A Practical Guide to Public Key Infrastructure

Introduction to Public Key Infrastructure

Building Trust on the Internet

Every day companies and individuals use the Internet to complete thousands of online transactions. Employees share files and confidential information via e-mail or computer networks, bank customers update their accounts or pay bills from their home computers, and products of every imaginable shape and function are ordered and paid for via online order forms.

As a business tool, the Internet has the potential to replace telephones and fax machines in daily transactions. In the wrong hands, Internet technology can also be used to intercept and forge messages, steal sensitive information, eavesdrop, and defraud organizations and individuals.

These threats exist because Internet transactions are inherently anonymous and public. In face-to-face transactions, there is a high level of trust between participants. For example, your bank can trust you because it has examined your government-issued identification and has your signature on file. When you apply for a personal loan, you are confident that the transaction is secure and confidential because you shared information only with the bank employee.

As more organizations use the Internet to do business, it is necessary to build trust between people who have never met and cannot meet each other. These people may only be identified by an e-mail address. E-mail addresses, however, are not a trusted form of identification.

A public key infrastructure provides a stronger form of identification. A public key infrastructure is a trust framework that organizations can build into their network systems (Internet, intranet, extranet) and security policies. Public key infrastructures can make Internet transactions as secure as face-to-face transactions.

What is a Public Key Infrastructure?

A public key infrastructure (PKI) makes it possible for you to identify and trust another Internet user. This can be another person, a computer, or some other electronic entity.

In a PKI, digital identification is used to prove the identity of Internet users. This digital ID, called a *digital certificate*, is like a driver's license. It contains the Internet user's name and some credentials. (The content of every digital certificate varies, depending on organizational policies and privacy issues.)

A digital certificate can also be used to verify a digital signature that is attached to e-mail messages or electronic forms. The signature itself is created using public key cryptography. Digital signatures and public key cryptography are explained in more detail in the following sections.

Public Key Cryptography

PKIs are built upon a security solution called *public key cryptography*.

In public key cryptography, a mathematical algorithm and a value are used to transform information into a form that is unreadable. That is, the information is encrypted. The information can be transformed back into a readable form (decrypted) only by using a complementary mathematical algorithm and a second, related value. These values are called *keys*.

Public key cryptography allows PKI participants to make one key publicly available while keeping the other key secret. The key that is publicly available is called the *public key*. The key that is kept secret is called the *private key*. Together, public and private keys are called a *keypair*.

Private keys can be stored on a computer's hard disk or on special cryptographic hardware such as a token, smart card, or hardware security module (HSM). If a private key is stored on a token, it can be used only while the token is inserted into the computer.

Public keys are usually associated with digital certificates. Digital certificates are easy to distribute, either as an attachment to an e-mail message or through the Web. By associating the public key value with a digital certificate, the identity of the person, computer, or other entity identified in the digital certificate can be strongly associated with the public key value.

Digital Signatures: Public Key Cryptography at Work

Digital signatures are one of the primary ways that public key cryptography can be used to make Internet communication safer for organizations and individuals.

To create a digital signature for an e-mail message, for example, a *hash* of the message[1] is encrypted using a private key. This encrypted information is called the *digital signature*. The digital signature is sent, along with the e-mail message and the sender's digital certificate, to another person. The digital signature can be decrypted and verified only by using the public key associated with the sender's digital certificate.

Figure 1: Creating a Digital Signature Using a Private Key

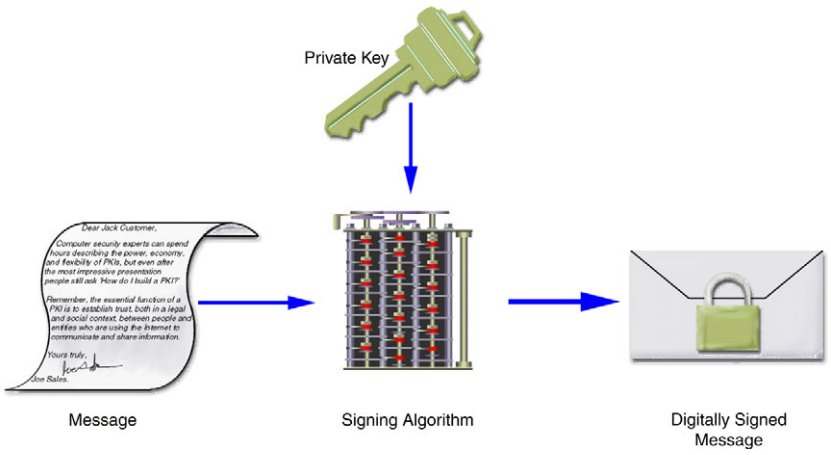

With a digital certificate, a digital signature can be used to identify a person, computer, or other entity. A digital signature can also be used to ensure that a message or file has not been tampered with.

For example, if you add a digital signature to an e-mail message and send it to your boss, he or she can confirm that you made the signature by decrypting it with your public key. If the digital signature cannot be verified (decrypted) with your public key, your boss will know that either it is not your signature or the message has been altered during transmission. This is because you are the only person who can use your private key to create your digital signature. Your public key will verify only signatures created with your private key.

1.The hash is a message fingerprint. Hash functions, fingerprints, and digital signatures are described in more detail in "Keys and Keypairs" on page 17.

There are other ways to use public key cryptography in a PKI, but digital signatures are a popular application of public key cryptography.

Components of a PKI

A PKI usually has the following components:

- Certificate Authority (CA)
- Registration Authority (RA)
- PKI-Enabled Applications

This section also describes two methods for checking the status of certificates issued by a CA. Publishing a certificate's status is the most important task that a CA performs after a certificate is issued.

Certificate Authorities

A certificate authority (CA) is a trusted authority that is responsible for creating, distributing, and revoking digital certificates.

A digital certificate identifies an Internet user by a name and contains a public key value. The process of binding a public key value to a person, computer, or other entity is called *certification*.

A CA is like a licensing authority. Digital certificates are issued only to users who can prove their identity and credentials to the CA. The CA may examine traditional forms of identification, such as a driver license or company records, before issuing a digital certificate. This process is called *vetting*.

CAs also respond to queries about the validity of certificates that they have issued. CAs always include a validity period in certificates. A certificate's validity period specifies how long the CA expects the certificate's contents to remain valid. A certificate may become unexpectedly invalid if information about the certificate subject changes—for example, if the person is no longer employed at the company or their name changes. A certificate is valid if it has not expired and the information in the certificate is true. The process of confirming that the information contained in a digital certificate is still valid is called *validation*.

Certificate Status Checking

CAs revoke certificates when information in the certificate becomes unexpectedly invalid or when it is necessary to revoke the PKI privileges of a user. The CA cannot delete the certificate or retrieve it from the user because the certificate is a public

document that is used by thousands of PKI participants (in addition to the person or entity identified in the certificate). Instead, the digital certificate is marked as "revoked" in the CA's database. This process is called *revocation*. A similar process occurs when CAs *suspend* certificates. CAs suspend certificates when they want to temporarily invalidate them. The certificates are then marked as "suspended" in the CA's database.

PKI users can determine if a digital certificate has been revoked or suspended by looking up the certificate's validity in the CA's database. This process is called *real-time online certificate status checking*. This process is extremely valuable to companies and organizations because it ensures that certificate validity information is always current and accurate.

Figure 2: Real-Time Online Certificate Status Checking

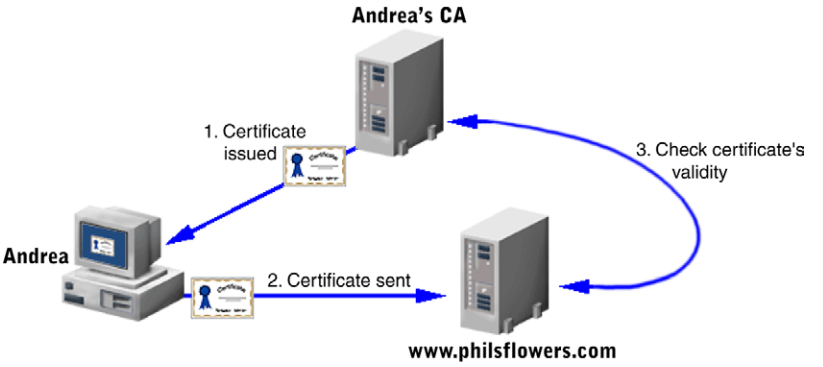

An older, but less reliable, method of certificate status checking requires PKI users to download a certificate revocation list (CRL). A CRL is a list of certificates that have been revoked by the CA. CRLs are generated periodically by a CA. One problem with CRLs is that they can be difficult to download and use. If a CA issues a CRL every day, all PKI participants (including servers and other electronic entities) must download the CRL on a daily basis.

A more significant problem is that CRL information is always stale. If a CA issues a CRL daily and it revokes a certificate immediately after issuing a CRL, then the revocation will not be known to PKI users until the next day, when the next CRL is issued. Stale CRLs are a serious risk for organizations that use the Internet to do business. If you are completing high-value transactions, not knowing the correct certificate status for even a small window of time can lead to fraudulent transactions.

An alternative to downloading a CRL is using the Online Certificate Status Protocol (OCSP) to make a standard call to an OCSP responder. This is an online method of obtaining certificate status. Leading vendors have implemented OCSP so that it gives real-time status.

Real-time online certificate status checking provides up-to-the-second certificate status information to PKI users. Instead of downloading a CRL to check the validity of a certificate, PKI users simply look up the certificate's status in the CA's database. This method of certificate status checking has several advantages over CRLs. Real-time online status checking is fast, easy to use, provides accurate and fresh status information, and reduces risk.

Registration Authorities

Registration authorities (RAs) are used to enroll new users into a PKI.

RAs are responsible primarily for vetting certificate requests. Approved certificate requests are sent to a CA, which then creates the requested digital certificates. Digital certificates may be distributed to users through either the CA or the RA.

There are three distinct advantages to using RAs:

- With RAs, organizations can set up local or standalone enrollment centers at distributed geographic locations. In an international company, for example, employees may be enrolled into a PKI via RA centers in Singapore, Germany, and Canada. Digital certificates are issued to those employees from the company's CA, which may be located in the United States.
- Organizations can separate the PKI operations performed by the CA and the RA. This is necessary if the organization wants to separate the certificate request process from the certificate issuing process, or if it wants to set up an audit trail.
- RAs enable better security because users interact only with the RA.

PKI-Enabled Applications

One of the greatest advantages to using a PKI is that it can be supported through a variety of off-the-shelf software programs such as:

- Web browsers
- E-mail clients
- VPN software and hardware

The two most widely-used Web browsers, Microsoft Internet Explorer and Netscape Navigator, are already PKI-enabled. They provide users with the ability to generate a keypair, download a digital certificate, and strongly authenticate to a Web server.

Popular e-mail programs such as Microsoft Outlook and Netscape Messenger are also PKI-enabled. Users instruct their e-mail program to digitally sign a message simply by clicking a button.

Companies frequently use an extranet to distribute information to customers and business partners. To extend the PKI beyond the firewall, many companies PKI-enable their extranet or create a virtual private network (VPN). Extranets and VPNs can use digital certificates to authenticate users and provide access control.

Almost any program that is used to connect to and share information over the Internet can be PKI-enabled. The flexibility and robustness of PKI technology allows organizations to create secure PKI solutions that meet their business needs.

PKI Security Policies

Organizations usually implement a PKI to resolve network security issues. Network security is only half of the equation, however. PKI applications operate within a framework of legal and social responsibilities which must be addressed through policy.

Before setting up a network security policy, organizations should establish policies governing user access privileges, administrative duties, system maintenance, and how the organization will prevent (and react to) security breaches. These policies are the common-sense foundation of any security system because they provide clear guidelines for operating the PKI.

Once these issues have been addressed, organizations can establish PKI security policies. These frequently include:

- Certification Practice Statement (CPS)
- Certificate Policy (CP)
- Liability and other legal considerations

Addressing policy issues is usually the first step in creating and implementing a PKI. Policy drives the design of the PKI and defines the constraints under which the PKI will operate. These three PKI policy concerns are described in more detail in the following sections.

Certification Practice Statement

A certification practice statement (CPS) is a legal document, created and published by a CA. It explains the CA's certificate issuance and revocation policies. Individuals and organizations use the CA's CPS to determine the level of trust they can place in a CA.

Certificate Policy

A certificate policy (CP) explains the conditions and limitations of use for a digital certificate. In other words, a CP defines the level of trust that a user can place in someone else's digital certificate. A CP can be embedded in or referenced in a digital certificate.

Liability and Other Legal Considerations

Organizations should determine who assumes responsibility for losses if a PKI entity engages in or falls victim to fraud. For example, does the CA bear full responsibility for all losses? Or do members of the PKI share responsibility? After determining an organization's legal rights and obligations, policies can be put in place to ensure that these responsibilities will be met.

Choosing a PKI Solution

A wide range of PKI solutions are already available in the market. Comparisons between product offerings can be confusing and frustrating because every PKI solution has different strengths. This section highlights some of the challenges facing anyone who is shopping for a PKI solution.

Interoperability

Interoperability is a major concern for anyone setting up an organization-wide system. Modern computer networks are composed of hardware and software purchased from several (often hundreds) of different vendors. An organization's network system works only because these third-party hardware and software products support the same protocols and standards.

For PKI vendors, interoperability means two things: Does the PKI interoperate with other applications, such as e-mail clients, Web browsers, and VPN software? And can two CAs from different PKIs interoperate with each other (that is, can they establish trust relationships and share information about certificate status)?

Choosing a PKI solution is particularly difficult because many standards for PKI technology are still in draft form. In the absence of clearly defined standards, some PKI vendors have developed proprietary PKI solutions. The best approach is to select a PKI vendor that embraces PKI standards and best practices.

Every organization has unique PKI needs. For this reason, many PKI vendors offer PKI development toolkits and application programming interfaces (APIs) that can be used by programmers to PKI-enable an organization's software programs.

Even if a PKI vendor offers a software development kit (SDK), it is still preferable to purchase a PKI solution that offers proven interoperability. Implementing a PKI is a significant investment and a PKI solution should meet as many of an organization's needs as possible.

Scalability

The issue of scalability is strongly related to interoperability. When we say that a PKI should be scalable, we mean:

- The PKI's services should extend not only throughout an organization, but also beyond it—to provide the appropriate level of security for all of the organization's internal users, as well as for any external entities that the organization deals with.

- The PKI should operate efficiently and effectively with all of the organization's users. A PKI that cannot provide reliable, high-volume service will be a liability for a growing organization.

Usability

A PKI must be easy to use. No one will use a PKI if the enrollment process is complicated or it is difficult to use the PKI in tandem with everyday Internet or work activities. The PKI should integrate seamlessly into an organization's existing network system and software programs. It should require little or no special training to use. For example, a user should be able to create or validate a digital signature by simply clicking a button.

A PKI must also be easy to manage. The administrative interface for the CA or RA should be an intuitive graphical user interface (GUI) that can be used to process high volumes of certificates and certificate requests. The interface should also be flexible and customizable. For example, it should be possible to automate routine tasks and configure the GUI to meet unique administrative needs.

Implementing a PKI should be as easy as possible. The everyday use of the PKI should be intuitive and should not require special training. A competent system administrator should be able to easily install the PKI software and quickly configure an organization's network systems to use a PKI.

Usability should be a concern for anyone choosing a PKI solution. If the PKI is too cumbersome to use, it will find little acceptance within the organization. If the PKI is used inconsistently or incorrectly, the organization's messages, files, and electronic paperwork will not be protected by the PKI' security and trust framework. To be effective, a PKI must integrate seamlessly into an organization's day-to-day operations.

Summary

All around the world, organizations are using the Internet and other network systems to do business. Surprisingly few of these organizations adequately protect themselves from eavesdropping, data theft, and other forms of fraud.

A PKI is a trust framework that can be extended to include every person, computer, and electronic entity in your organization. Among the advantages a PKI can offer to an organization are:

Authentication — Everyone in the PKI is identified by a digital certificate. Anyone who receives a digital certificate can verify the identity of the certificate holder and the validity of the certificate.

Non-repudiation — The author of a message cannot later deny having created the message (that is, repudiation cannot occur). Digital signatures can be used to establish the non-repudiation of transactions.

Data Integrity — Files and information have not been altered or tampered with during transmission. Data integrity is assured using public key cryptography.

Confidential Communication — Only the intended recipient is able to read the file or message.

Access Control — Access to sensitive information is controlled through the use of authenticated identities.

PKIs are flexible and can be adapted to meet the business needs of any organization, small or large. By choosing a PKI solution that interoperates with a wide range of third-party applications and is easy to use and manage, organizations can quickly and seamlessly integrate the PKI into their day-to-day business operations.