



**KAISER PERMANENTE INFORMATION TECHNOLOGY  
ENTERPRISE PLANNING AND ARCHITECTURE**

# **PUBLIC KEY INFRASTRUCTURE CONCERNS IN HEALTHCARE SETTINGS**

Dave Barnett  
Systems Architect  
Enterprise Planning and Architecture  
Kaiser Permanente  
*February 26, 2000*

**Copyright © 2000 by Kaiser Permanente Medical Care Program**

**Summary:** This document discusses issues regarding Public Key Infrastructure (PKI) implementations in healthcare settings. It is based on the experience of Kaiser Permanente during preliminary design of an Enterprise PKI for multiple applications.

Issues addressed include:

- Technical and operational PKI interoperability between healthcare providers, partners, affiliates, and patients.
- Privilege management in healthcare
- Long-term storage of electronic medical records.

Please direct your comments and questions to:

Author: Dave Barnett  
Systems Architect, Enterprise Planning and Architecture  
Kaiser Permanente  
Address: 25 North Via Monte SH80, Walnut Creek, CA 94598  
Phone: (925) 926-3520  
E-mail: [dave.barnett@kp.org](mailto:dave.barnett@kp.org)

**Copyright © 2000 by Kaiser Permanente Medical Care Program**

# Table Of Contents

<b>INTRODUCTION .....</b>	<b>1</b>
<b>INTEROPERABILITY ISSUES .....</b>	<b>2</b>
CERTIFICATE POLICY .....	2
<i>Background</i> .....	2
<i>Issue</i> .....	3
<i>Recommendations</i> .....	4
ASSURANCE OF IDENTITY .....	4
<i>Background</i> .....	4
<i>Issue</i> .....	5
<i>Recommendations</i> .....	5
PROFILE PROLIFERATION .....	6
<i>Background</i> .....	6
<i>Issue</i> .....	10
<i>Recommendations</i> .....	10
TRUST MODELS .....	11
<i>Background</i> .....	11
<i>Issue</i> .....	13
<i>Recommendation</i> .....	13
<b>PRIVILEGE MANAGEMENT.....</b>	<b>14</b>
BACKGROUND .....	14
<i>Attribute Certificates</i> .....	14
<i>Authorization API</i> .....	17
ISSUE .....	18
RECOMMENDATION .....	18
<b>LONG TERM STORAGE.....</b>	<b>19</b>
BACKGROUND .....	19
ISSUE .....	20
RECOMMENDATION .....	20
<b>SUMMARY OF ISSUES AND RECOMMENDATIONS .....</b>	<b>21</b>
CERTIFICATE POLICIES .....	21
ASSURANCE OF IDENTITY .....	21
PROFILE PROLIFERATION.....	21
TRUST MODELS .....	22
PRIVILEGE MANAGEMENT.....	22
LONG TERM STORAGE.....	22

## INTRODUCTION

This paper presents several concerns, which although not exclusively those of Healthcare, none-the-less represent issues which could have wide-ranging consequences to those who conduct business with, or are direct providers of, healthcare services.

With the passage of Public Law 104-191 (Health Insurance Portability and Accountability Act of 1996), and requirements for electronic clinical information systems and services, the protection of patient records has become increasingly complex and more critical. Although identification, authentication, authorization, and integrity services have traditionally been provided on an application-by-application basis, the requirements of today's health care and business world make "point" solutions impractical and ineffective.

A comprehensive security system for distributed networks needs to provide integrated, extensible, and flexible services, based on public interfaces. One effective way to provide this is through infrastructure components, shared among applications.

PKI, or Public Key Infrastructure, services are the solution of choice for shared encryption/decryption, digital signature, and authentication services. These services can be used in a variety of settings, including web-enabled patient and clinical information services, Virtual Private Networks for e-business, and secure E-mail.

PKI is a standards-based, flexible and extensible set of services. One of the downsides of flexible standards is that the choice of options implemented may have a significant impact on interoperability. Unless everyone agrees on which options are to be supported, and how those options are implemented, interoperability is not assured.

Although there is some utility in a PKI only used internally, i.e., within an organization, the benefits are significantly greater when used to support secure and trusted electronic communications with partners, affiliates, and patients. Healthcare is an integrated industry that requires secure and trusted electronic communication between many business partners. According to the California Medical Association, physicians typically do business with 50 to 100 different healthcare organizations.

Although interoperability concerns are not restricted to healthcare, it is crucial that standardized options and agreements on interoperability be developed for use within the healthcare community of interest. We have common interests and concerns. The more we can agree on how to implement PKI, the easier it will be to work together for the common benefit of our patients and our organizations.

# INTEROPERABILITY ISSUES

## CERTIFICATE POLICY

### BACKGROUND

In a PKI, there are two formal policy statements considered integral to deployment. These are the Certificate Policy (CP) and the Certification Practice Statement (CPS). There are no commonly agreed upon definitions of exactly what these two are, and they are often confused.<sup>1</sup>

Although there are some technical issues associated with CP and CPS, their primary function is a declaration of organizational policy about how PKI will be implemented and used. The Certificate Policy and Certification Practice Statement are legal statements limiting liability. X.509 certificates may contain a reference to these documents, binding the liability constraints to the certificate. For example, each VeriSign Certificate contains a field that directs the user to a web site that lists the policy describing the acceptable use of that certificate.<sup>2</sup>

The Internet Engineering Task Force (IETF) informational publication RFC 2527<sup>3</sup> states:

“A Version 3 X.509 certificate may contain a field declaring that one or more specific certificate policies applies to that certificate. According to X.509, a certificate policy is ‘a named set of rules that indicates the applicability of a certificate to a particular community and/or class of application with common security requirements.’ A certificate policy may be used by a certificate user to help in deciding whether a certificate, and the binding therein, is sufficiently trustworthy for a particular application.

A more detailed description of the practices followed by a CA [Certification Authority] in issuing and otherwise managing certificates may be contained in a certification practice statement (CPS) published by or referenced by the CA. According to the American Bar Association Digital Signature Guidelines (hereinafter "ABA Guidelines"), ‘a CPS is a statement of the practices which a certification authority employs in issuing certificates.’”

The Canadian Government Treasury Board clarifies the relationship between Certificate Policy, Certification Practice Statements, and interoperability in one of their PKI web sites:<sup>4</sup>

---

<sup>1</sup> Adams, Carlisle, and Sam Lloyd. *Understanding PKI*. Indianapolis; Macmillan, 1999

<sup>2</sup> <http://www.verisign.com/repository/CPS/>

<sup>3</sup> Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework, <ftp://ftp.isi.edu/in-notes/rfc2527.txt>

<sup>4</sup> [http://www.cio-dpi.gc.ca/pki/Documents/Certificate%20Policy/aboutCP\\_e.html](http://www.cio-dpi.gc.ca/pki/Documents/Certificate%20Policy/aboutCP_e.html)

“A CP states what assurance can be placed in a certificate. A CPS states how a CA establishes that assurance. A certificate policy may apply more broadly than to just a single organization; a CPS applies only to a single CA.

Certificate policies best serve as the vehicle on which to base common interoperability standards and common assurance criteria industry-wide (or possibly more global). A detailed CPS alone does not form a suitable basis for interoperability between CAs operated by different organizations.”

One of the issues in implementing a PKI is the effort required to create these documents. First, they are policy statements, and require high level review and decision making. Secondly, they are legal documents as well, and need to be approved by the organization’s legal department.

An implication for PKI interoperability is that organizations need to review each other’s CP and CPS. Since there are liability issues, attorneys from both organizations agree about who is responsible for what, and other legal obligations. This could be an expensive and time-consuming process.

One way to handle the problem of creating compatible CP and CPS documents is to develop a common policy for a community of interest. This is known as a “Model Policy” and, once published, can save all of the prospective users a great deal of effort and time. There may be several variations of the policy, depending on security and policy requirements of particular situations. These are usually referred to as Assurance Levels, such as Rudimentary, Basic, Medium, and High Assurance.

One of the technical aspects to the use of Certificate Policy and Certification Practices Statements is that policies, or sections of policies, should be assigned a unique object identifier (OID) which can be embedded in a certificate. OID’s from different organizations can also be mapped to each other within a certificate. This permits applications to make programmatic decisions about appropriate use of the certificate, and what privileges to grant the holder. For example, a PKI enabled application can look at an external researcher’s certificate, and see that there is an approved non-disclosure agreement on file, and that the holder is a licensed physician. The application can then grant access to clinical research information that would otherwise be confidential.

Many of the issues discussed in this document are policy issues, and can be included as part of a Model Policy statement. Some of the topics, however, are worthy of their own sections, and are presented as such.

There are many legal and organizational issues that can be solved by the model policy approach. The value of a Model Policy, and agreements based on that policy, cannot be over-emphasized.

## **ISSUE**

Inter-organizational agreements can be time consuming and expensive. Certificate Policies and Certification Practice Statements are legal documents identifying limits of liability. Conflicts, misalignment, and differing interpretations of requirements can result in significant interoperability problems.

## RECOMMENDATIONS

- **Develop and publish a Model Policy for the Healthcare Community of Interest, containing appropriate levels of assurance**
- **Use ASTM E31.20 draft Model Certificate Policy for Healthcare PKI as basis<sup>5</sup>**

## ASSURANCE OF IDENTITY

### BACKGROUND

Not all certificates are issued equally. Some certificates are issued based on a valid e-mail address. Other certificates require presentation of photo ID and must be vouched for by a trusted party. There have been proposals for group practice certificates. It is also possible to download a demo copy of a CA (Certification Authority) package, and create “roll-your-own” certificates. Each has an appropriate use. However, the amount of trust we can place in each of these types of certificates varies considerably.

The degree of trust that we can place in a certificate is an important part of the Certificate Policy and Certification Practice Statements, and is closely related to the Assurance Level associated with that policy.

Although the nature of signed certificates prevents “spoofing” (assuming that the certificates are validated), the levels of trust that can be placed in each certificate varies depending on the policies of the issuer. The vulnerability is not in the certificate itself, but in the method used to determine the identity of the requestor. For example, someone skilled in “social engineering<sup>6</sup>” might be able to acquire a healthcare provider digital certificate fraudulently. If the imposter has sufficient details about a legitimate practitioner, and can plausibly pose as that person, the door is opened for impersonation. Even though this may be a felony, if the acquisition can be accomplished remotely (e.g., by telephone or on-line using a phony, redirected e-mail address), the risk of getting caught may be low enough, and the potential gain attractive enough, that this situation is an inducement to misrepresentation.

Determining the degree of trust associated with a certificate is, in many cases, a business decision, based on best practices and liability/risk factors. Due diligence would require that prior to accepting a certificate issued by an external CA, an organization must understand and accept the conditions under which the certificate was issued. (This should be covered in the Certificate Policy and Certification Practice Statements.)

In other cases, there may a need to rely on a high degree of assurance surrounding the identity of the person who was issued the certificate. For example, if the DEA were to permit digitally signed electronic prescriptions, it would be necessary to specify the level of assurance required. Presumably, it would be more than a valid e-mail address.

In addition to “proof of identity,” there are other assurance requirements. For example, there is also a need to know something about the security practices around certificate

---

<sup>5</sup> A public copy can be viewed at <http://www.tunitas.com/pages/PKI/pki.htm>

<sup>6</sup> The ability to plausibly represent oneself as an authorized individual, and obtain information or access by subterfuge.

issuance and protection. If an organization did not provide adequate physical security for securing their Root CA server, we would have low assurance and trust in those certificates. If a policy required strong authentication, and that private keys of users never be outside of their control, this would exclude software based “roaming user” solutions. These types of assurance are covered in Certificate Policy and Certification Practice Statements.

## ISSUE

There is a lack of commonly understood and accepted terms describing various levels of assurance for certificate issuance. Each certificate binds an identity to a public key, and can be used to authenticate that individual. However, the proof of identity that is required varies by certificate issuer. This can result in situations that may cause increased risk of harm to patients and increased liability.

## RECOMMENDATIONS

- **Leverage the work of IETF PKIX Qualified Certificates<sup>7</sup> (a certificate whose primary purpose is identifying a person with high level of assurance in public non-repudiation ) and appropriate Model Policies**
- **Develop, define, and publish a set of standard assurance levels for Healthcare X.509 certificates.**

Encourage the use of these descriptions in Certificate Policy and Certification Practice Statements, and incorporate these descriptions into Federal and State regulations governing Healthcare PKI use.

- **Integrate these standards into a Model Certificate Policy for Healthcare, or create Model Policies for several assurance levels.**

---

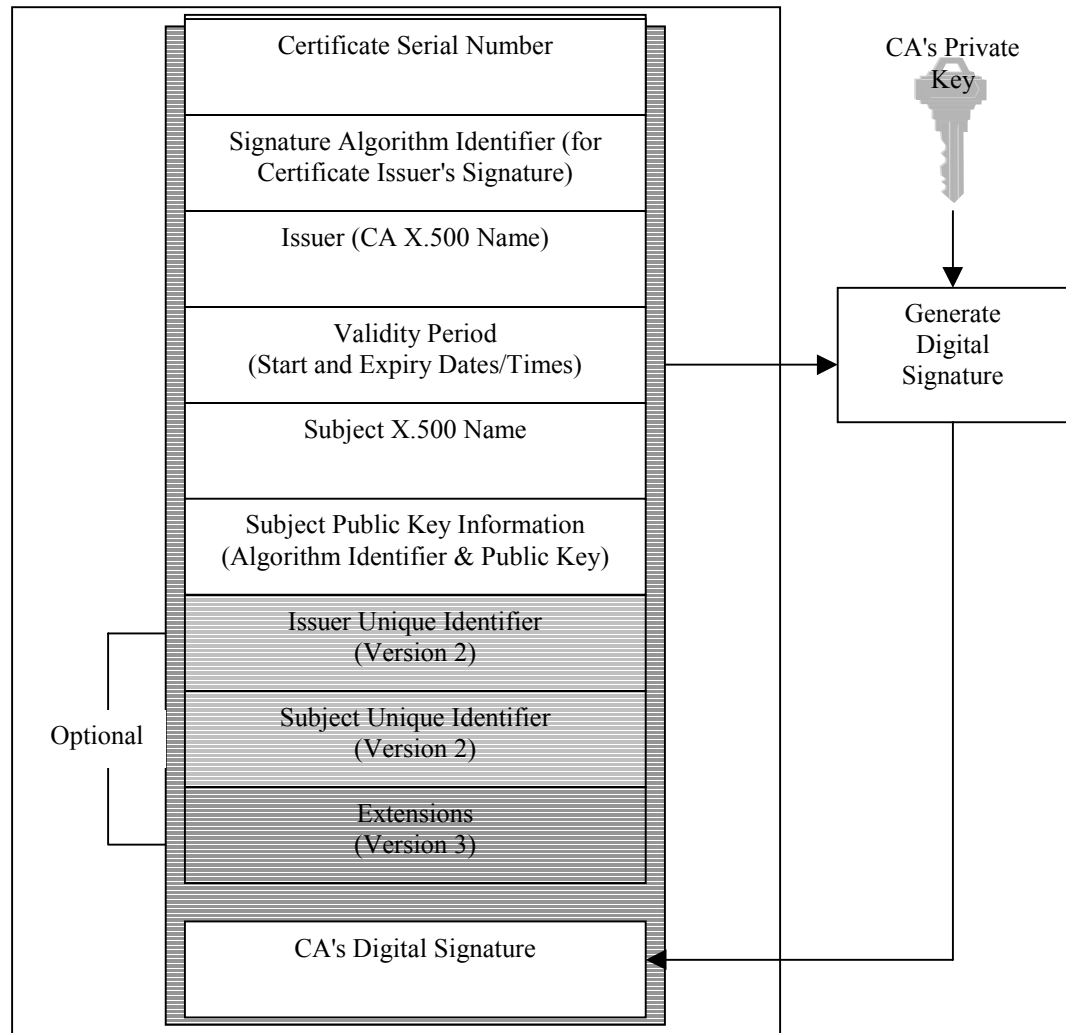
<sup>7</sup> <http://www.ietf.org/internet-drafts/draft-ietf-pkix-qc-02.txt>

## PROFILE PROLIFERATION


### BACKGROUND

X.509 is the generally agreed upon standard for digital certificates. Versions 1 and 2 provided the basic functionality, but little beyond that. Because of the lack of extensibility, V1 and V2 were not deployed widely. Version 3, the most recent and most popular, is much more powerful and flexible. It also introduced a great deal of complexity.

The standard format for an X.509 certificate is:



In an X.509 V3 certificate, there can be 0 to  $n$  extensions. Extensions have the format:



There are several types of extensions, including those for keys, for policy use, for subject and issuer, and for constraints and limitations.

Any of the extensions can be marked critical or non-critical. If it is marked critical, the extension is required and the application accessing it *must* be able to parse that extension, or the certificate will be non-usable. If an extension is marked non-critical, it may be ignored if the application does not support it.

The standard options and extensions for a X.509 certificates are:

Field	Purpose	Options	Comments
Version	Version of Certificate format (v1, v2, or v3)	<ul style="list-style-type: none"> <li>• 0</li> <li>• 1</li> <li>• 2</li> </ul>	<ul style="list-style-type: none"> <li>• V1</li> <li>• V2</li> <li>• V3 (typical)</li> </ul>
Certificate Serial Number	Unique identifier by CA	Serial number (sequential by CA)	
Signature Algorithm	OID of algorithm	<ul style="list-style-type: none"> <li>• 1 3 14 3 2 29</li> <li>• 1 2 840 10040 4 3</li> <li>• 1 2 840 113549 1 1 4</li> </ul>	<ul style="list-style-type: none"> <li>• RSA with SHA-1</li> <li>• DSA with SHA-1</li> <li>• RSA with MD5</li> </ul>
Issuer (CA) X.500 name	Globally unique name	Country = Organization = Common Name = (and others, such as Organization Unit, Location, etc.)	C= US O= Kaiser Permanente CN= CA1
Validity Period		Start Date = Expiry Date =	The longer the period, the easier to manage, but the greater the risk and longer the CRL.
Subject X.500 name	Globally unique name	Country = Organization = Common Name = (and others, such as Organization Unit, Location, etc.)	C= US O= Kaiser Permanente CN= Yongtian Zhang  Or CN=dave.barnett@kp.org OU= KP-IT OU= Enterprise Planning OU= Regional Offices L= SH80
Subject Public Key	Type and value of public key	Algorithm identifier = (OID) PK value =	Algorithm used and public key value
Issuer Unique Identifier	Organizational unique identifier (to avoid name collisions)		V2 Option (rarely used) Deprecated in RFC 2459

Field	Purpose	Options	Comments
Subject Unique Identifier	Organizational unique identifier (to avoid name collisions)		V2 Option (rarely used) Deprecated in RFC 2459
Authority Key Identifier	Hash of CA's public key, used for more efficient lookups		V3 Extension
Subject Key Identifier	Hash of subject public key, used for more efficient lookups		V3 Extension
Key Usage	Identifies purpose, or acceptable usage, of key (can be combinations)	<ul style="list-style-type: none"> <li>• Digital signature</li> <li>• Key Encryption</li> <li>• Data Encryption</li> <li>• Key agreement</li> <li>• CRL signing</li> </ul>	V3 Extension
Extended Key Usage	Finer grained detail on key use.	Sequence of OIDs identifying specific use, such as code signing, time stamping, SSL/TLS, etc.	V3 Extension
CRL Distribution Point	Location of CRL partition for revocation information about this certificate		V3 Extension
Private Key usage period	Allows validity period of key to be less than certificate (for key updates)	(start/expiry)	V3 Extension (RFC 2459 recommends not to use)
Certificate Policies	Identifies policies or practices that the CA supports for this certificate	<p>Defined by CA organization, and identified with OID.</p> <p>May include a qualifier, which can, for example, use a URL to point to the formal statement.</p>	V3 Extension  A named set of rules, defined by the organization and issued an OID, that indicates the acceptable use of this certificate. For example, may include community of use, key management, authentication policies, CRL update frequency, etc.
Policy Mappings	Allow policies of one organization to be mapped to policies of another	OID equivalencies between policies	V3 Extension
Subject Alternative name	Allows aliases for subject name	<p>One or more of:</p> <ul style="list-style-type: none"> <li>• X.500</li> <li>• X.400</li> <li>• RFC 822 (email)</li> <li>• DNS</li> <li>• EDI</li> <li>• URI/URL</li> <li>• IP address</li> <li>• OID</li> </ul>	V3 Extensions

Field	Purpose	Options	Comments
Issuer Alternative name	Allows aliases for issuer name	One or more of: <ul style="list-style-type: none"> <li>• X.500</li> <li>• X.400</li> <li>• RFC 822 (email)</li> <li>• DNS</li> <li>• EDI</li> <li>• URI/URL</li> <li>• IP address</li> <li>• OID</li> </ul>	V3 Extensions
Subject Directory Attributes	Additional information for X.500 directory	Can be Title, telephone number, address, access control, etc	V3 Extensions
Basic Constraints	Limit exposure to unauthorized certificates	<ul style="list-style-type: none"> <li>• Is this certificate for a CA (can issue certificates) or an end-user?</li> <li>• If a CA, what are the max hops for signature validation?</li> </ul>	V3 Extension for Certification Path Constraints  If used, must be critical
Name Constraints	Limit exposure to unauthorized certificates	CA Certificate only, applies to names subordinate to (or with wildcard/range): <ul style="list-style-type: none"> <li>• X.500 name</li> <li>• DNS</li> <li>• IP address (or subnet)</li> </ul>	V3 Extension for Certification Path Constraints  If used, must be critical <ul style="list-style-type: none"> <li>• C=US,O=KPMCP</li> <li>• kp.org</li> <li>• 10.9.8.0</li> </ul>
Policy Constraints	Limit exposure to unauthorized use	<ul style="list-style-type: none"> <li>• CA can set requirement that all certificates in path must have policy xyz.</li> <li>• CA can set requirement that no certificates in path may use policy mapping</li> </ul>	V3 Extension for Certification Path Constraints  If used, must be critical
Custom	Proprietary use and custom extensions	Need OID	V3 allows ability to create new extensions as needed.

In implementing a PKI, an important design decision is determining which algorithms, extensions, and options should be included. Understanding what the options are, how they work together, and how they are implemented in products and other PKI's is a time consuming task. Many extensions permit multiple options. For example, the name fields may use an email address, an X.400 address, a URL, an IP address, or an X.500 directory address, in almost any combination. In addition to the standard extensions, X.509 permits adding customized extensions.

The set of extensions and how they are implemented is a “certificate profile.” Part of a PKI design is determining what is required in the certificate profile.

There is a tendency, when first embarking on the PKI route, to identify all the extensions that sound interesting, and mark the important ones as “critical.” This is an acceptable design, as long as you never intend to use those certificates to interoperate with anyone outside of your organization. There is also a tendency for each community of users to create their own profile. This results in the same problem of interoperability as the community expands. Applications, too are subject to this tendency. For example, the certificate used for S/MIME may not work very well for VPN.

## ISSUE

The extensibility of X.509 V3 certificates is a blessing and a curse. The more options used, the more likely there will be interoperability problems. There is a tendency to create specialized profiles (sets of options), by application, by organization, and by community of interest. Each profile represents a significant amount of potentially duplicated effort. The end result, if unchecked, could be a need to issue and manage dozens of certificates for each end user. This substantially increases maintenance efforts, cost, complexity, and presents significant barriers to interoperability.

## RECOMMENDATIONS

- **Converge on the smallest number of profiles that will meet the requirements.**

This may mean development of industry profiles. For example, both the automotive industry ANX (Automotive Network Exchange), and the Federal Government PKI Technical Working Group have published profiles for interoperability within their respective communities of interest.

- **Develop a standard Healthcare X.509 V3 certificate profile. This will enhance interoperability and reduce implementation efforts.**

Where possible, distinguish between certificates used for signatures and those used for encryption. In most corporate settings, encryption certificates will need to be escrowed. Signature certificates should never be escrowed. (Escrow of digital signature certificates creates a situation allowing repudiation. It can be viewed as a theft of identity.)

- **Leverage existing work.**

The IETF PKIX (Public Key Infrastructure using X.509) Working Group<sup>8</sup> has published an Internet draft standard (RFC 2459) for X.509 Profiles. This is a very generic profile, with widespread recognition.

ASTM E-31 (Healthcare Informatics) has established a subcommittee (E31.20) to draft PKI profiles and policies for healthcare.<sup>9</sup>

---

<sup>8</sup> <http://www.ietf.org/html.charters/pkix-charter.html>

The Federal PKI Technical Working Group has published white papers and architecture documents for building a complex, distributed PKI. Included in this body of work is a Federal PKI Profile.<sup>10</sup>

## TRUST MODELS

### BACKGROUND


One of the major problems to overcome in setting up inter-organizational PKI's is cross-certification and trust models. A trust model describes the relationship between CA's. That is, whether they can recognize and validate each other's signatures. Cross-certification is the technical method used to establish trust by verifying signatures.

Within an organization, a hierarchical trust model can work well, especially if there is little potential for divestiture. In the hierarchical mode, there is one root CA, and it signs all certificates subordinate to it. There can be a single level (all end user certificates are signed by the root CA), or multiple levels (The root CA signs the certificate of first level subordinate CA's, which in turn sign the end user certificates or other lower level subordinate CA's).

In the real world, there is not one Root CA for all organizations. In order to validate the signature of a certificate signed by a CA outside of a strict hierarchy, CA's use cross-certificates. A cross-certificate is special type of certificate, in that both the subject and issuers are CA's. In mutual cross-certification, two peer CA's each sign each other's certificates. This allows signature validation to traverse organization boundaries. In other words, it establishes trust relationships in which one organization vouches for another, by allowing digital signatures by be verified outside of a direct hierarchy.

Although peer-to-peer trust relationships work well in settings in which there are a small number of peers, they don't scale well. Each peer has to cross-certify every other peer. For example, three peers could cross certify with just three bilateral agreements.

- CA 1 and CA 2 trust each other
- CA 1 and CA 3 trust each other
- CA 2 and CA 3 trust each other



However, the number of agreements increases rapidly.<sup>11</sup> For six peers, fifteen cross-certifying

<sup>9</sup> A public copy of the Model Certificate Policy for Healthcare PKI can be viewed at <http://www.tunitas.com/pages/PKI/pki.htm>


<sup>10</sup> <http://csrc.nist.gov/pki/twg/welcome.html#documents>

<sup>11</sup> The number of agreements required for  $n$  peers is  $(n^2-n)/2$ .

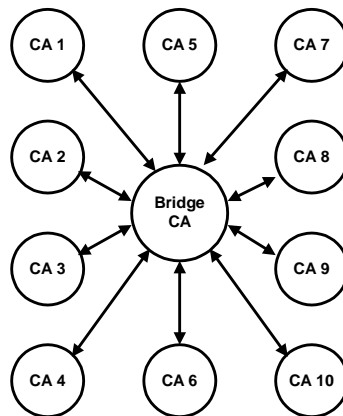
agreements are needed. For twenty peers, one hundred and ninety agreements would be needed.

In complex environments, such as those in healthcare, these extensive many-to-many relationships are not practical. Instead, a spoke and hub topology greatly simplifies the trust relationships. In these, each CA sets up a bilateral trust agreement (i.e., a cross-certification pair) with the hub. As new CA's are added, they only need to create an agreement with the hub. The hub creates the trust "paths" to every other CA that is attached.

This trust model is known as the Bridge CA. The Federal PKI, another instance of a complex environment, has developed and demonstrated a Bridge CA pilot. As an example of the simplification of the Bridge CA model, assume we have an environment with ten organizations each with its own CA. This modest number of CA's would require 45 bilateral agreements to interoperate:



In contrast, a Bridge CA, using the hub and spoke format, would only need to create ten bilateral trust agreements.



A Bridge CA can provide easier management trust relationships with both other Bridge CA's and peer CA's.

## ISSUE

If PKI's are created within organizations, and then are allowed to grow and establish *ad hoc* trust relationships with other organizations, the inter-organization agreements (including CP mapping) can rapidly become unmanageable.

## RECOMMENDATION

- Large organizations should consider developing their own Bridge CA to provide for a controlled trust interface to all external CA's
- Small organizations should consider using the services of an umbrella Bridge CA service provider

Kaiser Permanente is implementing a Bridge CA to provide an external trust interface to other CA's. The California Medical Association is providing a Bridge CA for use by California physicians. Kaiser Permanente and the CMA are cooperating in developing a healthcare Bridge CA to Bridge CA interoperability model, within the Tunitas Group Healthcare PKI demonstration.<sup>12</sup>

---

<sup>12</sup> <http://www.tunitas.com/pages/PKI/pki.htm>

# PRIVILEGE MANAGEMENT

## BACKGROUND

Management of access control and privileges in medical settings can become very complex. Clinicians, Business Office/Billing staff, Appointment Clerks and Receptionists, Psychiatric Social Workers, Advice Nurses, and Pharmacists have widely varying requirements for access to patient records. As we move to electronic medical records, the simple physical separation of types of charts and other paper documents is no longer applicable.

Each application accessing part of an electronic medical record needs a set of access controls. Role based authentication and access control makes this work easier, but it still represents a significant amount of duplicated effort. From an architectural view, it makes sense to move the authorization routines into the infrastructure. This allows component re-use as well as centralized and consistent management of privileges. The common authorization management function is referred to as a Privilege Management Infrastructure, or PMI.

The ITU defines PMI as "the complete set of processes required to provide an authorization service." The goal of PMI is to move redundant (and in some cases, conflicting) authorization functions out of individual applications and into an enterprise infrastructure.

During the past year, there has been increasing interest and work in PMI standards and products. However, standards are still in draft status and current products are immature and proprietary. Early product entries in this space include Baltimore Technologies Attribute Certificate Server, Entrust Secure Control, and DASCOS IntraVerse products. (IBM SecureWay uses the DASCOS technology, and has recently acquired the company).

There are two emerging approaches to PMI. Neither claims to fully define PMI, and each is represented as a subset of some of the functions that would be used in a full authorization framework. The ITU and IETF are using Attribute Certificates, while The Open Group is promoting the Authorization API (aznAPI). Since the approaches are at different layers, and represent different components, they are not mutually exclusive. The aznAPI allows for use of Attribute Certificates.

## ATTRIBUTE CERTIFICATES

The ITU notes that "Public-key certificates are principally intended to provide an identity service upon which other security services, such as data integrity, entity authentication, confidentiality and authorization, may be built." While X.509 certificates bind an identity to a public key, Attribute Certificates ("Att-Certs" or AC's) bind an identity to a

role, a set of privileges, or an attribute. The IETF draft on Attribute Certificates<sup>13</sup> presents the following analogy of the differences between X.509 PKC's (Public Key Certificates) and AC's:

"A PKC can be considered to be like a passport: it identifies the owner, tends to last for a long period and shouldn't be too easy to get. An AC is more like an entry visa in that it is typically issued by a different authority and doesn't last as long. As acquiring an entry visa typically requires presenting a passport, getting a visa can be a simpler process."

Since Attribute Certificates are digitally signed, they can be distributed over non-secure networks and be stored in distributed repositories, with the same advantages (and drawbacks) as X.509 certificates. One benefit that results from this technology is, unlike most authorization and access control, Attribute Certificates can be cached in memory or on disk without fear of compromise. This may result in substantial performance and availability improvements. In addition, the applications using on-line validation, such as Online Certificate Status Protocol (OCSP), can be readily adapted to use Attribute Certificates instead of authentication certificates.

Situations where the use of Attribute Certificates might be useful include electronic commerce, web access control, VPN access control, and interactions with affiliates and partners.

Privileges *can* be embedded in a standard X.509 Public Key Certificate. However, by decoupling the access control information from the authentication information (e.g. by using both PKC's and AC's), we can manage them separately. The Certification Authority (CA) issues Public Key Certificates and the Attribute Authority (AA) issues ACs. The validity period for a privilege may be far shorter than the validity period for authentication. We wouldn't want to have to revoke a PKC every time a privilege changes.

Attribute Certificates support both role and rule based authorization. Role Certificates are defined in the as a type of Attribute Certificate. The Role Certificate can be signed and managed by another authority (like a profile). The Attribute Certificate can bind an identity to a Role Certificate, which in turn provides authorization attributes. Alternatively, the Attribute Certificate can bind an identity to a specified set of privileges or attributes.

The International Telecommunication Union (ITU) has recently published the final proposed draft amendments (FPDAM) to X.509 Certificate extensions. Two major revisions are included -- certificate revocation enhancements and PMI using Attribute Certificates (v2). The description of PMI in the FPDAM is:

"A Privilege Management Infrastructure (PMI) is the set of processes required to provide authorization services. This Specification defines a framework for obtaining and trusting attributes of an entity in order to determine whether or not they are authorized to access a particular resource. The framework includes the issuance of an attribute certificate by an Attribute Authority (AA) and the validation of that attribute certificate by a verifier. The validation includes ensuring that the privileges in the

<sup>13</sup> <http://www.ietf.org/internet-drafts/draft-ietf-pkix-ac509prof-01.txt>

attribute certificate are sufficient to access the resource based on a predefined privilege policy, establishing a trusted delegation path of attribute certificates if necessary, verifying the digital signature on each certificate in the path, ensuring that each issuer was authorized to delegate privileges, and validating that the attribute certificates have not expired or been revoked by their issuers.”

Although PKI and PMI are separate infrastructures and may be established independently, they are related. This specification recommends that owners and issuers of attribute certificates be identified within attribute certificates by pointers to their appropriate public-key certificates. Authentication of the attribute certificate issuers and owners, to ensure that entities claiming privilege and issuing privilege are who they claim to be, is done using the normal processes of the PKI to authenticate identities. This authentication process is not duplicated within the attribute certificate framework.

The IETF PKIX working group has also published a draft (April 1999) on Attribute Certificates, and has indicated in its mission statement that it is pursuing additional work in this area. This draft states:

“The provision of authentication, data integrity and confidentiality services for current Internet protocols is well understood and many secure transports are defined (e.g. TLS, IPSEC, etc.). In many applications these services are not sufficient (or too cumbersome to administer) to provide the type of authorization services required.

When considering authorization, one is often less interested in the identity of the entity than in some other attributes, (e.g. roles, account limits etc.) which should be used to make an authorization decision.

In many such cases, it is better to separate this information from the identity for management, security, interoperability or other reasons. However, this authorization information also needs to be protected in a fashion similar to a public key certificate - the name for the structure used is an attribute certificate (an AC) which is a digitally signed (certified) set of attributes.

An AC is a structure that is similar to an X.509 public key certificate [RFC2459] with the main difference being that it contains no public key. The AC typically contains group membership, role, clearance and other access control information associated with the AC owner. The base syntax for ACs is also defined in the X.509 standard (making the term X.509 certificate ambiguous!). This document specifies a profile of the X.509 AC suitable for authorization purposes in Internet protocols.”

Entrust, SpyruS, and Baltimore Technologies endorse the Attribute Certificate approach. With the current interest from both the IETF and ITU, the use of Attribute Certificates for authorization is likely to become very significant within a few years. Attribute Certificates are highly compatible with X.500 Directory Services, and X.509 Public Key Certificates, and the IETF PKIX directions.

## AUTHORIZATION API

The Authorization Service API (aznAPI) is on The Open Group standards fast track.<sup>14</sup> DASCOM has been doing most of the work, and is the current developer of the aznAPI.

The main function of the aznAPI is to return a yes/no decision on access. This could greatly simplify the coding needed in applications for complex authorization and access control.

In earlier drafts of the aznAPI, there was discussion on how it provides a high level interface that can act as a facade for a variety of underlying technologies, including PKI, DCE, and Kerberos, or more complex authorization technologies such as a Policy Rule Evaluator engine. Architecturally, the aznAPI can be used as a generic PMI API supporting Attribute Certificates or other technologies. Both the X.509 Attribute Certificate model and the aznAPI are compatible with the ISO/IEC 10181-3 | ITU-T Rec. X.812 Authentication Framework.

The stated design goals of the aznAPI are:

- Define a simple, flexible API through which authorization functionality can be invoked.
- Enable application-transparent evaluation of policy rules to arrive at access decisions.
- Enable central management of policy independent of applications.
- Transparently support any reasonable authorization policy rule syntax (e.g. ACL's, labels, etc.)
- Separate authentication from authorization
- Permit derivation of authorization attributes from authentication data.
- Transparently support any reasonable authorization attribute type (groups, roles, etc.)
- Facilitate authorization in multi-tiered applications
- Permit externalization of authorization attributes for use in multi-tier application configurations
- Enable applications to access security policy (entitlement) information applicable to their resources.
- Support a variety of access control mechanisms as implementations of the API.
- Enable simultaneous use by a single application of multiple authentication and authorization services.
- Support application access to audit data related to the operation of the authorization service.

The aznAPI provides a level of abstraction which could hide dependencies on underlying technologies. Both tactically and strategically, this makes sense.

---

<sup>14</sup> <http://www.opengroup.org/publications/catalog/c908.htm>

## ISSUE

PMI seems a worthwhile approach for managing enterprise healthcare authorization services, particularly those that justify an infrastructure approach because of size or complexity. The endorsement of Attribute Certificates by the IETF indicates that PMI may become a significant component in electronic commerce applications. An Attribute Certificate framework could be closely integrated with PKI and Directory services, and could fit well with many current efforts. The use of aznAPI is also an attractive component of a PMI, and could be used in conjunction with Attribute Certificates, if supported by vendors. The aznAPI would also allow very complex authorization decisions to be made by a rules engine, and a simple yes/no result returned to applications. This has obvious value for the complex authorization requirements of healthcare. On the other hand, Attribute Certificates can be freely published and cached without fear of compromise, providing response time improvements and improved availability.

Because the products are at early stages of development, and the standards have not stabilized, current attempts at implementation would be risky, and require care. The growing interest and body of knowledge in the industry indicates that this is an architectural direction worth pursuing.

## RECOMMENDATION

- Healthcare organizations should encourage and become involved in further development of standards in this area.
- Healthcare organizations should ask PKI and security product vendors about plans for PMI.

# LONG TERM STORAGE

## BACKGROUND

In many settings, a record retention policy will require that after a certain period it is mandatory to destroy records. During investigations, discovery of old or even supposedly deleted e-mails can result in significant damage to a legal case and cause public embarrassment. A mandatory destruction requirement can save storage and maintenance and contain the effort needed to meet discovery orders. This limit on record retention might be very appropriate for in-patient facilities, and most organizations.

On the other hand, for long term care, it would be a disservice to both patients and physicians to destroy records after a given period. Our current policy at Kaiser Permanente is to retain records for at least the lifetime of the patient. For example, we might need to retain the electronic health records of an infant born today for a hundred years.

For long term retention, aside from the obvious technical issues of storage capacity<sup>15</sup>, there is also a problem with media deterioration over time, and with changing technology.<sup>16</sup> Tapes, CDs, and electronic media can deteriorate, and a restore or read of the original data may no longer have full integrity.

If an archived electronic record is included in a discovery process during litigation, the integrity of the record will be an issue. One way to assure the integrity of the record, and be able to convince a jury that the record has not been altered, is to use a write-only media. Adding a digital signature strengthens this proof. However, it can also become a vulnerability. Eventually, all electronic media deteriorate. However, where the loss of a bit or two might not be noticeable to visual inspection (e.g., it might appear as a “typo”), it would invalidate a digital signature. Any bit change in a digitally signed record, results in proof that the record has changed. In these circumstances, adding a digital signature makes the record appear more sensitive to deterioration.

In addition to increased sensitivity to deterioration, there is also a problem with key lifetimes. The technical basis for electronic cryptography is that some mathematical functions are very easy to perform one-way, but extremely difficult to reverse. The security is determined by the computational infeasibility of “breaking” the encryption. Encryption is effective only because it would take so long for a brute force decryption that the cost exceeds the value or useful period of the information.

As computational power increases (Moore’s law<sup>17</sup>), and as new algorithms<sup>18</sup> for factoring are developed, the key length needs to be increased to compensate. For planning purposes, the key length is usually determined by estimates of a key lifetime of ten to

---

<sup>15</sup> We expect our initial implementation of an electronic clinical repository to be at least 10 terabytes.

<sup>16</sup> For example: DVD, CD, 3.5” diskettes, 5 ¼” diskettes, 8” diskettes, paper tape...

<sup>17</sup> “Computing power doubles every 18 months,” or Moore’s Law, has been true for many years. However, the doubling is not likely to continue forever – some physical limits are expected to be reached soon

<sup>18</sup> Not only new algorithms, but also new paradigms, such as quantum and DNA/cellular computing may radically change our approach.

twenty years. That is, the “computational infeasibility” will probably continue to hold true for at least ten years.<sup>19</sup>

Bruce Schneier, after discussing key lifetimes of ten years, also notes that:

“The very notion of predicting computing power 10 years in the future, let alone 50 years is absolutely ridiculous. ...If the past is any guide, the future will be vastly different from anything we can predict. Be conservative. If your keys are longer than you can imagine necessary, then fewer technological surprises can harm you.”

On the other hand, Hilarie Orman presents a more optimistic view.<sup>20</sup> Even allowing for Moore’s Law, and enormous resources (US\$1 trillion, adjusted for inflation) available to launch a “brute force” attack, she calculates that current key lengths of 112 bits (Triple DES) allow for a lifetime of about fifty years.

However, even if current key lengths are good for fifty years, the key lifetime is less than the projected retention of some patient records. When the record retention period exceeds the key lifetime, the digital signatures and encryption strength lose effectiveness. That is, the signature can be forged, and the encryption can be broken. The ability to argue in court that an electronic record has not been altered since its creation, based on the digital signature, vanishes.

A key lifetime that is designed for ten or even fifty years is clearly inadequate for records that are to be retained for a hundred years. The simple solution – just increase the key size until it will last longer than the retention period – suffers from Moore’s law inverted: current systems are not powerful enough to handle significantly longer key lengths without significantly greater performance impact.

## ISSUE

Long term storage (greater than 50 years) of medical records poses several problems, especially in light of any potential litigation. (1) It is difficult to prove that an electronic record has not been altered. (2) Media deteriorates over time, destroying information and invalidating digital signatures. (3) Storage media technology changes over time, which may render archives obsolete and unreadable. (4) Key lengths are generally determined by an estimated lifetime of “computational infeasibility” of ten to fifty years. This is not long enough for electronic medical records, or others which require long-term archival.

## RECOMMENDATION

- Create a process and standard that allows for digitally signed records to be copied onto new media (and new technology) on a periodic basis (e.g., every 10 or 50 years). The period should not be fixed, but should be based on an index based on cryptographic technology advances, plus considerations for media deterioration. During this copy, the records should be “wrapped” or signed again with a key length that is appropriate for current technology. This period can be altered as the technology advances.

<sup>19</sup> Schneier, Bruce. *Applied Cryptography* (2<sup>nd</sup> Ed), New York: Wiley, 1996

<sup>20</sup> <http://search.ietf.org/internet-drafts/draft-orman-public-key-lengths-00.txt>

## SUMMARY OF ISSUES AND RECOMMENDATIONS

### CERTIFICATE POLICIES

Inter-organizational agreements can be time consuming and expensive. Certificate Policies and Certification Practice Statements are legal documents identifying limits of liability. Conflicts, misalignment, and differing interpretations of requirements can result in significant interoperability problems.

#### *Recommendations*

- Develop and publish a Model Policy for the Healthcare Community of Interest, with appropriate assurance levels.
- Use ASTM E31.20 draft Model Certificate Policy for Healthcare PKI as basis

### ASSURANCE OF IDENTITY

There is a lack of commonly understood and accepted terms describing various levels of assurance for certificate issuance. Each certificate binds an identity to a public key, and can be used to authenticate that individual. However, the proof of identity that is required varies by certificate issuer. This can result in situations that may cause increased risk of harm to patients and increased liability.

#### *Recommendations*

- Leverage the work of IETF PKIX Qualified Certificates<sup>21</sup> (a certificate whose primary purpose is identifying a person with high level of assurance in public non-repudiation ) and appropriate Model
- Develop, define, and publish a set of standard assurance levels for Healthcare X.509 certificates.
- Integrate these standards into a Model Certificate Policy for Healthcare, or create Model Policies for several assurance levels.

### PROFILE PROLIFERATION

There is a tendency to create specialized X.509 certificate profiles (sets of options), by application, by organization, and by community of interest. Each profile represents a significant amount of potentially duplicated effort. The end result, if unchecked, could be a need to issue and manage dozens of certificates for each end user. This substantially increases maintenance efforts, cost, complexity, and presents significant barriers to interoperability.

#### *Recommendations*

- Converge on the smallest number of profiles that will meet the requirements.

---

<sup>21</sup> <http://www.ietf.org/internet-drafts/draft-ietf-pkix-qc-02.txt>

- Develop a standard Healthcare X.509 V3 certificate profile. This will enhance interoperability and reduce implementation efforts.
- Leverage existing work, such as IETF PKIX and ASTM E31.20.

## **TRUST MODELS**

If PKI's are created within organizations, and then are allowed to grow and establish *ad hoc* trust relationships with other organizations, the inter-organization agreements (including Certificate Policy mapping) can rapidly become unmanageable.

### *Recommendation*

- Large organizations should consider developing their own Bridge CA to provide for a controlled trust interface to all external CA's
- Small organizations should consider using the services of an umbrella Bridge CA service provider

## **PRIVILEGE MANAGEMENT**

A Privilege Management Infrastructure (PMI) seems a worthwhile approach for managing enterprise healthcare authorization services, particularly those that justify an infrastructure approach because of size or complexity. The endorsement of Attribute Certificates by the IETF indicates that PMI may become a significant component in electronic commerce applications. An Attribute Certificate framework could be closely integrated with PKI and Directory services, and could fit well with many current efforts. The use of the Authorization API (aznAPI) is also an attractive component of a PMI, and could be used in conjunction with Attribute Certificates, if supported by vendors. The aznAPI would also allow very complex authorization decisions to be made by a rules engine, and a simple yes/no result returned to applications.

Because the products are at early stages of development, and the standards have not stabilized, current attempts at implementation would be risky, and require care. The growing interest and body of knowledge in the industry indicates that this is an architectural direction worth pursuing.

### *Recommendation*

- Healthcare organizations should encourage and become involved in further development of standards in PMI.
- Healthcare organizations should ask PKI and security product vendors about plans for PMI.

## **LONG TERM STORAGE**

Long term storage (greater than 50 years) of medical records poses several problems, especially in light of any potential litigation. (1) It is difficult to prove that an electronic record has not been altered. (2) Media deteriorates over time, destroying information and invalidating digital signatures. (3) Storage media technology changes over time, which

may render archives obsolete and unreadable. (4) Key lengths are generally determined by an estimated lifetime of “computational infeasibility” of ten to fifty years. This is not long enough for electronic medical records, or others which require long-term archival.”

*Recommendation*

- Create a process and standard that allows for digitally signed records to be copied onto new media (and new technology) on a periodic basis (e.g., every 10 years). The period should not be fixed, but be based on some cryptographic technology advancement index plus considerations for media deterioration. During this copy, the records should be “wrapped” or signed again with a key length that is appropriate for current technology. This period can be altered as the technology advances.