OFFICE OF THE MANAGER
NATIONAL COMMUNICATIONS SYSTEM

TECHNICAL NOTES

December 2004 *TECHNOLOGY AND PROGRAMS DIVISION* Volume 11, Number 3

Public Key Infrastructure

by Dale Barr

In the mid-1970s, Whitfield Diffie and Martin Hellman documented an approach for encryption in which ciphers were asymmetric, i.e., they used a separate key for encryption and decryption, based on vector-matrix multiplication. The two researchers postulated that these keys were so different that one could be publicized without danger of deriving or computing the other; an unintended recipient would need to perform the more difficult task of matrix inversion to recover the plaintext. Their paper, "New Directions in Cryptography," contained evidence that such ciphers could be constructed.

Public Key Infrastructure Architecture Components

The basic Public Key Infrastructure (PKI) architecture model has remained largely unchanged since it was originally published in the Internet Certificate and Certificate Revocation List (CRL) Profile RFC2459. The latest model is reflected in the most recent version of the Internet Certificate and CRL Profile RFC3280. Table 1 identifies the name and purpose of each component defined in RFC 3280.

Services of Public Key Cryptography

The discovery of public key cryptography has made a number of services available, some of which were either unknown or unachievable with symmetric ciphers.[1] The principal services are as follows, shown in Table 1:

[1] Ciphers in which the exact same key is used to encipher a message, and then decipher the resulting ciphertext.

COMPONENT	PRIMARY ROLE
End Entity	Although sometimes considered only as an end-user, the term "End Entity" has a more generic definition. It can be an end-user; a device, such as a router or server; a process; or any item that can be identified in the subject name of a public key certificate. End Entities can also be consumers of PKI-related services and, in some cases, providers of PKI-related services. For example, a Registration Authority (RA) is considered to be an End Entity from the point of view of the Certification Authority (CA).
Certification Authority	Public keys are distributed in the form of public key certificates. The CA is the foundation of the PKI because only CAs can issue public key certificates. The issuing CAs digitally sign public key certificates, which effectively binds the subject name to the public key. CAs are also responsible for issuing Certificate Revocation Lists (CRLs) unless delegated to a separate CRL Issuer. They may also be involved in a number of administrative tasks such as end-user registration, although these tasks are often delegated to the RA. CAs are often thought of as the "source of trust" in a PKI.
Registration Authority	An RA is an optional component that can be used to "offload" many of the administrative functions that a CA ordinarily assumes. The RA is normally associated with the End Entity registration process. This includes the verification of the identity of the End Entity attempting to register with the PKI. An RA can never be the issuer of a public key certificate.
Repository	Repositories are often associated with a directory. In the context of a PKI, however, a repository is a generic term used to denote any method for storing and retrieving PKI-related information, such as public key certificates and CRLs. A repository can be an X.500[2]-based directory with client access via the Lightweight Directory Access Protocol (LDAP). It also can be something simple, such as a means for retrieval of a flat file on a remote server via the File Transfer Protocol (FTP) or the Hyper Text Transfer Protocol (HTTP).
CRL Issuer	A CRL Issuer provides a CRL. Typically, the CA that issues a given set of certificates is also responsible for issuing revocation information associated with those certificates. However, it is possible for a CA to delegate that function to another entity. CRLs that are issued by another entity are referred to as indirect CRLs.

Table 1. PKI Components

[2]ISO and ITU standards define how global directories should be structured. X.500 directories are hierarchical with different levels for each category of information, such as country, state, and city.

Security between Strangers

A driving motivation behind public key cryptography was the need to enable secure communication between strangers in a public or open environment using an encryption key that can be transmitted over non-secure communications. For example, this key could be stored in a public repository; therefore, Party A can look up Party B's public key, use it to encrypt a message to Party B, then transmit it to Party B. Figure 1 summarizes the PKI process (i.e., asymmetric cryptography).

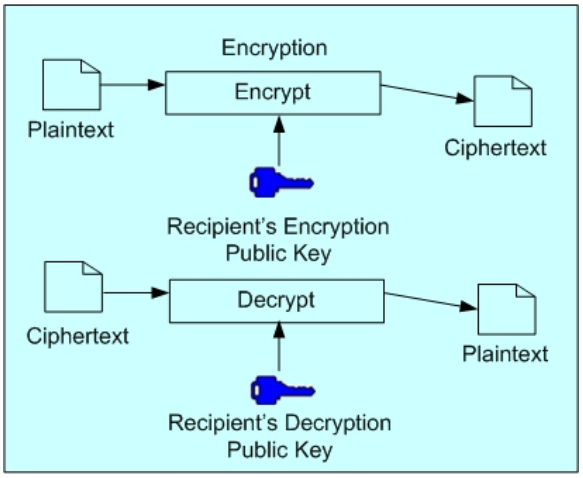

Figure 1. Asymmetric Cipher Model

The caveat in this process is that Party A must be confident that the public key retrieved from electronic repositories really does belong to Party B. The public repository must be trusted to return correct information. Trust may result from various circumstances: Party A may own the repository, have control over it, or the repository is located and controlled on a network trusted by Party A.

In general, public repositories should not be trusted without a third-party process to independently verify the data. A common mechanism to achieve this is the public key certificate.[3]

Digital Signature

A digital signature is a cryptographic function employing a hash function[4] and the signer's private key. The hash function creates a digital representation, or fingerprint, unique to the document. The fingerprint combined with the sender's private key assure the receiver that the digitally signed document has not been altered. Figure 2 illustrates digital signature operations. In addition, all digital signature mechanisms specify particular padding conventions to be applied to the data as part of the signature process. These padding bits are examined and removed as part of the verification process.

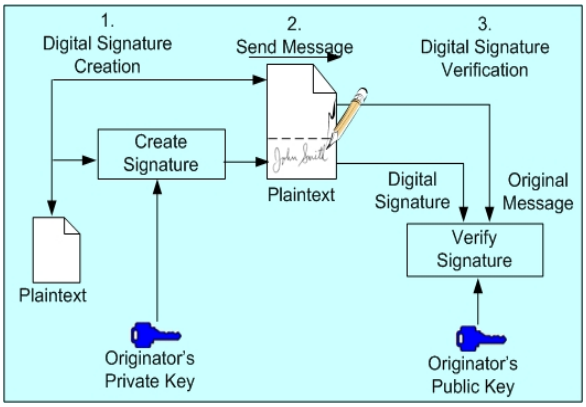

Figure 2. Generic Digital Signature Process

[3] A public key certificate is a digitally signed document that serves to validate the sender's authorization and name. The document consists of a specially formatted block of data that contains the name of the certificate holder (either a user or a system name) and the holder's public key, as well as the digital signature of a CA for authentication.

[4] This is a classic computer operation that forms a fixed-size result from an arbitrary amount of data. Ideally, even the smallest change to the input data will change about half of the bits in the result.

Data Integrity

A digital signature provides both data origin authentication and data integrity (i.e., evidence that the data has not been altered) because no two inputs hash to the same output. Any alteration to the data will lead to a different hash output value and will cause a failure in the signature verification process. If the signature verification is successful, the recipient can be confident that data integrity has been preserved.

Key Establishment

Public key cryptography can also be used to perform key establishment (sometimes called key exchange) between two entities, i.e., a process of two entities creating and sharing a secret symmetric key known only to them. This shared key establishment occurs in two ways:

- **Key Transfer.** Party A generates a symmetric key and sends it to Party B using the public key process.
- **Key Agreement.** Both entities jointly contribute to the generation of a symmetric key.

The resulting common key can be used by the two parties to protect data transmitted between them because only they know the key used for encryption and decryption.

Benefits of PKI Deployment

PKI is a security infrastructure that can be used across multiple applications and multiple environments. It can enable confidentiality, data integrity, authentication, and non-repudiation services in numerous contexts, including one or more of the following:

- Secure e-mail
- Secure Electronic Data Interchange (EDI)
- Secure electronic forms
- Secure desktop (for example, encryption of sensitive information on a laptop or PC)
- Secure intranets
- Secure extranets
- End-user access control (whether the user is a person or a piece of equipment)
- Secure remote access (e.g., in support of mobile users or work-at-home users)
- Secure Web applications

It is important to recognize that security threats originate from both external and internal sources. PKI is especially useful in distributed networks such as those which exist in multi-national corporations.

Barriers to PKI Deployment

Cost

No single formula can be applied to all organizations for determining the cost of deploying a PKI. Some resources within an organization can be leveraged to offset some of the costs. Many factors must be evaluated to help determine the Total Cost of Ownership (TCO) within a given organization, including the following:

- **Number of Hardware Components.** What is required to meet the demands of the target community? The number of components may depend on a variety of factors, including the scale of the community, geographic elements, and the amount of autonomy afforded to the departments or communities of interest.

- **Cost of Software and Support Tools.** Both initial software procurement and ongoing software maintenance costs should be considered.
- **Use of the Existing Corporate IT Infrastructure.** How much can be exploited to support the target community? For example, is a separate directory or repository product required, or can an existing corporate directory service be utilized?
- **Costs of Planning, Deployment, Operation, and Maintenance of the infrastructure.** These should be calculated for the life of the program.
- **Costs of Defining the Policies and Procedures.** Are new policies or procedures required to support external users and/or external organizations?
- **Additional Facilities.** Is the current capability to house the infrastructure components adequate? If not, what is required, and how much will it cost?
- **Availability of Required Components.** Are the PKI components available? Is full component redundancy necessary?
- **Training Costs.** These apply to administrators and operational personnel, as well as end users.
- **Level of Administrative Support.** For example, help desk support, End Entity registration procedures, and data maintenance.
- **Multi-Vendor PKI Deployments.** Will the deployed PKI interoperate with other PKIs from different vendors? Adopting standards-based technology has shown to be an essential cost saving factor.
- **Law and/or Policy-Related Doctrine.** Liability protection is essential in many cases, especially when interoperability is required with external users or other PKI domains.

The key to success is to plan ahead. Understanding as many of the issues as possible will help lead to the development of a solid business strategy and minimize the associated costs.

Lack of Maturity

In recent years, PKI has changed dramatically. While the technology has been around for a while, trade journal articles and conference presentations have suggested that PKI is still an "emerging technology," and they have called for exercising caution when making a deployment decision. They argue that the technology is still fairly new, and the standards and testing facilities necessary to guarantee multi-vendor interoperability have only begun to mature. At the same time, these articles have suggested that PKI is a "must have" technology.

Complexity/Uncertainty

Using a different perspective, some writers more recently claim that PKI has been hyped and is too complicated and expensive to be viable — a concern of government and corporate decision makers.

Because of the degree of uncertainty still associated with PKI technology, the trend for the past few years, which is expected to continue, has been to launch small-scale PKI pilots. These pilots typically focus on a single application (e.g., secure e-mail), and they limit the size of the end-user community typically to no more than a few hundred end users. The purpose of these pilots is to:

- Educate administrators and operations personnel through controlled, hands-on experience
- Establish a small core of key players within the organization to help promote corporate-level acceptance
- Allow a graceful rollout of new services over time
- Protect the initial PKI investment as new services are offered
- Determine whether or not PKI technology is viable and if it can offer significant cost savings
- Allow additional time to achieve corporate-level "buy-in"

It is expected that most enterprises will continue to proceed with caution. PKI vendors will need to work more diligently in the deployment area in order to help these small-scale pilots evolve into a more comprehensive, enterprise-wide security solution.

Repository Issues

Many enterprise domains utilize an on-line repository to allow for the timely and robust dissemination of certificates, certificate revocation information, and other PKI-related information, such as policies and procedures. Early experience in PKI deployment has demonstrated that it is not without problems. These issues are expected to be corrected as the products offered by the vendor community continue to evolve.

Lack of an Industry-Accepted Standard

One concern with directory services is a single accepted industry standard for offering these services has not been developed. Some market segments have adopted the ITU Recommendation, "Information Technology, Open Systems Interconnection, The Directory: Overview of Concepts, Models and Services," but numerous repository-related standards have been or are in the process of being developed. For example, LDAP (based on a simplified X.500 DAP) was developed under auspices of the Internet Engineering Task Force (IETF) and defines an access protocol between a client and a remote repository.

Issues remain open on functions associated with both client-to-server and server-to-server interaction and information exchange. For example, the IETF LDAPext Working Group (WG)[5] is developing additional standards, such as access control mechanisms and access control models. Also, the IETF LDAP Duplication/Replication/Update Protocol (LDUP) is under development, which may compete with the X.500 counterpart, the Directory Information Shadowing Protocol (DISP). Certificates and CRLs can also be distributed as part of a Domain Name System (DNS) function. Although any one of these solutions may be well suited for a given organization's needs, the breadth of choices can lead to interoperability difficulties when inter-organizational communication is required. Selecting

[5] The LDAPext group falls under the heading of the Applications Group of the various IETF working groups. This was concluded in 2003 when the various internet drafts were published into RFCs. The LDAPbis WG (also under applications) is now charged with getting the various RFCs through the standards track. See http://www.ietf.org/html.charters/ldapbis-charter.html and http://www.ietf.org/html.charters/OLD/ldapext-charter.html for more details.

standards-based solutions will help reduce some of these problems.

Multi-vendor Interoperability

In addition to the standards issue, multi-vendor interoperability is still a problem. Not all directory products are created equal. Experience has demonstrated that all vendors do not implement all functions within a standard, nor are functions implemented in a consistent manner from one vendor to another. These variances usually decrease with maturity, and most vendors appear eager to eliminate them by cooperating with their technology partners and customers.

Scalability and Performance

While not an actual barrier, scalability and performance issues associated with the deployment of a repository service are largely unknown at this time. Given the limited number of large-scale PKI deployments, implementation experience is not deep with respect to the data architecture and how many repository servers are required for effective performance for a specific organization.

Conclusions

PKI offers many services that can benefit the Federal Government. It can enable the Government to implement services such as Virtual Private Networks and 802.11 wireless networks without the fear of unauthorized users being able to intercept data.

Deployment of PKI may not be reasonable until an industry standard, including a multi-vendor interoperability standard, are developed and documented. Given the very few large-scale PKI deployments, solid lessons learned information is scarce making it difficult to anticipate potential deployment issues. Decision makers will resist deploying PKI until they know how much a deployment will cost in terms of dollars and other resources, such as infrastructure and personnel.

References

[1] Adams, Carlisle, and Lloyd, Steve, "Understanding PKI: Concepts, Standards, and Deployment Considerations," Second Edition, Addison Wesley Publishers, 2002.

[2] Leung, June, et al., "PKI Deployment – Business Issues," A PKI Forum Note, February 2003.

[3] "The Starter PKI Program," Thawte Technologies, July 1, 2004.

[4] Housley, R., et al., "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," IETF RFC 3280, April 2002.

[5] Wahl, M., et al., "Lightweight Directory Access Protocol (v3)," IETF RFC 2251, December 1997.

[6] "Email Security – The IBE Advantage," Voltage Security, Inc., July 1, 2004.

[7] Kiran, S., et al., "PKI Basics – A Technical Perspective," A PKI Forum Note, November 2002.s

[8] ITU-T X.509 Recommendation, "Information Technology – Open Systems Interconnection – The Directory Public Key and Attribute Certificate Frameworks," June 2000 (equivalent to ISO/IEC 9594-8, 2000).

[9] ANS X9.30-1997, "Digital Signature Algorithm."

[10] ANS X9.31-1998, "Digital Signatures Using Reversible Public Key Cryptography for the Financial Services Industry (rDSA)."

[11] ANS X9.62-1999, "Elliptic Curve Digital Signature Algorithm (ECDSA)."

[12] Diffie, Whitfield, "The First Ten Years of Public-Key Encryption," Proceedings of the IEEE, Vol. 76, No. 5, May 1988.

[13] Diffie, Whitfield, and Hellman, Martin, "New Directions in Cryptography," IEEE Transactions on Information Theory 22 (pp. 644–654), 1976.

[14] Bailey, Stuart, "Gaining Control of Your Network Identity Infrastructure," Infobox, April 6, 2004.

[15] Federal Information Processing Standards Publication (FIPS) 186-2, "Digital Signature Standard (DSS)," U.S. Department of Commerce, May 19, 1994.

[16] Lareau, Patricia, "PKI Basics – A Business Perspective," A PKI Forum Note, April 2002.

[17] ITU-T X.500 Recommendation, "Information Technology – Open Systems Interconnection – The Directory: Overview of Concepts, Models and Services," February 2001.

[18] Federal Information Processing Standards Publication (FIPS) 113, "Computer Data Authentication," U.S. Department of Commerce, May 30, 1985.

For further information, please contact:

National Communications System
Technology and Programs Division (N2)
P.O. Box 4052
Arlington, VA 22204-4052
(703) 607-6200