1String subset.

### [5.2](#) POP

Proof of Possession, or POP, means that the CA is adequately convinced that the entity requesting a PKC containing a public key Y has access to the private key X corresponding to that public key.

POP is important because it provides an appropriate level of assurance in the correct operation of the PKI as a whole. At its lowest level, POP counters the "self-inflicted denial of service"; that is, an entity voluntarily getting a PKC that cannot be used to sign or encrypt/decrypt information. However, as the following two examples demonstrate, POP also counters less direct, but more severe, threats.

#### [5.2.1](#) POP for Signing Keys

It is important to provide POP for keys used to sign material, in order to provide non-repudiation of transactions. For example, suppose Alice legitimately has private key X and its corresponding

public key Y. Alice has a PKC from Charlie, a CA, containing Y. Alice uses X to sign a transaction T. Without POP, Mal could also get a PKC from Charlie containing the same public key, Y. Now, there are two possible threats: Mal could claim to have been the real signer of T; or Alice can falsely deny signing T, claiming that it was instead Mal. Since no one can reliably prove that Mal did or did not ever possess X, neither of these claims can be refuted, and thus the service provided by and the confidence in the PKI has been defeated. (Of course, if Mal really did possess X, Alice's private key, then no

POP mechanism in the world will help, but that is a different problem.)

One level of protection can be gained by having Alice, as the true signer of the transaction, include in the signed information her PKC or an identifier of her PKC (e.g., a hash of her PKC). This might make it more difficult for Mal to claim authorship - he would have to assert that he incorrectly included Alice's PKC, rather than his own. However, it would not stop Alice from falsely repudiating her actions. Since the PKC itself is a public item, Mal indeed could have inserted Alice's PKC into the signed transaction, and thus its presence does not indicate that Alice was the one who participated in the now-repudiated transaction. The only reliable way to stop this attack is to require that Mal prove he possesses X before his PKC is issued.

For signing keys used only for authentication, and not for non-repudiation, the threat is lower because users may not care about Alice's after-the-fact repudiation, and thus POP becomes less important. However, POP SHOULD still be done wherever feasible in this environment, by either off-line or on-line means.

#### [5.2.2](#) POP for Key Management Keys

Similarly, POP for key management keys (that is, keys used for either key agreement or key exchange) can help to prevent undermining confidence in the PKI. Suppose that Al is a new instructor in the Computer Science Department of a local University. Al has created a draft final exam for the Introduction to Networking course he is teaching. He wants to send a copy of the draft final to Dorothy, the Department Head, for her review prior to giving the exam. This exam will of course be encrypted, as several students have access to the

computer system. However, a quick search of the PKC repository (e.g., search the repository for all records with subjectPublicKey=Dorothy's-value) turns up the fact that several students have PKCs containing the same public key management key as Dorothy. At this point, if no POP has been done by the CA, Al has no way of knowing whether all of the students have simply created these PKCs without knowing the corresponding private key (and thus it is safe to send the encrypted exam to Dorothy), or whether the students have somehow acquired Dorothy's private key (and thus it is certainly not safe to send the exam). Thus, the service to be provided by the PKI - allowing users to communicate with one another, with confidence in who they are communicating with - has been totally defeated. If the CA is providing POP, then either no students will have such PKCs, or Al can know with certainty that the students do indeed know Dorothy's private key, and act accordingly.

CMP requires that POP be provided for all keys, either through on-line or out-of-band means. There are any number of ways to provide POP, and the choice of which to use is a local matter. Additionally, a PKC requester can provide POP to either a CA or to an RA, if the RA can adequately assure the CA that POP has been performed. Some of the acceptable ways to provide POP include:

- Out-of-band means:
  - For keys generated by the CA or an RA (e.g., on a token such as a smart card or PCMCIA card), possession of the token can provide adequate proof of possession of the private key.
  - For user-generated keys, another approach can be used in environments where "key recovery" requirements force the requester to provide a copy of the private key to the CA or an RA. In this case, the CA will not issue the requested PKC until such time as the requester has provided the private key. This approach is in general not recommended, as it is extremely risky (e.g., the system designers need to figure out a way to protect the private keys from compromise while they are being sent to the CA/RA/other authority, and how to protect them there), but it can be used.
- On-line means:

- For signature keys that are generated by an EE, the requester of a PKC can be required to sign some piece of data (typically, the PKC request itself) using the private key. The CA will then use the requested public key to verify the signature. If the signature verification works, the CA can safely conclude that the requester had access to the private key. If the signature verification process fails, the CA can conclude that the requester did not have access to the correct private key, and reject the request.
- For key management keys that are generated by the requester, the CA can send the user the requested public key, along with the CA's own private key, to encrypt some piece of data, and send it to the requester to be decrypted. For example, the CA can generate some random challenge, and require some action to be taken after decryption (e.g., "decrypt this encrypted random number and send it back to me"). If the requester does not take the requested action, the CA concludes that the requester did not possess the private key, and the PKC is not issued.

Another method of providing POP for key management keys is for the CA to generate the requested PKC, and then send it to the requester in

encrypted form. If the requester does not have access to the appropriate private key, the requester cannot decrypt the PKC, and thus cannot use it. After some period of time in which the PKC is not used, the CA will revoke the PKC. (This only works if the PKC is not made available to any untrusted entities until after the requester has successfully decrypted it.)

### [5.3](#) Key Usage Bits

#### [5.3.1](#) Key Usage Extension

The key usage extension defines the purpose (e.g., encipherment, signature, certificate signing) of the key contained in the PKC. This extension is used when a key that could be used for more than one operation is to be restricted. For example, when an RSA key should be used only for signing, the `digitalSignature` and/or `nonRepudiation` bits would be asserted. Likewise, when an RSA key should be used only for key management, the `keyEncipherment` bit would be asserted. When

used, this extension should be marked critical.

According to [[FORMAT](#)], bits in the KeyUsage type are used as follows:

- The digitalSignature bit is asserted when the subject public key is used to verify digital signatures that have purposes other than non-repudiation, certificate signature, and CRL signature. For example, the digitalSignature bit is asserted when the subject public key is used to provide authentication via the signing of ephemeral data.
- The nonRepudiation bit is asserted when the subject public key is used to verify digital signatures used to provide a non-repudiation service which protects against the signing entity falsely denying some action, excluding certificate or CRL signing.
- The keyEncipherment bit is asserted when the subject public key is used for key transport. For example, when an RSA key is to be used for key management, this bit must asserted.
- The dataEncipherment bit is asserted when the subject public key is used for enciphering user data, other than cryptographic keys.
- The keyAgreement bit is asserted when the subject public key is used for key agreement. For example, when a Diffie-Hellman key is to be used for key management, this bit must asserted.
- The keyCertSign bit is asserted when the subject public key is used for verifying a signature on certificates. This bit may only

be asserted in CA certificates.

- The cRLSign bit is asserted when the subject public key is used for verifying a signature on revocation information (e.g., a CRL).
- The meaning of the encipherOnly bit is undefined in the absence of the keyAgreement bit. When the encipherOnly bit is asserted and the keyAgreement bit is also set, the subject public key may be used only for enciphering data while performing key agreement.

- The meaning of the decipherOnly bit is undefined in the absence of the keyAgreement bit. When the decipherOnly bit is asserted and the keyAgreement bit is also set, the subject public key may be used only for deciphering data while performing key agreement.

PKIX does not restrict the combinations of bits that may be set in an instantiation of the keyUsage extension.

The discussion on the PKIX mailing list has centered on the digitalSignature bit and the nonRepudiation bit. The question has come down to something like: When support for the service of non-repudiation is desired, should both the digitalSignature and nonRepudiation bits be set, or just the nonRepudiation bit?

Note: Provision of the service of non-repudiation requires more than a single bit set in a PKC. It requires an entire infrastructure of components to preserve for some period of time the keys, PKCs, revocation status, signed material, etc., as well as a trusted source of time. However, the nonRepudiation key usage bit is provided as an indicator that such keys should not be used as a component of a system providing a non-repudiation service.

According to [[SIMONETTI](#)], the intent is that the digitalSignature bit should be set when what is desired is the ability to sign ephemeral transactions; e.g., for a single session authentication. These transactions are "ephemeral" in the sense that they are important only while they are in existence; after the session is terminated, there is no long-term record of the digital signature and its properties kept. When something is intended to be kept for some period of time, the nonRepudiation bit should be set. This generally implies that an application will digitally sign something; thus, some implementors turn on the digitalSignature bit as well. Other implementors, however, keep the two bits mutually exclusive, to prevent a single key from being used for both ephemeral and long-term signing.

While [[FORMAT](#)] is silent on this specific issue, the working group's

general conclusion is that a PKC may have either or both bits set. If only the nonRepudiation bit is set, the key should not be used for ephemeral transactions. If only the digitalSignature bit is set, the key should not be used for long-term signing. If both bits are set,

the key may be used for either purpose.

To actually enforce this requires that an application understands whether it is signing ephemeral transactions or for the long-term. The application developers will have to understand the difference, and set up their checks on the key usage bits field accordingly. For example, TLS implementors should check only the digitalSignature bit, and ignore the nonRepudiation bit. S/MIME implementors, though, will have a difficult choice to make, since their application could be used for either purpose, and they will generally accept signing using keys associated with PKCs having either or both bits being turned on. Certification Authorities should know what applications they are providing PKCs for, and provide PKCs according to the requirements of those applications. If CA's are tied into non-repudiation systems, they may treat PKCs differently when the nonRepudiation bit is turned on (e.g., store information associated with the PKC - like the user's identification provided during PKC registration, or PKC generation date/time stamps - for longer periods of time, in more secure environments).

The bottom line is that this is an area where PKI implementors are somewhat limited in what they can do. The onus is on developers of PKC-using systems to understand their requirements and request PKCs with the appropriate bits set.

#### [5.3.1](#) Extended Key Usage Extension

[Add in text to talk about the extended key usages!]

#### [5.4](#) Trust Models

An important design decision is where in the PKI the trust point for a particular EE should be located (i.e., where is the EE's Root CA) . There are two extremes: the Top CA or the CA who issues the EE's certificate. Of course, the trust point for a particular EE can be anywhere in the PKI, but the following presents the advantages and disadvantages of locating the the trust point at these two places.

Advantages of Top CA trust point:

- Path discovery is easier since all EEs trust the same CA.
- Certificate paths are potentially shorter between distant EEs, since the verifier need only trace back to the root, not back to

his local CA.

- Root can enforce adherence to a certificate policy by subordinate CAs.
- Cross certification with other PKIs can be controlled at a senior level.

Disadvantages of root trust point:

- Compromise of the root key is catastrophic, requiring a re-distribution of certificates to all EEs. Similarly trust point roll-over affects entire hierarchy.
- Users are required to trust a CA which may be remote from them.
- Distribution of the trusted point certificate to distant EEs may be non-trivial.
- Verification back to the root CA may be too onerous for low value transactions.
- Certificate paths are potentially longer for nearby EEs since the verifier must always trace back to the root, not back to the CA it shares with the other party.

Advantages of local trusted point:

- The trusted point certificate need only be distributed from the CA to its local (nearby) EEs.
- EEs are more likely to trust their local CA (which could be part of the same immediate organization) than a geographically remote CA.
- Compromise of the local CAs private key only affects its own EEs. Similarly for trusted point roll-over.
- Potentially shorter certification paths between nearby EEs, since the verifier may belong to the same CA as the other party.

Disadvantages of local trust point:

- Potentially longer certification paths between distant EEs, since the verifier must trace the path back to its local CA.

INTERNET DRAFT

PKIX Roadmap

22 October 1999

- Path discovery more complex and may not be supported in current products.
- More difficult for the root to control cross-certification or ensure adherence to the certificate policy.

## 6 Acknowledgements

A lot of the information in this document was taken from the PKIX source documents; the authors of those deserve the credit for their own words. Other good material was taken from mail posted to the PKIX working group mail list ([ietf-pkix@imc.org](mailto:ietf-pkix@imc.org)). Among those with good things to say were (in alphabetical order, with apologies to anybody we've missed): Sharon Boeyen, Santosh Chokhani, Warwick Ford, Russ Housley, Steve Kent, Ambarish Malpani, Matt Fite, Michael Myers, Tim Polk, Stefan Santesson, Dave Simonetti, and Paul Hoffman.

## 7 References

[AC] S. Farrell, R. HousleyMcNeil, M., and Glassey, T., "An Internet Attribute Certificate Profile for Authorization," <[draft-ietf-pkix-ac509prof-01.txt](#)>, October 1999.

[CMC] Myers, M., Liu, X., Fox, B., and Weinstein, J., "Certificate Management Messages over CMS," <[draft-ietf-pkix-cmc-05.txt](#)>, 14 July 1999.

[CMP] Adams, C., Farrell, S., "Internet X.509 Public Key Infrastructure Certificate Management Protocols", [RFC 2510](#), March 1999.

[CMP-HTTP] Tschal"ar, R., Kapoor, A., and Adams, C., "Using HTTP as a Transport Protocol for CMP", <[draft-ietf-pkix-cmp-http-00.txt](#)>, August 1999.

[CMP-TCP] Tschal"ar, R., Kapoor, A., and Adams, C., "Using TCP as a Transport Protocol for CMP", <[draft-ietf-pkix-cmp-tcp-00.txt](#)>, August 1999.

[INTEROP] Moskowitz, R., "CMP Interoperability Testing: Results and Agreements", <[draft-moskowitz-cmpinterop-00.txt](#)>, June 1999.

[CMS] R. Housley, "Cryptographic Message Syntax," [RFC 2630](#), July 1999.

[CRMF] Myers, M., Adams, C., Solo, D., and Kemp, D., "Internet X.509 Certificate Request Message Format," [RFC 2511](#), March 1999.

[DHPOP] Prafullchandra, H., and Schaad, J., "Diffie-Hellman Proof-of-Possession Algorithms," <[draft-ietf-pkix-dhpop-02.txt](#)>, 1 October 1999.

[DNS] Mockapetris, P.V., "Domain names - concepts and facilities," [RFC 1034](#), November 1987.

[DVCS] Adams, C., Sylvester, P., Zolotarev, M., Zuccherato, R., "Internet X.509 Public Key Infrastructure Data Certification Server Protocols", <[draft-ietf-pkix-dcs-03.txt](#)>, 15 October 1999.

[ECDSA] Bassham, L., Johnson, D., and Polk, W., "Internet x.509 Public Key Infrastructure: Representation of Elliptic Curve Digital Signature Algorithm (ECDSA) Keys and Signatures in Internet X.509 Public Key Infrastructure Certificates," <[draft-ietf-pkix-ipki-ecdsa-01.txt](#)>, 3 June 1999.

[ETNPT] Namjoshi, P., "Internet X.509 Public Key Infrastructure Extending trust in non repudiation tokens in time," <[draft-ietf-pkix-extend-trust-non-repudiation-token-00.txt](#)>, 28 May 1999.

[IP] Postel, J., "Internet Protocol," [RFC 791](#), September 1981.

[IPv6] Deering, S., and Hinden, R., "Internet Protocol, Version 6 [[IPv6](#)] Specification," [RFC 1883](#), December 1995.

[FORMAT] Housley, R., Ford, W., Polk, W., and Solo, D., "Internet X.509 Public Key Infrastructure Certificate and CRL Profile," [RFC 2459](#), January 1999.

[FTPHTTP] Housley, R., and Hoffman, P., "Internet X.509 Public Key Infrastructure Operational Protocols: FTP and HTTP," [RFC 2585](#), July 1998.

[KEA] Housley, R., and Polk, W., "Internet X.509 Public Key Infrastructure Representation of Key Exchange Algorithm (KEA) Keys in Internet X.509 Public Key Infrastructure Certificates," [RFC 2528](#), March 1999.

[LAAP] Farrell, S., Chadwick, C.W., "Limited AttributeCertificate Acquisition Protocol", <[draft-ietf-pkix-laap-00.txt](#)>, Octoboe 1999.

[LDAPv2] Yeong, Y., Howes, T., and Kille, S., "Lightweight Directory Access Protocol", [RFC 1777](#), March 1995.

[MISPC] Burr, W., Dodson, D., Nazario, N., and Polk, W., "MISPC Minimum Interoperability Specification for PKI Components, Version 1", <<http://csrc.nist.gov/pki/mispc/welcome.html>>, 3 September 1997.

[OCSP] Myers, M., Ankney, R., Malpani, A., Galperin, S., and Adams, C., "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP," [RFC 2560](#), June 1999.

[OCSPX] Myers, M., Ankney, R., Malpani, A., Galperin, S., and Adams, C., "OCSP Extensions," <[draft-ietf-pkix-ocsp-x-00.txt](#)>, 3 September 1999.

[PEM] Kent, S., "Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management," [RFC 1422](#), February 1993.

[PKCS10] RSA Laboratories, "The Public-Key Cryptography Standards(PKCS)," RSA Data Security Inc., Redwood City, California, November 1993 Release.

[PKI-LDAPv2] Boeyen, S., Howes, T., and Richard, P., "Internet X.509 Public Key Infrastructure Operational Protocols - LDAPv2," [RFC 2559](#), April 1999.

[PKI-LDAPv3] Chadwick, D.W., "Internet X.509 Public Key Infrastructure Operational Protocols - LDAPv3," <[draft-ietf-pkix-ldap-v3-01.txt](#)>, August 1999.

[POLPRAC] Chokhani, S., and Ford, W., "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework," [RFC 2527](#), March 1999.

[QC] Santesson, S., Polk, W., and Gloeckner, P., "Internet X.509 Public Key Infrastructure Qualified Certificates", <[draft-ietf-pkix-qc-01.txt](#)>, 6 August 1999.

[RFC-822] Crocker, D., "Standard for the Format of ARPA Internet Text Messages," [RFC 822](#), August 1982.

[SCHEMA] Boeyen, S., Howes, T., and Richard, P., "Internet X.509 Public Key Infrastructure LDAPv2 Schema," [RFC 2587](#), June 1999.

[SCVP] Malpani, A., Hoffman, P., "Simple Certificate Validation Protocol (SCVP)," <[draft-ietf-pkix-scvp-01.txt](#)>, 9 August 1999.

[SIMONETTI] Simonetti, D., "Re: German Key Usage", posting to ietf-pkix@imc.org mailing list, 12 August 1998.

[TSP] Adams, C., Cain, P., Pinkas, D., and Zuccherato, R., "Internet X.509 Public Key Infrastructure Time Stamp Protocols", <[draft-ietf-pkix-time-stamp-04.txt](#)>, September 1999.

[X.509] ITU-T Recommendation X.509 (1997 E): Information Technology -

Open Systems Interconnection - The Directory: Authentication Framework, June 1997.

[X9.42] ANSI X9.42-199x, Public Key Cryptography for The Financial Services Industry: Agreement of Symmetric Algorithm Keys Using Diffie-Hellman (Working Draft), December 1997.

[X9.55] ANSI X9.55-1995, Public Key Cryptography For The Financial Services Industry: Extensions To Public Key Certificates And Certificate Revocation Lists, 8 December, 1995.

[X9.57] ANSI X9.57-199x, Public Key Cryptography For The Financial Services Industry: Certificate Management (Working Draft), 21 June, 1996.

## [8](#) Security Considerations

TBSL

## 9 Editor's Address

Alfred Arsenault U. S. Department of Defense 9800 Savage Road Suite  
6734 Fort George G. Meade, MD 20755-6734 (410) 684-7114  
awarsen@missi.ncsc.mil

Sean Turner IECA, Inc. 9010 Edgepark Road Vienna, VA 22182 (703)  
628-3180 turners@ieca.com

## 10 Disclaimer

This work constitutes the opinion of the editors only, and may not  
reflect the opinions or policies of their employers.

Expires April 22, 2000