



Entrust

Trusted Public-Key Infrastructures

Date: August 2000
Version: 1.2



Entrust is a registered trademark of Entrust, Inc. in the United States and certain other countries. Entrust is a registered trademark of Entrust Limited in Canada. All other Entrust product names and service names are trademarks or registered trademarks of Entrust, Inc or Entrust Limited. All other company and product names are trademarks or registered trademarks of their respective owners.

© Copyright 2000-2003 Entrust. All rights reserved.

Introduction

As organizations look to gain competitive advantage by improving their products and services, technology related to digital signatures and security of information is an attractive option. In the past few years, public-key technology has become the preferred means for providing these capabilities. Public-key technology provides a variety of critical enabling capabilities for bringing trust to e-business transactions. Through encryption, public-key technology provides **confidentiality**. Through digital signatures, the technology provides the following features:

- **Strong authentication.** Strong authentication means users can securely identify themselves to other users and servers on a network without sending secret information (for example, passwords) over the network.
- **data integrity.** Data integrity means that the verifier of a digital signature can easily determine whether or not digitally signed data has been altered since it was signed.
- **support for non-repudiation.** Support for non-repudiation means that the user who signed data cannot successfully deny signing that data.

Each capability is required to effectively move business processes from the paper-based world to the electronic world, and to improve existing electronic processes. The rest of this paper elaborates on the elements that allow businesses to take advantage of public-key technology. In particular, this paper concentrates on the following items:

- the concept of a public-key infrastructure
- the requirements for implementing an effective, comprehensive public-key infrastructure.

What is a public-key infrastructure?

The comprehensive system required to provide public-key encryption and digital signature services is known as a *public-key infrastructure* (PKI).

The purpose of a public-key infrastructure is to manage keys and certificates. By managing keys and certificates through a PKI, an organization establishes and maintains a trustworthy networking environment. A PKI enables the use of encryption and digital signature services across a wide variety of applications.

Note: This paper assumes the reader has a basic understanding of public-key cryptography. To get a brief overview of cryptography, refer to the White Paper titled *An Introduction to Cryptography*, available on the Entrust Web site at <http://www.entrust.com>.

What is an effective public-key infrastructure?

There are a number of requirements that businesses have with respect to implementing effective public-key infrastructures. First and foremost, if users cannot take advantage of encryption and digital signatures in applications, a PKI is not valuable. Consequently, the most important constraint on a PKI is *transparency*. The term *transparency* means that users do not have to understand how the PKI manages keys and certificates to take advantage of encryption and digital signature services. An effective PKI is transparent.

In addition to user transparency, a business must implement the following items in a PKI to provide the required key and certificate management services:

-
- public key certificates
 - a certificate repository
 - certificate revocation
 - key backup and recovery
 - support for non-repudiation of digital signatures
 - automatic update of key pairs and certificates
 - management of key histories
 - support for cross-certification
 - client-side software interacting with all of the above in a secure, consistent, and trustworthy manner.

Note: In this paper, the term *client-side* refers to application clients and application servers. PKI requirements are the same for both application clients and servers, and both are “clients” of the infrastructure services described in this paper.

The remaining sections of this paper define each of the requirements listed above. All of these requirements must be met for an organization implementing a PKI to establish and maintain a trustworthy environment. All of these requirements must also be met to have an automatic, transparent, and usable PKI.

Certificates and Certification Authorities

For public-key cryptography to be valuable, users must be assured that the other parties with whom they communicate are “safe”—that is, their identities and keys are valid and trustworthy. To provide this assurance, all users of a PKI must have a registered identity. These identities are stored in a digital format known as a public key *certificate*. *Certification Authorities* (CAs) represent the people, processes, and tools to create digital certificates that securely bind the names of users to their public keys.

In creating certificates, CAs act as agents of trust in a PKI. As long as users trust a CA and its business policies for issuing and managing certificates, they can trust certificates issued by the CA. This is known as *third-party trust*. For more information on third-party trust, refer to the Entrust White Paper titled *The Concept of Trust in Network Security*. This White Paper is available on the Entrust Web site at <http://www.entrust.com>.

CAs create certificates for users by digitally signing a set of data that includes the following information (and additional items):

- the user’s name in the format of a distinguished name (DN). The DN specifies the user’s name and any additional attributes required to uniquely identify the user (for example, the DN could contain the user’s employee number).
- a public key of the user. The public key is required so that others can encrypt for the user or verify the user’s digital signature.
- the validity period (or lifetime) of the certificate (a start date and an end date).
- the specific operations for which the public key is to be used (whether for encrypting data, verifying digital signatures, or both).

The CA’s signature on a certificate ensures that any tampering with the contents of the certificate can be easily detected. (The CA’s signature on a certificate is like a

tamper-detection seal on a bottle of pills—any tampering with the contents of a certificate is easily detected) As long as the CA’s signature on a certificate can be verified, the certificate has integrity. Since the integrity of a certificate can be determined by verifying the CA’s signature, certificates are inherently secure and can be distributed in a completely public manner (for example, through publicly-accessible directory systems).

Users retrieving a public key from a certificate can be assured that the public key is valid. That is, users can trust that the certificate and its associated public key belong to the entity specified by the distinguished name. Users also trust that the public key is still within its defined validity period. In addition, users are assured that the public key may be used safely in the manner for which it was certified by the CA.

Key backup and recovery

A business must be able to retrieve encrypted data when users lose their decryption keys. This means that the enterprise to which the user belongs requires a system for backing up and recovering the decryption keys. There are two reasons why key backup and recovery are so important to businesses.

The first reason is that users forget passwords. It is potentially catastrophic for a business to lose data when users forget the passwords required to access their decryption keys. Valuable information would be lost forever if there was no ability to securely recover those keys. Furthermore, unless users know they can always recover their encrypted data (even if they forget their passwords), some users will not encrypt their most valuable and sensitive information for fear of losing it—even though that information needs to be protected the most.

The second reason is that users may lose, break, or corrupt the devices in which their decryption keys are stored. For instance, if a user’s decryption keys are stored on a magnetic card, the magnetic field on the card can become corrupted. Again, permanent loss of those decryption keys can be disastrous. Users are prevented from recovering encrypted data unless their decryption keys are backed up.

Which keys require backup?

Earlier, this paper introduced the notion of different functions for key pairs. One key pair is used for encrypting and decrypting data. This is called the “encryption key pair”. Another key pair is used for digitally signing data and verifying signatures. This is called the “signing key pair”. Note that there is no discussion above regarding backup and recovery of signing key pairs. The only keys requiring backup are users’ decryption keys. As long as a trusted agent (for example, the CA) securely backs up users’ decryption keys, security is not compromised and the user’s data can always be recovered. However, signing keys have different requirements from decryption keys. In fact, as the next section describes, backing up signing keys destroys a basic requirement of a PKI.

Support for non-repudiation

Repudiation occurs when an individual denies involvement in a transaction. (For instance, when someone claims a credit card is stolen, this means that he or she is repudiating liability for transactions that occur with that card anytime after reporting the theft). Non-repudiation means that an individual cannot successfully deny involvement in a transaction. In the paper-world, individuals’ signatures legally bind them to their transactions (for example, credit card charges, business contracts, ...). The signature prevents repudiation of those transactions. In the electronic world, the replacement for the pen-based signature is a digital signature. All types of e-business

require digital signatures because electronic commerce makes traditional pen-based signatures obsolete.

The signing private key

The most basic requirement for non-repudiation is that the key used to create digital signatures—the signing key—be generated and securely stored in a manner under the sole control of the user at all times. It is not acceptable to back up the signing key.

Unlike encryption key pairs, there is no technical or business requirement to backup or restore previous signing key pairs when users forget their passwords or lose, break, or corrupt their signing keys. In such cases, it is acceptable for users to generate new signing key pairs and continue using them from that time forward.

The need for two key pairs

It is difficult to simultaneously support key backup and recovery *and* non-repudiation. To support key backup and recovery (as discussed in a previous section), the decryption keys must be backed up securely. To support non-repudiation, the keys used for digital signature cannot be backed up and must be under the sole control of the user at all times.

To meet these requirements, a PKI must support two key pairs for each user. At any point in time, a user must have one current key pair for encryption and decryption, and a second key pair for digital signature and signature verification.

Over time, users will have numerous key pairs that must be managed appropriately, as discussed in the following section.

Key update and management of key histories

Cryptographic key pairs should not be used forever. They must be updated over time. As a result, every organization needs to consider two important issues:

- updating users' key pairs, and
- maintaining, where appropriate, the history of previous key pairs.

Updating users' key pairs

The process of updating key pairs should be transparent to users. This transparency means users do not have to understand that key update needs to take place and they will never experience a “denial of service” because their keys are no longer valid. To ensure transparency and prevent denial of service, users' key pairs must be automatically updated before they expire.

Maintaining histories of key pairs

When encryption key pairs are updated, the history of previous decryption keys must be maintained. This “key history” ensures that users can access any of their prior decryption keys to decrypt data. (When data is encrypted with a user's encryption key, only the corresponding decryption key—the paired key—can be used for decrypting). To ensure transparency, the client-side software must automatically manage users' histories of decryption keys.

The key history must also be securely managed by the key backup and recovery system. This ensures that encrypted data can be recovered securely, regardless of

what encryption public key was used to originally encrypt the data (and, by extension, regardless of when the data was encrypted).

When a signing key pair is updated, the previous signing key should be securely destroyed. This destruction prevents any other person from gaining access to the signing key and is acceptable because there is no need to retain previous signing keys.

Certificate repositories and certificate distribution

As mentioned earlier in this paper, the CA acts as a trusted third-party issuing certificates to users. Businesses also must distribute those certificates so they can be used by applications. Certificate repositories store certificates so that applications can retrieve them on behalf of users. The term *repository* refers to a network service that allows for distribution of certificates.

Over the past few years, the consensus in the information technology industry is that the best technology for certificate repositories is provided by directory systems that are LDAP (Lightweight Directory Access Protocol)-compliant. LDAP defines the standard protocol to access directory systems.

Several factors drive this consensus position:

- storing certificates in directories and having applications retrieve certificates on behalf of users provides the transparency required for use in most businesses
- many directory technologies supporting LDAP can be scaled to:
 - support a very large number of entries
 - respond efficiently to search requests due to their information storage and retrieval methods, and
 - be distributed throughout the network to meet the requirements of even the most highly-distributed organizations

In addition, the directories that support certificate distribution can store other organizational information. As discussed in the next section, the PKI can also use the directory to distribute certificate revocation information.

Certificate revocation

In addition to verifying the CA's signature on a certificate (as discussed in the section titled *Certificates and Certification Authorities*), the application software must also be sure that the certificate is still trustworthy at the time of use. The CA must revoke certificates that are no longer trustworthy.

There are numerous reasons why a certificate may need to be revoked prior to the end of its validity period. For instance, the private key (that is, either the signing key or the decryption key) corresponding to the public key in the certificate may be compromised. Alternatively, an organization's security policy may dictate that the certificates of employees leaving the organization must be revoked. In these situations, users in the system must be informed that continued use of the certificate is no longer considered secure.

The revocation status of a certificate must be checked prior to each use. As a result, a PKI must incorporate a scalable certificate revocation system. The CA must be able to securely publish information regarding the status of each certificate in the system. Application software, on behalf of users, must then verify the revocation information prior to each use of a certificate. The combination of publishing and consistently using certificate revocation information constitutes a complete revocation system.

A way to distribute certificate revocation information is for the CA to create secure certificate revocation lists (CRLs) and publish these CRLs to a directory system. CRLs specify the unique serial numbers of all revoked certificates. Another way to distribute revocation information is through Online Certificate Status Protocol (OCSP), where a OCSP server responds to client requests for revocation status. Regardless of the method use to check revocation status, prior to using a certificate, the client-side application must *automatically* check the revocation status to determine if the certificate is still trustworthy. Client-side applications must automatically check for revoked certificates consistently and transparently on behalf of users.

Cross-certification

Cross-certification, or PKI networking, extends third-party trust relationships between Certification Authority domains. For example, two trading partners, each with their own CA, may want to validate certificates issued by the other partner's CA via peer-to-peer cross-certification. Alternatively, a large, distributed organization may require multiple CAs in various geographic regions using hierarchical cross-certification to join regional subordinate CAs to a root CA. Cross-certification allows different CA domains to establish and maintain trustworthy electronic relationships.

The term *cross-certification* refers to two operations. The first operation, which is generally executed infrequently, is the establishment of a trust relationship between two CAs. In the case of peer-to-peer cross-certification, two CAs securely exchange their verification keys. These are the keys used to verify the CAs' signatures on certificates. To complete the operation, each CA signs the other CA's verification key in a certificate referred to as a "cross-certificate". In the case of hierarchical cross-certification, a superior CA certifies a subordinate CA by signing the subordinate CA's verification key.

The second operation is done by the client-side software. The operation, which is executed frequently, involves verifying the trustworthiness of a user certificate signed by a cross-certified CA. The operation is often referred to as "walking a chain of trust". The "chain" refers to a list of cross-certificate validations that are "walked" (or traced) from the CA key of the verifying user to the CA key required to validate the other user's certificate.

When walking a chain of cross-certificates, each cross-certificate should be checked to ensure that it is still trusted. User certificates must be able to be revoked; so must cross-certificates. This requirement is frequently overlooked in discussions regarding cross-certification.

For more information on third-party trust and cross-certification, refer to the Entrust white paper titled *The Concept of Trust in Network Security and Cross-Certification and PKI Policy Networking*. These white papers are available on the Entrust Web site at <http://www.entrust.com>.

Client-side software

When discussing requirements for PKI, businesses often neglect the requirement for client-side software. (For instance, many people only focus on the CA component when discussing PKIs). Ultimately, however, the value of a PKI is tied to the ability of users to use encryption and digital signatures. For this reason, the PKI must include client-side software that operates consistently and transparently across applications on the desktop (for example, e-mail, Web browsing, e-forms, file/folder encryption, ...). A consistent, easy-to-use PKI implementation within client-side software lowers PKI operating costs.

In addition, client-side software must be technologically enabled to support all of the elements of a PKI discussed earlier in this paper. The following list summarizes the requirements client-side software must meet to ensure that users in a business receive a usable, transparent (and thus, acceptable) PKI.

- *public key certificates*

To provide third-party trust, all PKI-enabled applications must use certificates in a consistent, trustworthy manner. The client-side software must validate the CA's signature on certificates and ensure that the certificates are within their validity periods.

- *key backup and recovery*

To ensure users are protected against loss of data, the PKI must support a system for backup and recovery of decryption keys. With respect to administrative costs, it is unacceptable for each application to provide its own key backup and recovery. Instead, all PKI-enabled client applications should interact with a single key backup and recovery system. The interactions between the client-side software and the key backup and recovery system must be secure, and the interaction method must be consistent across all PKI-enabled applications.

- *support for non-repudiation*

To provide basic support for non-repudiation, the client-side software must generate the key pairs used for digital signature. In addition, the client-side software must ensure that the signing keys are never backed up and remain under the users' control at all times. This type of support must be consistent across all PKI-enabled applications.

- *automatic update of key pairs*

To ensure transparency, client-side applications must automatically initiate updating of users' key pairs. This activity must be done in accordance with the security policies of the organization. It is unacceptable for users to have to know that their key pairs require updating. To meet this requirement across all PKI-enabled applications, the client-side software must update key pairs transparently and consistently.

- *management of key histories*

To ensure that users can easily access all data encrypted for them (regardless of when it was encrypted), PKI-enabled applications must have access to users' key histories. The client-side software must be able to securely recover users' key histories.

- *a scalable certificate repository*

To minimize the costs of distributing certificates, all PKI-enabled applications must use a common, scalable certificate repository. Transparent support for certificate retrieval from a common repository improves usability. It is costly and cumbersome for applications to require different certificate repositories.

- *certificate revocation*

PKI-enabled applications must interact with a common, scalable certificate revocation system. Because certificate revocation is central to trust management, client-side software must interact with the certificate revocation system in a consistent manner across all applications.

- *support for cross-certification*

To ensure consistent application of trust management policies, all PKI-enabled applications must use a common cross-certification model. As mentioned in the section titled "*Cross-certification*," cross-certification allows the PKI to determine the trustworthiness of a certificate issued by a foreign Certification

Authority. For example, the client-side software must check to ensure that none of the cross-certificates in a chain of trust is revoked.

Each of the issues discussed above relates to interactions between the client-side software and the infrastructure elements of a PKI. There are additional requirements of the client-side software that are independent of the infrastructure elements. For example, the client-side software should allow users to encrypt and decrypt information even when they are disconnected from the infrastructure elements of the PKI. To maximize usability and minimize cost, the client-side software should support multiple types of key storage devices (for example, smart cards, PC cards, secure files, etc.). Users should be able to use a single key storage device across all PKI-enabled applications.

Summary

The goal of a PKI is to establish and maintain a trustworthy networking environment. This goal is achieved by providing key and certificate management services that enable encryption and digital signature capabilities across applications. A fundamental constraint of a PKI is user transparency.

A comprehensive PKI must implement the following items:

- public key certificates
- a certificate repository
- certificate revocation
- key backup and recovery
- support for non-repudiation of digital signatures
- automatic update of key pairs and certificates
- management of key histories
- support for cross-certification
- client-side software interacting with all of the above in a secure, consistent, and trustworthy manner.

Only a comprehensive PKI can achieve the goal of establishing and maintaining a trustworthy networking environment, while at the same time providing an automatic and transparent system that is usable.

Entrust is a registered trademark of Entrust, Inc. in the United States and certain other countries. Entrust is a registered trademark of Entrust Limited in Canada. All other Entrust product names and service names are trademarks or registered trademarks of Entrust, Inc or Entrust Limited. All other company and product names are trademarks or registered trademarks of their respective owners.

© Copyright 2000-2003 Entrust. All rights reserved.