WP8 D8.4
PKI Challenge (pkiC) - Challenges for the PKI Industry

Authors: Kevin Blackman
David Chadwick

A two-year project

fully-funded by the European Commission and
the Swiss Government

run by EEMA as part of a 13-strong management consortium

to solve the current interoperability problems associated with the deployment of Public Key Infrastructure (PKI) technology

www.eema.org/pki-challenge

The pkiC Management Consortium: -

Baltimore + Belgacom + Royal Mail + EEMA + GlobalSign + KPMG + Makra Security & Standards + SmartTrust + University of Leuven + University of Salford + Utimaco + WISeKey

An EEMA Project

The European Forum for Electronic Business

EEMA

Funded by

European Commission

Bundesamt für Bildung und Wissenschaft
Office fédéral de l'éducation et de la science
Ufficio federale dell'educazione e della scienza
Uffizi federal da scolaziun e scienza
Federal Office for Education and Science

Swiss Government

Contents

1. Introduction

1.1. Workpackage Overview

The Best Practice Workpackage is one of the tasks of the pkiC that most deliberately attempts to meet a crucial demand of EU programmes, which is to translate conclusions and results from research work into practical advice and assistance for other regions and entities.

This paper attempts to perform this task for the PKI industry, which comprises standards bodies, the European Commission, users groups with an interest in this area, the PKI Forum and other participants. The definition of PKI industry must include manufacturers of PKI software and hardware, who are at the same time guiding the evolution of standards, but the reader should note that the Best Practice Guide for Vendors, which is another paper in this set, addresses the specific needs of this group.

The general objectives of this workpackage could be summarised as follows:

- To identify the best practice implications revealed by the pkiC interoperability trials and their associated survey work;
- To determine the challenges that the industry faces, giving practical advice on how to improve interoperability amid differing implementations of standards among pkiC products, each tailored to the needs of the different actors in the economy (citizens, enterprise users, corporations, government);
- To disseminate this information as widely as possible through identified mediums.

In order to do this the group

1. created draft structures for a set of best practice guidelines/recommendations and circulated them to all pkiC consortium members.
2. asked each partner to draft a set of initial ideas for inclusion in these guidelines.
3. assessed the initial ideas from the partners and produced a final set of guidelines/recommendations.

The workpackage calls for three deliverables, each targeted at a different audience:

1. **Guide for End Users**
   Intended for those who will be installing and running a PKI and expect to use it to do business with other companies who have their own PKIs.

2. **Recommendations to Vendors**
   Intended for anybody developing PKI / PKA software

3. **Challenges for the PKI Industry** (this document)
   Intended for Standards bodies, the European Commission (for future projects), Users groups with an interest in this area, PKI Forum etc.

1.2. Rationale for Identification of Best Practices

In order to clearly establish the rationale for this paper, it is necessary to propose a working definition of the concept of a *Public Key Infrastructure*, and thus establish a mutual understanding for the following sections.

1.2.1. Infrastructure

An infrastructure can be described as the basic facilities, services, and installations needed for the functioning of a community or society, such as transportation and communications systems, water and power lines etc. Of the examples, power and communications (TCP/IP, LAN etc.) infrastructures allow different entities to plug into them and use them on demand, when needed.

An infrastructure designed to meet the security needs of individuals, communities and societies must therefore recognise the same fundamental doctrine and benefits. A security infrastructure that underpins the needs of an organisation and their users must therefore be accessible by every application, object, or other entity that requires the security infrastructure's services. The infrastructure must have interfaces, access points, which are consistent and uniform. The security infrastructure avoids ad-hoc, proprietary, non-interoperable implementations, and introduces the possibility of manageable, consistent security across multiple applications and platforms.

1.2.2. PKI – A definition

A PKI by itself is not a security infrastructure, however it can form the basis upon which many of the essential features of a security infrastructure can be built and managed.

A PKI can thus be thought of as "the basis of a pervasive security infrastructure whose services are implemented and delivered using public-key concepts and techniques."[1]

A **PKI** that underpins a security infrastructure must offer a number of components and services:

- Certification Authority
- Certificate Repository
- Certificate Revocation
- Key backup and recovery
- Automatic Key Update
- Key history management
- Cross-certification
- Support for Non-repudiation
- Client Software
- Standard APIs and protocols that allow the components to interoperate with each other
- Support for Authentication, Integrity, and Confidentiality

[1] [ADAMS, LLOYD]

In addition to those mentioned above, a **Comprehensive PKI** would also offer the following additional components and services:

- Notarisation
- Secure Data Archive
- Secure Time Stamping
- Privilege / Policy Creation and Management tools

Such a comprehensive PKI does not currently exist, and no known product implements all of the above aspects. Certainly since PKIs are typically implemented to solve a particular problem, only a subset of the above services and features are necessary in a particular implementation.

1.2.3. Standards

There are many standards used in PKI that together offer a broad range of options for each aspect of the infrastructure. Users thus enjoy a wide range of choice, and flexibility to design an infrastructure that is able to provide a solution to their problems.

However the mere fact that so many options are available for each component or feature of a PKI means that the number of possible differing infrastructure implementations exponentially increases as more components are added to it. This increases the complexity of client and helper applications that need to be compatible with dissimilar, yet still standards compliant infrastructures serving different communities. This increases the implementation and support costs of the entire infrastructure.

By its very nature this problem is especially severe where interoperability is desired between PKIs serving different enterprises or communities. This interoperability can be as simple as ensuring trusted secure asynchronous communication between members of different communities, as in the case of secure e-mail; or as complex as supporting an extension of trust, and the ability to execute legally binding signed transactions between those same members, as in the case of qualified signatures.

This paper will concentrate mainly upon PKIs built around the X.509, IETF PKIX [PKIX], RSA Public Key Cryptography [PKCS] and related standards.

1.2.4. Rationale

Interoperability between PKI implementations that form the basis of security infrastructures - whether serving a loosely connected group of individuals, or a more cohesive and defined group, such as an enterprise, society, nation, community of nations - is an extremely important goal. The cost of failing to achieve this goal can be measured in real economic terms:

- Increased complexity in client and helper applications;
- Increased recurring support costs;
- The continued use of manual processes rather than migrating to automated electronic ones;
- Lost opportunity because of the failure to achieve communities of trust between borders, and to engage in legally binding electronic transactions of all kinds;

- Loss of confidence in the technology with the result that benefits take longer to realise.

In recognition of this goal this paper focuses on how the intended audience can help craft the future evolution of PKI standards to create a pervasive substrate, a Comprehensive PKI, which can be the basis of a Security Infrastructure serving a global community of users.

1.2.5. Structure of the paper

The paper primarily examines the experience of the pkiC project and identifies the major challenges facing the industry based on this experience, including referencing the other papers in the work package that provide more detail, and it makes broad suggestions that are meant to spur thought and provide some ideas, regardless of "that which inspires controversy", on serious steps towards achieving the promised benefits of PKI.

This paper does not pretend to be a comprehensive collation of all the challenges facing the PKI industry. An in-depth literature review of previously published material was not part of this work, and interested readers are encouraged to perform their own research in this respect.

2. Standards, standards, and more standards

The current standards relevant to PKI offer a broad range of options for each aspect of the infrastructure. The plethora of non-compatible options available for use by each component means that the number of possible differing infrastructure implementations grows exponentially as each new component is added. This increases the complexity of client and helper applications, which need to provide support for dissimilar, yet still compliant infrastructures serving different communities. To deal with this vendors have to either implement all of the different options, making their products configurable but large and expensive, or choose a subset of options for their range of products, making them smaller and cheaper to build, but at a loss of interoperability with other vendors' products.

The Best Practice papers mention a subset of standards that an entity implementing a PKI wishing to be interoperable with other PKIs in the present or future should consider. These standards have evolved through widespread acceptance by vendors and customers, and prescription by industry bodies. For further information on these standards please consult *pkiC Best Practice Paper I – Guide for End Users, and II – Recommendations to Vendors*.

This section considers some of the practical business problems that users encounter when attempting to use a PKI as the basis of a comprehensive interoperable security infrastructure.

2.1. Certificate Creation and Key/Certificate Distribution

The certificate request and retrieval process is addressed by a secure protocol mechanism. The IETF PKIX working group maintains a pair of competing specifications on the standards track that addresses this requirement in both online and offline modes:

- The Internet X.509 Public Key Infrastructure Certificate Management Protocols (CMP) [RFC2510]. and
- Certificate Management Messages over CMS (CMC [RFC2797])

Both of the above make use of

- The Internet X.509 Certificate Request Message Format (CRMF) [RFC2511]

whilst in addition CMC makes use of the Public Key Cryptography Standard (PKCS) 10 [RFC2986] and the Cryptographic Message Syntax standard [RFC 2630] (which is an enhancement of PKCS 7 [RFC2315]) message formats. Other proprietary protocols also exist, such as Cisco's Simple Certificate Enrolment Protocol.

Conclusions

CMP (in conjunction with CRMF) and CMC appear to be the two most comprehensive life-cycle management protocols among all the options. However the experience of the pkiC showed that there was poor support for these protocols by vendors, and interoperability between PKIs using these protocols is practically non-existent at this stage (except for the use of simple manual CMC).

Automatic certificate creation, enrolment, and renewal processes are essential to achieve economic efficiencies, and avoid heavy support costs.

Challenge

The security industry should consider consolidation and adopting a single set of standards that allow clients, regardless of the underlying operating system and network layers, to interface with an interoperable PKI in delivering certificate lifecycle management functions.

2.2. Certificate dissemination

After the generation of key pairs and certificates, a method must exist for disseminating the public key certificate. The need to retrieve an end-entity certificate can be driven from two separate requirements:

- The need to encrypt data for another end entity;
- The need to verify a digital signature received from another entity.

Typically a certificate is disseminated via one or more of the following methods:

- Physical out of band distribution (using storage media);
- The use of an on-line public database or repository (accessible via an established interface such as LDAP, DAP, HTTP);
- In-band protocol distribution (e.g. attaching the certificate to an outgoing SMIME message).

Verifying digital signatures is the easier of the two requirements to solve, since the certificate owner can distribute the certificate (chain) in-band along with the digitally signed message. The certificate owner thus presents the recipient with the certificates they need. Encrypting data is more difficult to solve, since the message originator must somehow first obtain the message recipient's certificate, perhaps without the certificate owner knowing that their certificate is needed. In very small communities the use of out-of-band or in-band certificate dissemination methods may be acceptable for encryption users. However any security infrastructure which wishes to scale to large communities, such as enterprises and the global Internet, must expose a consistent and reliable interface to its public certificate repositories to allow users to search for the certificates they require. Unfortunately, such an interface does not exist today.

It is a challenge for the PKI industry to establish a consistent and reliable interface for public key certificate distribution and retrieval, and this section explores this issue further.

Conclusions

LDAP

The first challenge the pkiC project faced was establishing a common repository format. Initial attempts were made to use the latest version of the most popular repository interface, LDAP v3. However since all participants were not up to date, LDAP v2 was accepted as the minimum requirement. This lack of support for LDAPv3 by all vendors meant that the pkiC was undertaken in the knowledge that it could be compromised by the long-established issues that the PKI industry has with the inconsistent use of LDAP directories by PKI software. Some of the issues with LDAP v2 (and LDAPv3) are:

- The mandatory-to-implement authentication mechanism between an LDAP v2 client and the repository is based on userid and password transmitted in the clear;
- A standard LDAP access control scheme does not exist;
- No standard mechanism for data replication between LDAP repositories exists;
- Search filters for certificates are undefined;
- The X.500 attributes used to store PKI items such as Certificates and CRLs have more than one accepted method of referencing them (i.e. either with or without the; binary description);
- No agreed upon signed operations capability exists;
- The X.500 attributes used to store X.509 objects have no consistent location within the Directory Information Tree and the precise structure and location is vendor dependant. This is a problem because LDAP servers cannot yet support searches for particular certificates and CRLs. Taken together, this means that it is highly unlikely that a vanilla implementation of PKI software from one vendor will be able to successfully query the LDAP directory of another;
- There are still a large number of interoperability problems caused by the structure and flexibility allowed in the Directory components, particularly, in this case, in the Distinguished Name. In the case of cross-certification, this causes difficulties where one vendor’s directory must support certificates with DN components that are not in its schema;
- LDAPv2 has no mechanisms to support a distributed directory, and LDAPv3 only has limited mechanisms (i.e. support for referrals);
- Given the above, the pkiC had a significant problem in locating the certificate repositories of the different PKIs. This issue causes problems in the real world as well as in the challenge interoperability testing, for both certificate and CRL

retrieval. The pkiC chose to use an LDAP redirector as a pragmatic solution. However this product is not useful in all real world situations, and as the existence of the redirector must still be made known to the client application via some method, problems were still encountered;

Certificate Profiling

The base criteria for pkiC were largely based on RFC2459, which was superseded by RFC3280 in April 2002. RFC3280 introduced extensions that allow the basic structure for X.509 certificates to provide information that allows relying software to locate the CA Issuers and the CA Repository. However use of these extensions, the Authority Information Access – CA Issuer method, and the Subject Information Access – CA Repository method were not in widespread use as yet.

It is clear that the minimum set of attributes defined in the pkiC was not enough to accommodate the major PKI systems that are currently available. It is equally clear that some of those systems also have problems supporting that minimal list.

Challenge

LDAP has been adopted as the defacto standard for certificate repositories, but currently suffers from a multitude of problems, chief among which are a lack of support for X.509 certificates, distributed directories and differing and incompatible directory schemas.

- LDAPv2 suffers from a number of problems that are universally recognised, and which are being addressed in LDAPv3. The core specifications of LDAPv3 include RFCs 2251-2256 and 2829-2830. LDAPv3 should be strongly recommended over LDAPv2. LDAPv3 includes support for stronger authentication mechanisms, secure transports, and referrals;
- An industry standard LDAPv3 Schema for X.509 certificates should be adopted by the security industry. Possible candidates include:
  - Klasen, N., Gietz, P. "An LDAPv3 Schema for X.509 Certificates", <draft-klasen-ldap-x509certificate-schema-01.txt>, March, 2003
  - D. W. Chadwick, S. Legg. “LDAP Schema and Syntaxes for PKIs”, <draft-ietf-pkix-ldap-pki-schema-00.txt>, June 2002;

RFC3280 introduced extensions that allow the basic structure for X.509 certificates to provide information that allows relying software to locate its CA Issuers and CA Repository. Interoperability testing using these extensions, the Authority Information Access – CA Issuer method, and the Subject Information Access – CA Repository method, between PKI communities and client applications should be pursued, and once successful widespread adoption should be encouraged.

All PKI systems should support a minimum set of components in the Distinguished Name (DN) (for more information see pkiC Best Practice Paper I – Guide for End Users).

Global Public Repository

The experience of the pkiC also raised the issue of locating a user's certificate in the absence of any other information than perhaps his name, company name, and/or email address. At the present time the first contact between correspondents leads to an exchange of email addresses, and then signed emails thus leading to an exchange of public key signature verification certificates, and possibly certificate chains up to trusted anchor CAs. From this point forward the CA's public repository can possibly be located, and the user's encryption certificate and the certificates of other correspondents in his community retrieved. This is possible via two ways, firstly if the certificates are held in the same repository as the CRLs pointed to by the CRL Distribution Point extension in the user's certificate, or secondly, if a certificate chain is available, by using the CA Repository access method of the Subject Information Access extension in the immediately superior certificate in the chain. Outside of such an exchange, there is no automatic method of finding a public repository for a company based on possibly well-known facts such as the company name or its domain name. (Note that in many cases companies do not want their certificate repositories to be public, so they are not concerned about this problem, and probably prefer certificate exchanges to be made via personal contact.)

There are a number of projects that attempt to address global directory services, and some specialise in establishing a network of public key servers, primarily for PGP (e.g. the PGP Key Server System), which even supports X.509 certificates. However most global directory service and universal communication identifier projects in planning are not considering support for public keys or certificates e.g. ENUM [ENUM]

Challenge
It is a challenge for the PKI industry to ensure that a global repository system for public key certificate distribution and retrieval is established. This challenge is compounded when one considers the large number of companies who might prefer their certificates to remain confidential. It is a further challenge to ensure that future global directory services that are currently being planned or built should support certificate distribution and retrieval.

2.3. Certificate validation

Certificate validation is the process by which the integrity of public-key certificates is verified, and includes determining the following:

- The certificate has been issued by a recognised trust anchor or its trusted subordinate, and this may include certificate path processing;
- The digital signature of the certificate is valid;
- The certificate is within its stated validity period;
- The certificate has not been revoked;
- The certificate is being used in a manner, which is consistent with its policy constraints, name constraints, and intended usage restrictions.

There are many possible methods of publishing certificate revocation information e.g. full CRLs, CARLs (Certification Authority Revocation Lists), EPRLs (End entity public key revocation lists), CRL Distribution Points, Delta and indirect CRLS, Indirect CRLs, OCSP, Redirect CRLs, and SCVP.

Conclusions

pkiC testing found that the CRL Distribution Points and the Authority Information Access OCSP Access Method extensions are essential elements of an end-entity certificate, as these can be used by relying parties to obtain certificate revocation information. Furthermore, the Subject Information Access attribute is an essential element of a CA's certificate, as this can be used to locate the CA's repository for CRLs and certificates.

pkiC testing did not test for support of the Authority Information Access CA Issuers access method which retrieves information about CAs that are superior to the issuer of the certificate. Support for this extension can resolve some of the problems, which affected certificate path construction and path validation.

Challenge

Industry and regulatory certification bodies should promote greater use and support of the following services and extensions in end-entity and CA certificates:

- CRL Distribution Points;
- Authority Information Access – CA Issuers, and OCSP services.
- Subject Information Access – CA Repository access method should be a mandatory extension.

3. Broad Conclusions

Standards and Certificate Profiling

One of the implications from the pkiC interoperability testing project is that the existing standards are far too complex, and have far too many options, to ensure that different vendors can build fully interoperable systems from them. Profiling the standards would seem to be essential. But even then, this may not be enough. The base pkiC interoperability specification [D2.2] was a simple profile of the PKI related standards, but even this proved to be too complex for any of the vendors to fully test against e.g. no CMP testing was done at all. Given that the EESSI specifications add yet another layer of complexity to the existing international standards, one must question if these specifications will ever be implemented in their entirety, if at all. It would appear that currently the cost is too great, and the market is too small to make this economically feasible. The implication for the EU is that fewer and simpler, rather than more and complex, standards are needed.

The EU should consider a policy of more stringent profiling, with a much greater level of detail than is currently employed e.g. specifying exactly the bits that should be contained in fields such as the subject key identifier, rather than simply saying that a field should be present, and mandating the usage of the Authority Information Access and Subject Information Access extensions for public certificates.

Directory Services

In a similar vein, the problems that are experienced with directory services, would indicate that more stringent requirements should be placed on directory names and DIT structures, rather than allowing the current “anything goes” policy, which is a major cause of interoperability problems.

Furthermore there are strong indications that any Public Key Infrastructure will not achieve its lofty title unless a global directory public key/certificate service, integrated with Internet protocols, is established. Furthermore, support for certificate storage, retrieval and searching must be part of the Internet’s evolving future global directory services solutions, before they become widely implemented.

Interoperability Testing

What has been clearly borne out by the pkiC is that the process of testing implementations and refining the standard specifications to produce interoperable implementations should be more of an integral part of the standardisation process. The support of trials such as pkiC early in the development process of standards would greatly improve the quality of the resulting standards and greatly improve the chances of suppliers developing products that are truly interoperable. Thus, it is suggested that the importance of trials such as the pkiC is brought to the attention of EU policy makers and that suppliers are encouraged to participate in such trials as early as possible in the standardisation process.

4. References

Directive 1999/93/EC of the European Parliament and of the Council (PDF)
http://europa.eu.int/comm/internal_market/en/media/sign/Dir99-93-ecEN.pdf
Source: European Commission
Includes: Full text of the Electronic Signature Directive

Electronic Signature Directive Overview
http://europa.eu.int/comm/internal_market/en/media/sign/99-915.htm
Source: European Commission
Includes: Brief overview of Electronic Signature Directive components

[EESSI] The European Electronic Signature Standardization Initiative
http://www.ict.etsi.org/eessi/EESSI-homepage.htm
Source: EESSI
Includes: Information on the EESSI, a standards group established by the European ICT Standards Board with the support of the European Commission

[ADAMS, LLOYD] Understanding PKI, Concepts, Standards and Deployment Considerations 2003
Addison Wesley
Carlisle Adams, Steve Lloyd

[D2.2] WP2 N013 – Interoperability Test Criteria (D2.2)
pkiC Interoperability Test Criteria Specifications
https://www.eema.org/pki-challenge/criteria.asp

[ENUM] E.164 number and DNS
http://www.ietf.org/rfc/rfc2916.txt

[RFC3280] Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List (CRL) Profile
http://www.ietf.org/rfc/rfc3280.txt

[RFC2459] Internet X.509 Public Key Infrastructure Certificate and CRL Profile
http://www.ietf.org/rfc/rfc2459.txt

[PKIX] The Internet Engineering Task Force Public Key Infrastructure (X.509) Working Group
http://www.ietf.org/html.charters/pkix-charter.html

[PKCS] Public-Key Cryptography Standards
http://www.rsasecurity.com/rsalabs/pkcs/

[X509] ITU-T X509 Recommendation
Information technology - Open Systems Interconnection - The Directory: Public-key and attribute certificate frameworks
http://www.itu.int/rec/recommendation.asp?type=folders&lang=e&parent=T-REC-X.509

[ANSIX9F1] ANSI X9F1 Standards Projects

5. Definitions

Subject Information Access a private extension used in end-entity and CA certificates, which indicates how information and services offered by the subject in the certificate can be obtained. Access methods that have already been defined are: for CAs - CA Repository, and End-entities - time stamping. The CA Repository access method identifies the location of the repository where the CA publishes certificate and CRL information. Other access methods may be defined in the future. [RFC3280]

Authority Information Access The authority information access extension indicates how to access CA information and services for the issuer of the certificate in which the extension appears. Information and services may include on-line validation services and CA policy data. RFC 3280 defines two access methods: CA Issuers and OCSP.
id-ad-caIssuers lists CAs that have issued certificates superior to the CA that issued the certificate containing this extension. The referenced CA Issuers is intended to aid certificate users in selecting a certification path that terminates at a point trusted by the certificate user.
The id-ad-ocsp OID is used when revocation information for the certificate containing this extension is available using the Online Certificate Status Protocol (OCSP) [RFC2560]

PKI Public Key Infrastructure

PGP Pretty Good Privacy

ENUM E.164 number and DNS, is a protocol that is the result of work of the Internet Engineering Task Force's (IETF's) Telephone Number Mapping working group. The charter of this working group was to define a Domain Name System (DNS)-based architecture and protocols for mapping a telephone number to a Uniform Resource Identifier (URI), which can be used to contact a resource associated with that number. (http://www.ietf.org/rfc/rfc2916.txt)