

Introduction to Public Key Infrastructure (PKI)

PKI is a security architecture that has been introduced to provide an increased level of confidence for exchanging information over an increasingly insecure Internet.

This guide provides the reader with a basic introduction to key terms and concepts used in a PKI including Certificates, Keys and Authorities. It mentions the features and services used by the PKI and the techniques involved in public key cryptography.

1.1 Introduction

The term PKI can be very confusing, even to a technologist, because it is used to mean several different things. On the one hand PKI may mean the methods, technologies and techniques that together provide a secure infrastructure. On the other hand, it may mean the use of a public and private key pair for authentication and proof of content.

A PKI infrastructure is expected to offer its users the following benefits:

- certainty of the quality of information sent and received electronically
- certainty of the source and destination of that information
- assurance of the time and timing of that information (providing the source of time is known)
- certainty of the privacy of that information
- assurance that the information may be introduced as evidence in a court or law

These facilities are delivered using a mathematical technique called public key cryptography, that uses a pair of related cryptographic keys to verify the identity of the sender (signing) and/or to ensure privacy (encryption).

PKI facilities have been developed principally to support secure information exchange over insecure networks - such as the Internet - where such features cannot otherwise be readily provided. PKI facilities can, however, be used just as easily for information exchanged over private networks, including corporate internal networks. PKI can also be used to deliver cryptographic keys between users (including devices such as servers) securely, and to facilitate other cryptographically delivered security services.

1.2 How the Public Key Cryptography concept works

Public-key cryptography uses a pair of mathematically related cryptographic keys. If one key is used to encrypt information, then only the related key can decrypt that information. If you know one of the keys, you cannot easily calculate what the other one is. As a result, in a 'public key system' you have the following:

- A public key. This is something that you make public - it is freely distributed and can be seen by all users.
- A corresponding (and unique) private key. This is something that you keep secret - it is not shared amongst users. Your private key enables you to prove, unequivocally, that you are who you claim to be.

The next sections describe how these keys are used in practice. They consider using separate public and private keys for encryption than those used for signing.

The Public Key used for Encryption

Another person uses your public encryption key when they want to send you confidential information. The information to be sent is encrypted using your public key*. You can provide your public key to the sender, or it can be retrieved from the directory in which it is published.

* Note: In normal practice, the actual information being sent is encrypted using a secret key algorithm (symmetric cryptography). Symmetric algorithms are much faster than public/private key algorithms (asymmetric cryptography). A random key (the session key) is generated, and it is used with the symmetric algorithm to encrypt the information. The public key is then used to encrypt that key and both are sent to the recipient.

The Private Key used for Decryption

A private key is used to decrypt information that has been encrypted using its corresponding public key. The person using the private key can be certain that the information it is able to decrypt must have been intended for them, but they cannot be certain who the information is from.

* Note: In normal practice the private key is used to decrypt the session key, and that key is used to decrypt the actual information rather than the private key decrypting all the information.

The Private Key for Signature

If the sender wishes to prove to a recipient that they are the source of the information (perhaps they accept legal responsibility for it) they use a private key to digitally sign a message (a digital signature). Unlike the handwritten signature, this digital signature is different every time it is made. A unique mathematical value, determined by the content of the message, is calculated using a 'hashing' or 'message authentication' algorithm, and then this value is encrypted with the private key – creating the digital signature *for this specific message*. The encrypted value is either attached to the end of the message or is sent as a separate file together with the message. The Public Key corresponding to this private key may also be sent with the message, either on its own or as part of a certificate.

* Note: Anyone receiving information protected simply by a digital signature can check the signature and can read and process the information. Adding a digital signature to information does not provide confidentiality.

The Public Key for Signature

The receiver of a digitally signed message uses the correct public key to verify the signature by performing the following steps. A non-technical example is given after these steps.

1. The correct public key is used to decrypt the hash value that the sender calculated for the information
2. Using the hashing algorithm (where certificates are in use it will be stated in the public key certificate sent with the message), the hash of the information received is calculated
3. The newly calculated hash value is compared to the hash value that the sender originally calculated. This was found in step 1 above. If the values match, the receiver knows that the person controlling the private key corresponding to the public key sent the information. They also know that the information has not been altered since it was signed

4. If a public key certificate was sent with the information it is then validated with the CA that issued the certificate to ensure that the certificate has not been falsified and therefore the identity of the controller of the private key is genuine
5. Finally, if one is available, the revocation list for the CA is checked to ensure that the certificate has not been revoked, or if it has been revoked, what the date and time of revocation were.

As an example, suppose you are sent a Word document by e-mail. The sender has signed it by calculating a hash value for that Word document, and then encrypted that value with their private key. You receive the Word document, and calculate the hash value for it. You decrypt the hash value that the sender encrypted and compare the two. If they are equal, the document hasn't changed and you are certain who sent the document. (If they don't match you know that the document has changed or the sender is not who they claimed.)

If no errors have been found, the receiver can now be certain of the authenticity and accuracy of the information that has been received.

The following table summarizes who uses public and private keys and when:

Key Function	Key Type	Whose Key Used
Encrypt data for a recipient	Public key	Receiver
Sign data	Private key	Sender
Decrypt data received	Private key	Receiver
Verify a signature	Public key	Sender

To encrypt information that will be stored for your own use (that is, you will be the only person able to read it), you must use your own Public Key as the recipient's key (you are the recipient) in order to be able to decrypt and read the information later. If you use someone else's Public Key, then only they will be able to decrypt and read the information. (To avoid the difficulty associated with not being able to read encrypted messages if you are not one of the recipients, some systems do not delete the original message after encryption whilst others store a copy of the key used for encryption either under the sender's Public Key or under a System Recovery Key. These latter methods are also referred to as key escrow or key recovery.)

Public Key Cryptography is therefore used for the encryption/decryption and signing/verification of information. Encrypting information ensures privacy by preventing unintended disclosure, and signing messages authenticates the sender of the message and ensures the message has not been modified since it was sent. It has to be remembered that only the information signed/encrypted has been protected. Commonly in e-mail systems headers, addresses and body messages may have no protection at all and should not be considered secure or part of the protected information.

1.2.1 The certificate

In the section on public and private keys, references were made to certificates. A certificate is information referring to a public key, that has been digitally signed by a Certification Authority (CA). The information normally found in a certificate conforms to the ITU (IETF) standard X.509 v3. Certificates conforming to that standard include information about the published identity of the owner of the corresponding private key, the key length, the algorithm used, and associated hashing algorithm, dates of validity of the certificate and the actions the key can be used for.

A certificate is not essential to operating a PKI, however, some scheme is necessary to locate information about the controller of a private key, and the X.509 certificate is the most commonly implemented scheme.

1.2.2 Controlling Key Usage

One of the fields in a public key certificate (certificate) is the key usage field. It is used by the CA to state the uses the CA has approved. It does not mean that the corresponding private key cannot be used in any other ways. There is no certificate with a private key. People receiving information protected using a public key system should check, where a certificate is provided, that the key usage stated in the certificate corresponds to the actual use.

1.3.2 Storing methods for Public and Private Keys

Certificates

Public keys are stored within digital certificates along with other relevant information (user information, expiration date, usage, who issued the certificate etc.). The CA enters the information contained within the certificate when it is issued and this information cannot be changed. Since the certificate is digitally signed and all the information in it is intended to be publicly available there is no need to prevent access to reading it, although you should prevent other users from corrupting, deleting or replacing it.

Protection

If someone gains access to your computer they could easily gain access to your private key(s). For this reason, access to a private key is generally protected with a password of your choice. Private key passwords should never be given to anyone else and should be long enough so that they are not easily guessed. This is the same as looking after your ATM CARD and its PIN. If someone manages to get hold of your card then the only thing that prevents him or her using it is the PIN (password) protecting it. If someone has your PIN then they can take your money and you can't stop them.

Different vendors often use different and sometimes proprietary storage formats for storing keys. For example, Entrust uses the proprietary **.epf** format, while Verisign, GlobalSign, and Baltimore, to name a few, use the standard **.p12** format.

1.3 The components of a PKI

A public key infrastructure is created by combining a number of services and technologies:

1) Certification authority (CA)

A CA issues and verifies certificates (see above). The CA takes responsibility for identifying (to a stated extent) the correctness of the identity of the person asking for a certificate to be issued, and ensures that the information contained within the certificate is correct and digitally signs it.

Generating key pairs

The CA may generate a public key and a private key (a key pair) or the person applying for a certificate may have to generate their own key pair and send a signed request containing their public key to the CA for validation. The person applying for a certificate may prefer to generate their own key pair so as to ensure that the private key never leaves their control and as a result is less likely to be available to anyone else.

Issuing Certificates

Unless you generate your own certificate (some applications software will enable you do this) you will generally have to purchase one from a CA. Before a CA issues you with a certificate they will make various checks to prove that you are who you say you are.

The CA could be thought of as the PKI equivalent of a passport agency – the CA issues you a certificate after you provide the credentials they require to confirm your identity, and then the CA signs (stamps) the certificate to prevent modification of the details contained in the certificate.

A CA may also state the quality of the checks that were carried out before the certificate was issued. Different classes of certificate can be purchased that correspond to the level of checks made. There are three or four general classes of certificate: Class 1 certificates can be easily acquired by supplying an email address, Class 2 certificates require additional personal information to be supplied, and Class 3 certificates can only be purchased after checks have been made as to the requestors identity. A 4th class may be used by governments and organizations needing very high levels of checking.

Using certificates

An individual may have any number of certificates issued by any number of CAs. Different Web applications may insist that you use certificates issued only by certain CAs. For example, a bank may insist that you use a certificate issued by them in order to use their services, whereas a public Web site may accept any certificate you offer (just as some allow free choice of ID and password).

The CA can be a unit within your organization, a company (i.e. a bank or a post office), or an independent entity (VeriSign).

Verifying Certificates

The public key certificate is signed by the CA to prevent its modification or falsification. This signature is also used when checking that the public key is still valid. The signature is validated against a list of 'Root CAs' contained within various 'PKI aware' applications (e.g. your browser).

Some CA certificates are called 'Root Certificates' as they form the root of all certificate validation. Certificate validation occurs automatically using the appropriate public certificate contained within the root CA list.

Note: PGP (Pretty Good Privacy) users normally act as their own issuing authority, so you accept their certificate on the basis that they are who they say they are without further verification. This method is called the 'Web of trust' because it is based upon people you trust rather than liability by contract.

2) Revocation

Where a system relies upon publishing certificates so that people are able to communicate with each other, there has to be a system for letting people know when certificates are no longer valid. It can be done in one of two ways. Certificates can be deleted from the Directory or database in which they should be found. As a result, any attempt to find them to check that they still exist will fail and anyone looking for them would know that they have been revoked.

There are two problems with this approach. The first is that a denial of service attack on the Directory or database might create the appearance of a failed certificate. The second is that the Directory was designed to optimize the time to read information, so deleting information is normally avoided, as is updating. Also, deleting the record does not tell the person asking for the information why it is not there, and they may need to know why and when it was removed.

As a result, a system of revocation lists has been developed that exists outside the Directory or database. This is a list of certificates that are no longer valid (for whatever reason), equivalent to a lost or stolen ATM card list. There are currently two different methods for checking for certificate revocation – 'CRL' or 'OCSP'. Revocation lists may be publicly available even when the matching Directory or database is not. This is because certificates may have been distributed for use beyond the private network of the organization involved.

3) Registration Authority (RA)

A CA may use a third-party – a Registration Authority (RA) – to perform the necessary checks on the person or company requesting the certificate to ensure that they are who they say they are. That RA may appear to the certificate requestor as a CA, but they do not actually sign the certificate that is issued.

4) Certificate publishing methods

One of the fundamentals of PKI systems is the need to publish certificates so that users can find them. (You must be able to get hold of the public encryption key for the recipient of encrypted information.) There are two ways of achieving this. One is to publish certificates in the equivalent of an electronic telephone directory. The other is to send your certificate out to those people you think might need it by one means or another. The commonest approaches are listed below.

Directories

Directories are databases that are X.500/LDAP-compliant. This means that they contain certificates in the X.509 format, and that they provide specific search facilities as specified in the LDAP standards published by the IETF.

Directories may be made publicly available or they may be private to a specific organization – i.e. a company may have its own directory where it holds the certificates for its users and only its users can access this directory. A Directory is kept private when it contains information that the owner does not wish to be publicly available. Public directories on the other hand can be read by anyone with access to them.

Databases

A database can be configured to accept X.509 format certificates. This may be done for private systems where the search methods for locating certificates do not follow the LDAP structure. Because it is essentially proprietary, this method is not used for public systems.

Email, floppy discs etc.

Certificates may be sent within an e-mail so that the recipient can add them to their own collection on their server or desktop, depending upon the way their security systems have been configured. They may also be put onto floppy discs, or any other medium.

5) Certificate Management System

This term refers to the management system through which certificates are published, temporarily or permanently suspended, renewed or revoked. Certificate management systems do not normally delete certificates because it may be necessary to prove their status at a point in time, perhaps for legal reasons. A CA (and perhaps an RA) will run certificate management systems to be able to keep track of their responsibilities and liabilities.

6) 'PKI aware' applications

This term usually refers to applications that have had a particular CA software supplier's toolkit added to them so that they are able to use the supplier's CA and certificates to implement PKI functions. The term does not mean that the applications have any 'knowledge' built into them about what the security requirements really are, or which PKI services are relevant to delivering them. These issues are quite separate from having PKI services available.