

The development of public key infrastructures; Are we on the right path?¹

Audun Jøsang*, Hein Meling* and Manyi Lu**

*Norwegian University of Science and Technology

**SINTEF Informatics

October 1999

Abstract

Public key cryptography allows a user to digitally sign messages with a private signature key so that recipients in possession of the corresponding public verification key can check the correctness of the signature. The main problem with this method is to authenticate public keys in order to know that the digital signature is authentic. A Public Key Infrastructure refers to a network where the authenticity of public keys is certified by Certification Authorities in a hierarchy, and it is believed that this will be an important element in the development of electronic commerce. Certification services are already available from commercial Certification Authorities, and we are starting to see the emergence of national and international public key infrastructures. This paper describes public key infrastructures in general and discusses Norwegian as well as some international implementation efforts.

1 Introduction

Internet was originally designed with a closed community of researchers in mind and misuse of the Internet services was not thought to be a problem. Now that Internet is a global and open network with several hundred million users connected many types of misuse have emerged.

These problems calls for the use of secure communication protocols that enable Internet users to establish trust in organisations as well as in other users. Several frameworks and schemes has been suggested for providing authentication and secure communication, mainly through the use of public key infrastructures (PKI) and certification authorities (CA). These frameworks will likely pave the way for electronic commerce.

Several commercial CAs already provide customers with signed public keys, with varying degree of authentication from email verification to physical appearance verification.

2 Public Key Certification

Digital signatures are based on a pair of mathematically linked cryptographic keys called the private key (or signature key) and the public key (or verification key) with the property that the private key can not be deduced from the public key. The private key is private and confidential so that only the user in possession of the key can generate the digital signature. The corresponding public key allows anyone to verify that a digital signature produced by the private key is correct, i.e. that indeed the private key was used in the signature process. However, a digital signature must not only be correct, it must also be

¹Appeared in the Proceedings of the Norwegian Informatics Conference (NIK'99), Trondheim, Norway, 1999

authentic and this is where certification comes in. The crucial point is how to know that the owner of public key really is whom he claims to be, or in other words to be certain about the authenticity of the public key. One must not forget that Internet users can not base their judgements on faces or familiar voices, and that electronically received public keys a priori can not be trusted.

2.1 Certification Hierarchies

The problem of authenticating electronically received public keys can be solved by a hierarchy of CAs with mutual trust relationships. Alternatively an anarchic web of trust relationships such as in PGP [Zim95] can be used, but in general a hierarchic solution based on the hierarchic X.500 Directory [ITU93] and on X.509 Certificates [ITU97] will be required.

Certification between nodes is directed. As such one possibility is to have a hierarchy that is strictly top down, as illustrated in Fig. 1.




Figure 1: A strict hierarchy

A strict hierarchy can in principle not be used for user-to-user authentication when it is assumed that the user only knows the public key of its local CA. That is because there exists no certification path from the user upwards to the root. However, if it can be assumed that the user always knows the public key of the top CA, the certification path to other users can always start at the root.

Another and better possibility is to have a general hierarchy that includes two-way certification between CAs, as illustrated in Fig. 2.




Figure 2: A general hierarchy

When certification takes place in both upwards and downwards direction, each user will only need to obtain an authentic copy of a single CA's public key, while still being able to establish a certification path to every other user in the network.

The opposite of a hierarchic structure is an anarchic structure where each CA (and user) is free to choose which other CAs (and users) it wants to certify, as illustrated in Fig. 3.




Figure 3: Anarchic structure

The anarchic structure corresponds to the web of certification on which PGP [Zim95] is based. The disadvantage of an anarchic certification network compared to a hierarchic structure is that there exists no simple algorithm for identifying certification paths between all globally identifiable users of an anarchic network, whereas such algorithms exist for hierarchic networks.

A general hierarchy with two-way certification is thus the best way to establish certification paths from user to user, and to have a truly open PKI. In general, a PKI can consist of several commercial certification hierarchies which are interlinked either in a hierarchic way or by cross certification. CAs on the top plane should all be cross certified so that no CA alone can represent the top root as illustrated in Fig. 4.




Figure 4: Cross certification between hierarchies

In the X.509 framework, a certification authority produces the certificate of a user by signing a collection of information, including the user's distinguished name and public key, as well as an optional unique identifier containing additional information about the user. For example the public key certificate of a user with distinguished name Y and unique identifier UY produced by the certification authority with name X and unique identifier UX , has the following form:

$$X \ll Y \gg = X \{V, SN, AI, X, UX, Y, UY, Y_p, T^Y\} \quad (1)$$

where V is the version of the certificate, SN is the serial number of the certificate, AI is the identifier of the algorithm used to sign the certificate, UX is the optional unique identifier of the certifier X , UY is the optional unique identifier of the user Y , Y_p is the public key of Y , T^Y indicates the period of validity of the certificate. The signature in the

certificate can be checked for validity by any user with knowledge of X 's public key X_p , and thereby obtain an authenticated copy of Y 's public key Y_p . This process is denoted by

$$Y_p = X_p \bullet X \ll Y \gg. \quad (2)$$

In order for two users to verify the authenticity of each others public keys it is sufficient that there exists a certification path between them. We will use an example from the X.509 standard to illustrate this. Fig. 5 ([ITU97] Figure 4) illustrates a fragment of a Directory Information Tree (DIT) as defined in [ITU93] where the CAs form a hierarchy. Besides the information shown in the boxes we assume that each user knows the public key of its certification authority, as well as its own public and private keys. In particular it can be seen that each CA has verified and certified the node above and below.

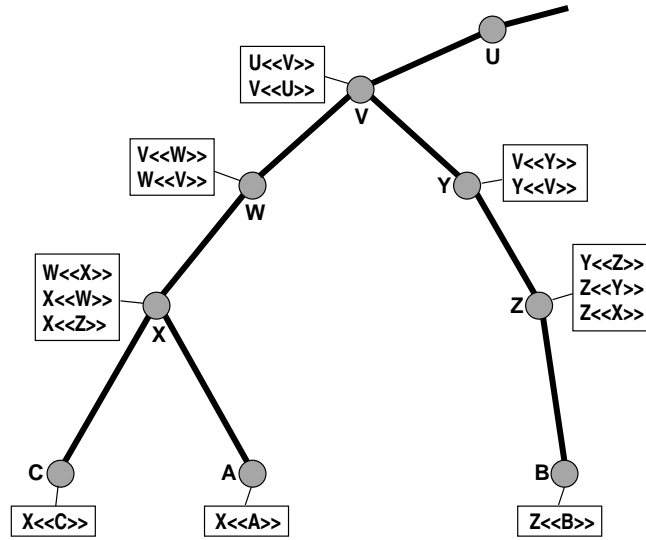


Figure 5: CA hierarchy [ITU97]

If for example A wants to verify the authenticity of B 's public key he must first acquire the certificates from the CAs in the hierarchy according to the principle described in the X.500 Directory Standard [ITU93]. The following certificates must be acquired: $X \ll W \gg$, $W \ll V \gg$, $V \ll Y \gg$, $Y \ll Z \gg$, $Z \ll B \gg$.

The chain of certificates can be resolved by the principle of Eq.(2) to obtain an authenticated copy of B 's public key:

$$B_p = X_p \bullet X \ll W \gg W \ll V \gg V \ll Y \gg Y \ll Z \gg Z \ll B \gg \quad (3)$$

A public key certificate is nothing more than the CA's digital signature on a user's (or another CA's) public key. In order to verify that the CA's certificate is correct, the recipient must obtain an authentic copy of the CA's public key. A crucial problem is thus how the user receives the CA's public key. Another problem is how users (or CAs) store their private keys. This will be discussed in the sections below.

2.2 Propagating Trust

For the distribution of public keys in open networks it is not conceivable to have a single global authority that is trusted for key generation and distribution because there will

always be different administrative domains which typically will have conflicting economical and political interests. In this situation, each agent has to decide for herself which other CAs she wants to trust for key distribution and certification, and based on this determine the legitimacy of received certificates and the authenticity of keys.

Technically seen, humans do not sign cryptographic certificates, keys do. However, it is usually assumed that human agents are using cryptographic keys as a tool to make certificates so that practically speaking humans do sign certificates. For this assumption to be correct it is essential to explicitly express trust in the binding between the key used for certification and it's owner, because failing to do so would deprive any authentication scheme of it's relationship to humans, and would turn the scheme into authentication for and by keys. The key-to-owner binding can not be objectively assessed, and necessarily becomes a subjective measure, meaning that two individuals can have different opinions about any particular binding.

To have established the binding between a key and it's owner is not enough for accepting certificates produced by it if for example the key owner deliberately certifies flawed keys. Another essential element is therefore to consider the trustworthiness of the certifying CA itself for the purpose of recommending keys by certification. As for the key-to-agent binding, the CA trustworthiness also becomes a subjective measure, meaning that an agent who is trusted by me does not have to be trusted by you.

In [Jøs96] we argued that trust simply is a human belief, involving a subject (the trusting party) and the object (the trusted party). Trust in the key-to-owner binding can for example be expressed as believing that: "the key is authentic", whereas trust in the certifier is to believe that "he will only certify keys that he considers authentic".

Levels of trust in the key-to-agent binding is to some degree captured by the concept of certification classes described in the next section, but present PKI implementation efforts do not consider CA trustworthiness at all, they simply assume that they can be fully trusted. These issues are presently the subject of much research, see e.g. [RS97, Jøs98, JY98, Jøs99].

3 PKI for the World Wide Web

Commercial Public Key Infrastructures for the World Wide Web such as for example VeriSign's PKI [Ver97] are strict hierarchies, and this will be briefly described below.

3.1 VeriSign's PKI

VeriSign provides a certification framework where certification only is done in a downwards direction. Applied to the example of Fig. 5 the Certification Authority X is for example not supposed to issue the certificate $X \ll W \gg$. The consequence of this is that it is only possible to establish a certification path from user to user if everyone knows the public key of VeriSign's root CA. What then can a PKI such as VeriSign's be used for?

The most common applications of certificates in the Internet is presently for establishing encrypted connections between web browser and server, and for certifying active Web components such as Java applets and Microsoft's ActiveX components.

Encrypted connections are needed to protect confidential information, such as credit card numbers, sent from the user to the web server. This is currently implemented with the SSL protocol [Net95] which uses the web servers public key for establishing a secret session key between web browser and server. The server's public key has been certified by a CA.

Regarding active components, the security problem users are facing is whether such imported programs can safely be executed. One way this can be solved in web browsers is that the components are digitally signed by the software component manufacturer's public key which previously has been certified by a CA.

Users who want to verify that the software manufacturer is authentic must verify the authenticity its public key, and this requires an authentic copy of public key of the CA that certified the SW manufacturer. This problem is solved by pre-installing and shipping the public keys of the most well known CAs with the web browser when it is installed. However, this introduces an additional element that also needs to be trusted; namely the manufacturing and distribution process of commercial software. We feel that there is too little discussion around this point, and that this kind of 'certification by plastic shrink-wrapping' can have serious weaknesses.

Fig. 6 illustrates the VeriSign PKI hierarchy and the subdivision into classes, which will be discussed in Section 3.2.

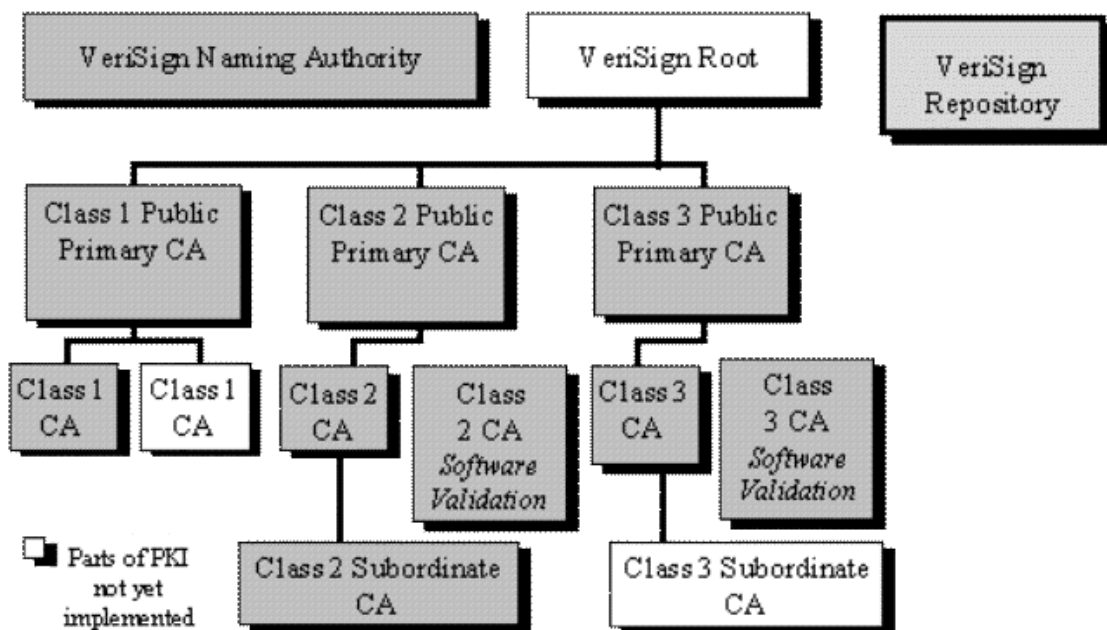


Figure 6: VeriSign PKI hierarchy [Ver97]

In VeriSign's certification scheme all public keys are certified from the top of the hierarchy through the intermediate CAs' public keys to the bottom users' public keys. However, the top root public key is not certified by anyone, which causes the whole hierarchy to hang in the void. In VeriSign's Certification Practice Statement it is for example stated that root public keys should be "self-signed" ([Ver97] Sec.2.5.1) but this is in our view meaningless.

3.2 Certification Classes

It is absolutely mandatory that the CA verifies the identity of the owner of the public key before issuing a certificate. Failing to do so will undermine the whole purpose of the certificate: It shall certify that the owner of the certified public key really is the user associated with the identity specified as part of the certificate. There are different methods for verifying the owner identity, each associated with a particular trust policy, as will be described below.

Most commercial CAs issue different classes of certificates depending on how the identity of the owner is verified. The purpose of having different certification classes is to give a hint regarding how much a certificate should be trusted, or in other words how strongly one should believe that the owner of a public key really is who she claims to be.

VeriSign uses 3 different classes, where Class 1 represent the lowest assurance level and Class 3 the highest. The description of VeriSign's certification classes below is taken from [Ver97], Ch.2.

- **Class 1** certificates confirm that a user's name (or alias) and E-mail address form an unambiguous subject within the VeriSign repository. Class 1 certificates are communicated electronically to subscribers and added to his or her set of available certificates.

Class 1 certificates do not facilitate the authentication of the identity of the subscriber. Rather, they merely represent a simple check of the non-ambiguity of the subject name within the VeriSign repository, plus a limited verification of the E-mail address. The subscriber's common name (and, when submitted, Registration Field Information) contained in a Class 1 certificate is considered non-verified subscriber information (NSI).

- **Class 2** certificates confirm that the application information provided by the subscriber does not conflict with information in well-recognized consumer databases.

Class 2 certificates may provide reasonable, but not foolproof, assurance of a subscriber's identity, based on an automated on-line process that compares the applicant's name, address, and other personal information on the certificate application against widely referenced databases. Confirmation is based upon VeriSign proprietary matching criteria of third-party databases against the information in the application.

- **Class 3** certificates provide important assurances of the identity of individual subscribers by requiring their personal (physical) appearance before a Class 3 Local Registration Authority (LRA) or its delegate (such as a notary).

Class 3 certificate processes utilize various procedures to obtain probative evidence of the identity of individual subscribers. These validation procedures provide stronger assurances of an applicant's identity than Class 2 certificates. The practical uses and reliability of Class 3 certificates are bolstered by utilizing notaries (an existing, important, and legally-recognized authentication process). For business entity Class 3 certificates, the requirement for "out-of-band" communication with the business organization and confirmation of business entity information via third parties provide further assurance of trustworthiness.

Only Class 3 certification thus requires physical appearance for providing evidence about a person's identity, and is the only class that can provide high assurance regarding a certificate's correctness. At the same time this is also the most expensive form of providing evidence. Organisations with a large network of offices, such as banks and national postal services, are in a good position to perform this type of evidence checking. New Internet companies are able to perform Class 1 and 2 verification, but usually lack the infrastructure for Class 3 identity verification.

4 The Development of PKI in Norway

In Norway, Posten SDS¹ and FellesData² started providing PKI services 3-4 years ago. However, their solutions were proprietary with no inter-working possible. Recently, Telenor Nett³ and UNINETT⁴ have also started operating PKIs. The question is whether they can agree on cross certification and inter-working.

4.1 Forvaltningsnett

The Public Administration Network Project *Forvaltningsnett*⁵ has specified standards and policies for PKIs to be used by all public service organisations, and one of the essential criteria was that all PKIs must be cross certified. In September 1999, a contract was signed with Telenor Nett and Posten SDS for providing PKI services, and these two companies are required to establish cross certification by December 1999. Other players which are able to comply with the specifications will be able to join later. Forvaltningsnett requires private keys to be stored on smart cards, and has presently only specified one certification class equivalent to VeriSign's Class 3, i.e. requiring physical appearance before a certificate can be issued.

4.2 Telenor Nett and Posten SDS

The PKI services offered by Telenor Nett and Posten SDS depend to a certain degree on the SW solutions on which their respective PKIs are based. Presently Telenor Nett uses Entrust's⁶ SW whereas Posten SDS uses iD2's⁷ SW. The main difference between Telenor Nett's and Posten SDS's solutions is that Telenor Nett first issues an authorisation code (like a one-time password) to the user which is used to generate the private key and the certificate, whereas Posten SDS issues the private key and certificate directly.

Both Telenor Nett and Posten SDS store certificates on their respective X.500 directory servers. As mentioned the users' private keys can be stored in smart cards, but they also provide solutions where the private keys can be stored on disk encrypted under a password. Forvaltningsnett only allows the former solution.

Both Telenor Nett and Posten SDS issue smart cards with the users private key and with the CA's public key stored securely on the card. In that way, the user is able to authenticate the CA's public key during the hand-over of the smart card.

Presently, the respective PKI hierarchies of Telenor Nett and Posten SDS are flat, that is, there is in principle no distinct root CA, so the question of top-down or two-way certification between CAs is irrelevant. However, this may change in the future, and the only viable solution is then to implement a general hierarchy as illustrated in Fig. 2.

Telenor Nett's PKI is presently not linked to any superior CA whereas Posten SDS is connected to the root CA of the Universal Postal Union. This will allow users of Posten SDS to establish certification paths with users of other national post organisations around the world.

Telenor Nett is in a good position to issue Class 2 certificates to the general public due to customer relationship and large customer databases. Posten SDS is in a good position to issue Class 3 certificates to the general public due to the large network of post offices.

¹Posten SDS is the computer services branch of the Norwegian national postal service

²FellesData is a company providing computer services mainly to banks in Norway

³Telenor is the main public telephony operator in Norway

⁴UNINETT provides computer network services to universities and research organisations in Norway

⁵Information about Forvaltningsnett at: <http://www.forvaltningsnett.dep.no>

⁶Information about Entrust at <http://www.entrust.com>

⁷Information about iD2 at <http://www.id2.se>

4.3 FellesData

PKI services offered by FellesData are mainly aimed at the banking and financial institution market in Norway. Over 50 savings banks have already become subscribers of its PKI services.

Currently, FellesData operates three self-certified CAs which cover three different business fields. As the operation grows, FellesData will consider building a hierarchy of CAs.

FellesData aims at participating in the Forvaltningsnett project. As a result of this, cross-certification with Posten SDS and Telenor Nett is expected to be implemented by Spring year 2000.

There are two methods for users to obtain their keys, certificates and necessary software. Both methods are considered to be equally secure. The first method is through registered mail. In this case, subscriber has to show their ID when collecting the package from the post office. The second method requires subscribers' personal appearance at the Local Registration Authority.

FellesData operates a X.500 directory with LDAP support, through queries to the directory, certificates of subscribers and CAs may be easily obtained. In addition, the public keys of CAs are distributed together with the software package to subscribers.

4.4 UNINETT

UNISA is service offered by UNINETT for the exchange of public cryptographic keys.

Similarly to the certification link between Posten SDS and UPU, the root CA of UNINETT is certified by the ICE-TEL Certification Authority. ICE-TEL is supposed to be the top-level certification authority for the European Public Key Infrastructure. At the moment, this infrastructure comprises a network of the following countries: Denmark, Germany, Italy, Norway, Slovenia, Spain, and United Kingdom. Domestically, UNINETT currently does not support cross-certification with other PKI operators.

UNISA implements a top-down certification hierarchy. The public key of the top level CA is delivered to each subscriber, and it resides in the confidential area of each subscriber's local machine. In addition, the top-level authority is required to widely publish its public key, and this is in practice the way users can verify the authenticity of the CA's public key.

Each user is responsible for the confidentiality of his private key. In UNISA system, a user needs to store sensitive information such as private key in a confidential area in his computer. Information in this area will be encrypted and in order to access the confidential area, the corresponding password which may be the same as the log-in password, has to be given.

The certificate issuing policy of UNINETT requires subscribers' personal appearance at the Local Registration Authority and providing necessary identification. It also suggests a much simpler procedure for issuing of certificates collectively, to for instance students.

5 Conclusion

We have described design issues for public key infrastructures such as top-down or two-way certification, and the use of trust relationships. Further we have described how web browsers use public key infrastructures to authenticate active components and to estab-

lish confidential sessions between browser and server. Finally we have briefly described Norwegian efforts for establishing public key infrastructures.

It must be recognised that computer networks are global, and a PKI will only get its full potential if it is connected to a global infrastructure. Cross certification between CAs will allow users to choose which CA they want to use, while still being able to authenticate any other user on the planet as long as he or she is connected to the same global hierarchy.

An issue which has received little attention in present PKI implementations is the problem of assessing trust through certification chains. This becomes a complex problem with cross certification between CAs and with hierarchies crossing cultural and political borders. This is something that PKI providers will have to take seriously in the future.

References

- [ITU93] ITU. *Recommendation X.500, Data Communication Network Directory (also known as ISO/IEC 9594: Information Technology - Open Systems Interconnection - The Directory)*. International Telecommunications Union, Telecommunication Standardization Sector (ITU-T), 1993.
- [ITU97] ITU. *Recommendation X.509, The Directory: Authentication Framework (also ISO/IEC 9594-8, 1995)*. International Telecommunications Union, Telecommunication Standardization Sector (ITU-T), June 1997.
- [Jøs96] A. Jøsang. The right type of trust for distributed systems. In C. Meadows, editor, *Proc. of the 1996 New Security Paradigms Workshop*. ACM, 1996.
- [Jøs98] A. Jøsang. A Subjective Metric of Authentication. In J. Quisquater et al., editors, *Proceedings of ESORICS'98*, Louvain-la-Neuve, Belgium, 1998. Springer.
- [Jøs99] A. Jøsang. An Algebra for Assessing Trust in Certification Chains. In J. Kochmar, editor, *Proceedings of the Network and Distributed Systems Security Symposium (NDSS'99)*. The Internet Society, 1999.
- [JY98] M. Jakobsson and M. Yung. On assurance structures for WWW commerce. In *Proceedings of Financial Cryptography 98*, 1998.
- [Net95] Netscape. *The SSL (Secure Sockets Layer) 3.0 Protocol*. Netscape Communications Corp, 1995.
- [RS97] Michael K. Reiter and Stuart G. Stubblebine. Toward acceptable metrics of authentication. In *Proceedings of the 1997 IEEE Symposium on Research in Security and Privacy*, Oakland, CA, 1997.
- [Ver97] Verisign. *Verisign Certification Practice Statement (CPS), Version 1.2*. Verisign Inc, 1997. URL: <http://www.verisign.com/repository/CPS1.2/>.
- [Zim95] P.R. Zimmermann. *The Official PGP User's Guide*. MIT Press, 1995.