

realtimepublishers.comtm

The Shortcut Guidetm To



Managing Certificate Lifecycles



Kevin Behr

Introduction to Realtimepublishers

by Don Jones, Series Editor

For several years, now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtimepublishers.....	i
Chapter 1: The What and Why of PKI.....	1
It's a Matter of Trust.....	3
PKIs.....	4
Digital Certificates.....	5
What Does the Certificate Contain?.....	5
CAs.....	6
Are All CAs the Same?.....	8
Privacy.....	9
Encryption.....	10
How Encryption Is Used to Protect Privacy on the Web?.....	12
Authentication.....	13
Data Integrity.....	14
Non-Repudiation.....	15
Putting It All Together.....	17
Summary.....	18
Chapter 2: Root Management.....	20
Certificate Policy and Certificate Practice Statements.....	20
The Components of Certificate Policies and CPSs.....	25
Other Documents.....	26
Do You Need to Develop a Certificate Policy and/or CPS?.....	27
Who Should Be Involved?.....	28
Access Restriction.....	30
Backing Up.....	31
Audit.....	35
Summary.....	38
Chapter 3: The Certificate Lifecycle.....	40
Issuance.....	40
Contrasting Validation and Verification Procedures.....	42
Higher-Assurance Certificates.....	49
Managing Multiple Certificates.....	50
Re-Issuance.....	52
Expiry.....	52

Renewal.....	53
Revocation	54
Summary	55
Chapter 4: Managing PKI Infrastructures.....	56
The Requirements for an Effective PKI.....	56
PKI Requirements.....	57
The Threats to Your PKI.....	57
Building a Network of Trust.....	58
The Role of Certification Authorities	58
CA Systems.....	60
Digital Certificates	61
Elements in a PKI	62
CA.....	63
Registration Authority	63
Certificate Server	63
Time Stamping Service.....	63
Revocation Authority.....	64
Recovery Authority.....	64
Personal Security Environment.....	64
End Entities.....	64
Certificate Management.....	64
Initialization	65
Publication of Certificates and Revocation Lists.....	65
Key Recovery.....	65
Revocation	66
Certificate Issuance.....	66
Certificate Servers.....	67
Directory Services.....	68
X.500 Directory Specification	69
Certificate Servers and Revocation Lists.....	71
Checking Revocation Status Online	71
Using Revocation Lists	72
The Design and Implementation of a PKI	74

Requirements Analysis74

Design75

 In-House vs. Outsourced PKI75

 Choosing a CA76

 Choosing Components76

 Establish Practices76

Pilots77

Roll-Out and Monitoring77

Summary77

Download Additional eBooks from Realtime Nexus!77

Copyright Statement

© 2006 Realtimerepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimerepublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimerepublishers.com, Inc or its web site sponsors. In no event shall Realtimerepublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimerepublishers.com and the Realtimerepublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimerepublishers.com, please contact us via e-mail at info@realtimerepublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 1: The What and Why of PKI

Digital certificates are the central component in a Public Key Infrastructure (PKI) used to protect personally identifiable information, prove that online merchants are authentic, and protect the integrity of online transactions. Yet, many people have never even heard of digital certificates. They are buried deep inside many applications and technologies in today's Web-powered world, which most people take for granted.

If you have ever shopped for a certificate, you know that there is a wide selection of products and vendors from which to choose. Knowing what you need and, more importantly, why you need it, can be pretty confusing—even for a seasoned professional. This guide to managing the certificate lifecycle will cover a range of topics surrounding digital certificates, with an eye towards giving you the inside track when it comes to making decisions about PKI. This guide is for both those new to digital certificates and for technologists with extensive experience.

This first chapter provides an intro to PKI—think of it as PKI 101—and will introduce key technologies and concepts found throughout the rest of the guide. This chapter will explore:

- Public and private keys
- Roots and chaining
- Digital certificates
- Certificate Authorities (CAs)
- Secure Socket Layer (SSL)

Chapter 2 focuses on what you need to know about Root Authority Management, looking closely at the roles of

- Backup
- Access restriction
- Audits

Chapter 3 examines the entire lifecycle of a digital certificate, including:

- Issuance
- Reissuance
- Expiry
- Renewal
- Revocation

In addition, this chapter takes an in-depth look at best practices around authentication and verification of digital certificates.

Chapter 4 concludes the guide with a close look at the decision-making process when choosing between building an in-house PKI and third-party outsource management for your organization. This chapter covers topics such as

- PKI
- Processes and controls
- In-house management
- Outsourcing

As a reader of this guide, it is highly likely that you have information that you need to secure. It is also likely that, whether knowingly or not, you have used digital certificates during an e-commerce session. Every aspect of today's e-commerce that contains personally identifiable information should be protected by PKI.

At first glance, digital certificates may seem to be confusing and a subject best left to the hardcore specialists—but have no fear, this guide will break down the definitions and concepts into layman's terms. We will begin to unravel the secrets of this amazing set of technologies in this chapter by defining many of the key underpinning technologies and how they contribute to the overall network of trust.

If you have ever embarked on the path to building an e-commerce Web site, you may have been told by a consultant or vendor that you need to purchase a digital certificate. You probably have even asked someone why such a certificate is necessary. Let's start with the basics and build on them to create a more complete picture that explains the what and why of PKI.

It's a Matter of Trust

For two parties to engage in commerce, it is necessary to satisfy several basic needs beforehand. Many people like to get to know a retailer before they actually buy from it. Some take comfort in the size or ubiquity of a particular chain as a measure of comfort when doing business. These needs are particularly important when it comes to commerce over the Internet. There have been many studies done by think tanks and analysts to better understand Internet buyer behavior. A Webwatch study found that less than 30 percent of Internet users trust Web sites that sell products or services (see Figure 1.1).



Figure 1.1: Webwatch "It's a Matter of Trust" study (Source: "A Matter of Trust: What Users Want From Web Sites" at <http://www.consumerwebwatch.org/dynamic/web-credibility-reports-a-matter-of-trust-abstract.cfm>).

With recent headline disclosures of information theft and the constant lure of Internet criminals and phishers, there has been much focus on making the Internet a safer place to do business. The term *phishers* describes scam artists who use fake emails or instant messages to attempt to acquire sensitive information such as passwords and credit card information by masquerading as financial institutions or businesses you are familiar with. By using social engineering techniques, phishers attempt to defraud legitimate customers into disclosing confidential information. It is safe to say that everyone is concerned about the protection of personally identifiable information.

Fraud is still the largest impediment to the growth of e-commerce. Perception may be worse than reality, but the numbers are real. Online fraud overall is higher than its traditional counterpart. Management consulting firm Kinsey, and others, show that the barrier to online shopping is trust for many users. Ernst and Young, an international accounting and consulting firm, studied shoppers in nine European countries and found that "honesty, respect, and reliability" were the most important values concerning Internet shopping.

How can an infrastructure of trust be created on the Internet? Much thought has been focused around this concept. Today's PKI choices are a working solution for many of these issues.

PKIs

Let's take a closer look at the definition of PKI. PKI is the application of cryptography to ensure privacy, authentication of entities, assurance of transactional integrity, and guaranteed non-repudiation when transferring commercial information. The term is often used to describe everything from the discrete components that make up the whole package of PKI to the related vendors and the products that they sell.

PKI is the term commonly used to describe the following elements:

- Third-party vouching or vetting services—Essentially a resource that has investigated the party in question and has sufficient evidence to believe the party is in fact who it says it is.
- A method of binding these vetted or proved identities to keys, which are commonly referred to as digital certificates.
- Certificate Authorities (CAs)—The entities responsible for providing the digital certificates.
- Registration Authorities—The entities responsible for performing the due diligence on requesters of digital certificates, then issuing the certificate.
- Validation Authorities—Those responsible for looking up the validity of a particular certificate and, in essence, verifying its authenticity and validity.
- Applications—Various technologies, such as SSL, that depend on PKI.

Like many other Internet technologies, such as domain name services and email, the larger PKI is made up of many autonomous PKIs that interoperate according to industry standards. There is no actual global PKI; rather, PKI is used to refer to all the individual PKI systems. Some of these systems are commercial in nature and some are operated by government, research, and educational institutions. The standard that governs much of PKI is the International Telecommunications Union (ITU) standard X.509. The standard was defined by the Internet Engineering Task Force (IETF) PKI X.509 group or PKIX and adopted as a standard by the ITU. There are many PKIs that together form an infrastructure of trust on the Internet. PKI technology also is used to secure intranets or private networks designed for employee-only access. This guide will refer to PKI as a system based on public/private key cryptography that manages digital keys.

 For more information about X.509, it is published as ITU recommendation ITU-T X.509. You can find more information about PKIX through the PKIX IETF working group at <http://www.ietf.org/html.charters/pkix-charter.html>.

Digital Certificates


Digital certificates are often referred to as public key certificates. What does this term mean? Earlier, the chapter mentioned that PKI is rooted in cryptography. The subject of cryptography can be quite complicated, but for this guide's purposes, digital certificates are actually quite simple. Think of them as you would a passport document. To get a passport, you must satisfy the requirements of your country to prove your identity. Usually these demands are satisfied by presenting a photo ID and some sort of proof of birth. With digital certificates, the CA performs some basic checks to ensure that the requesters are indeed who they say they are.

After all, if it is the CA who is going to be asked to vouch for that particular requester, the requester must satisfy the CA's requirements in order to do so. No commercial CA would be in business for long if they vouched for criminals. Once the authority is satisfied with your credentials, they give you a digital certificate. It is an electronic document that is difficult (not impossible) to tamper with or use outside the scope of the person to whom they issue it (more on that later).

What Does the Certificate Contain?

The certificate contains a digital signature of the CA that says the CA has vetted the identity of the certificate holder for the particular uses outlined in the certificate. The certificate also contains the public key of the certificate holder. This key allows other parties to send the certificate holder encrypted messages or transactions that can be read only by the party with the corresponding private key, which is the certificate holder. Every digital certificate has two keys bound to it. The first is a public key, and the second is a private key held only by the certificate holder (see Figure 1.2).

Think of the keys as having separate functions. One is used to lock a door, and the other is used to unlock it. These types of keys are called asymmetric-keys.

 Certificate holders will be discussed in more detail later in this chapter, and the concept of asymmetric keys will be covered in-depth in the section on privacy and encryption.




Figure 1.2: The digital certificate registration process.

There are three classes of digital certificates:

- Class 1 certificates are commonly used for individuals needing a public key for low-value or non-commercial communications. Typically, the only validation of the requester is that the email address or user is not already associated with a key at the CA. These keys are the least expensive option as they have no commercial value only private value.
- Class 2 offers an increased level of assurances and are suitable for medium-value transactions and commercial communications. Typically, the validation process includes a comparison of submitted business information with records or databases containing business registration information. These certificates are available from most CAs for a modest fee as they still are not the most desirable for commerce applications due to their lower level of assurance.
- Class 3 certificates offer the highest level of assurances and are intended for individuals, organizations, and devices. Usually, the validation processes are much more stringent than those for class 1 or 2 certificates. Steps are taken by the CA to verify the existence of the organization, and the identity of the requester must be proved as well. Often, the checks include verification of the authorization of the requester by the company applying for the CA. In addition to these validation checks, the CA checks whether the applicant is entitled to use the domain name listed in the certificate application and whether the requester is authorized to manage the domain. These certificates are the most expensive of the three classes due to the extra work it takes to validate the certificate.

An X.509 version 3 digital certificate contains:

- Owner's identifying information
- Owner's public key
- Dates that the certificate is valid (starting and ending)
- Serial number
- Certificate type (level of assurance)
- The issuing CA's name and signature

CAs

CAs are third-party organizations that register digital certificates and bind them to individuals and companies. The CA binds the identity of an individual or an institution to a certificate by signing the certificate with the CA's public key. Much in the way a Notary Public would seal or sign a document and attest that the parties signing are all actually who they say they are, CAs use differing techniques to verify the validity of the certificate applicant at the time of application.

It is important to note that the methods of verification differ among the various CAs, with some offering additional assurances. If you click the button labeled "Issuer Statement" in the browser dialog box, you will notice a legal document outlining the responsibilities of the root CA, the subscriber, and you as the "Relying Party." Most CAs have strict legal language governing the certificates that they issue. It is also important to note that the CAs place the responsibility on you the consumer (or Relying Party) to verify that the information in the certificate is valid and that you feel comfortable trusting both the CA and the merchant.

A Word About Agreements

All CAs have several legal documents in place to both limit their liability and serve as roadmaps for navigating the work of PKI trust networks. The following list highlights some of the key documents you should read before purchasing a digital certificate (Figure 1.3 illustrates the digital certificate trust model):

- Relying party agreement—This agreement is between the person or organization that is depending on the CA to vet a company or individual. Most of the agreements start with a statement in all capital letters, such as YOU MUST READ THIS RELYING PARTY AGREEMENT BEFORE VALIDATING A (insert company digital certificate product name here) CERTIFICATE FROM XYZ. These agreements outline just what it is that the CA is warranting for the certificates it issues. It is very important to understand the applicable laws of your state and country, as they will affect how the agreement is interpreted.
- PKI disclosure statement—Many CAs have a statement of disclosure that outlines what actual corporate entity is providing which services and what levels of assurance are offered for which products the company offers. This statement is an important tool to use to compare the actual products because many CAs use different names for similar services. The statement will outline assurance or reliance limits as well as the obligations of subscribers and relying parties.
- Certification Practice Statement (CPS)—This document contains the assertions of the company management team with regard to the practice employed by a CA in providing its services. These documents can prove to be very lengthy—some are well over 70 pages long! The CPS is usually mentioned in both relying party agreements and often in the disclosure statements, so it is a good idea to look over the document to make sure you understand any liabilities it may convey to you as a subscriber. The CPS can be one way to compare CAs in terms of the policy and procedures that govern their respective PKI trust networks. These agreements are fairly easy to find on most CA Web sites; if you ever have trouble finding them, try doing an advanced Internet search limited to the domain name of the CA you are investigating. A search for “Certification Practice Statement” should give you what you need. If this search still proves fruitless, a call or email to the CA itself will get you what you need.




Figure 1.3: The digital certificate trust model.

Are All CAs the Same?

CAs are not all the same. Some CAs are known as root CAs, and they form the backbone of PKI. They are so central to the web of trust that no one needs to verify them. In fact, many third-party non-root CAs use what are called *chains of trust* to validate themselves as CAs (see Figure 1.4). Think of this idea in social terms. Imagine you had a friend named Billy that was acquainted with someone you wanted to date named Pat. You might ask Billy to arrange an introduction. Of course, you could always approach Pat directly and name drop Billy to borrow his credibility, but this option is seldom the preferred method. Instead, it is more likely that you would ask Billy to create a scenario in which he could introduce you to Pat. This situation is, in essence, the same concept as a chain of trust—Billy vouches for you to Pat. A CA that everyone trusts vouches for an unknown third party, legitimizing the party's claim. If that party is a CA, this chain of trust could extend below the third-party CA into the other CAs that they trust.





Figure 1.4: The chains of trust between CAs.

A Helpful Shortcut

When buying a certificate from a CA, it is a good idea to make sure you are dealing with a root CA. You can do so by looking at a certificate issued by the CA. All certificates contain a certification path tab, and if there is another CAs name above the tab in the certification path, you are looking at a chained CA—meaning the CA needs a root CA to vouch for it. To find a list of trusted root CAs in Internet Explorer (IE), for example, from the Tools menu, select Internet Options, then select the Content tab, click Certificates, and navigate to the tab labeled Trusted Root Certification Authorities. As the PKI information of all root CAs is built-in to most commercial Web browsers, your certificate is most likely to work as planned with this known and trusted PKI. CAs that depend on browser updates and chains of trust to offer their certificates are more prone to user experience issues than root CAs.

CAs differ in validation and verification methods as well. It is important to note that CAs all do a basic level of validation for certificates they issue. Some use automated verification methods that query public record databases to determine whether

- The company is legitimately registered
- The company owns the domain name on the application
- The requester is authorized to make changes to the domain name in question

 Beware! There are CAs that purport to check this automatically, and there is little evidence that it can be done reliably. Ask your CA how they authenticate these items and make sure you feel comfortable with the methods they describe.

Some CAs prefer to use more traditional manual methods such as fax and phone to gather this information. In addition, some CAs actually work to provide a higher level of assurance than just the basic validation mentioned. In the minds of many industry-savvy technologists, the assurance value of all digital certificates is not equal. The value (and cost) is higher for those that take extra actions in the various lifecycle steps, such as registration, to prove the true identity and legitimacy of the requesting entity and its right to assert control over the domains in the application. The value is also higher for those that take the appropriate steps to protect the CA's own private keys with elaborate processes and controls.

Privacy

All transactions that contain sensitive information must be kept confidential. Because the Internet is essentially a public place, a method must be set to ensure that the information exchanged is only human-readable by the parties engaged in the transaction. Continuing with the analogy of a public place, imagine two parties that need to communicate across a crowded room. They want their conversation to be unintelligible to the other people sharing the room. What should they do? In sports, it is a common practice to develop a secret code language, be it hand signals in baseball or code words used to describe plays in football. Some sort of cryptography is used to conceal this sensitive information so that it does not fall into the hands of the opponent. The method that emerged as the standard on the Internet was just that—a form of cryptography, or encryption.

You may remember stories from various points in history in which top-secret messages needed to be transferred to secret operatives in foreign territories. These messages were so confidential that if the information fell into the wrong hands, lives and even governments could be compromised. The answer to the spy problem was one and the same as the answer for the Internet privacy problem—encrypting the messages using cryptography.

The need for cryptography arises out of the fact that all data that traverses the Internet is carried in packets. Think of packets as postcards with your data written on the back. The information travels in the open and is available to anyone that can see the postcard. Like a postcard, the Internet is a public network, and just like any other public venue, people can eavesdrop on your communications. On the Internet, someone can intercept the packets that carry your information and inspect the contents for your sensitive information. The worst part is that you may never know this happened until it's too late!

Cryptography has several elements that are used in PKI. The first use is in converting the message to something meaningless if it is intercepted by a third party. In this case, the message must be transformed into something else during transit. But if the message is turned into something nonsensical during transit, how will the message convey anything of importance to the intended receiver? This process of creating an encrypted message uses a key—the “K” in PKI. Both sender and receiver have their own set of keys with which they encrypt the message. The method of encryption is called an algorithm, also known as a cipher, which is a way of scrambling and then descrambling the message.

A Simple Example of Encryption

A simple but effective cipher would be to shift the QWERTYUIOP keys on the keyboard one key to the left so that every letter is transposed to the letter immediately to the left. Every time the letter “I” is entered, the algorithm would decipher it into a letter “O.” The key would tell the receiver to undo this by shifting the letter back to the left, thus discovering the original message only intended for the receiver.

For example, the “I hope you are doing well” primitive cipher would look like “O Jp[e upi str fpmh er;”. Of course, this example is very simple and would easily be decoded by someone else, but provides an illustration of the basic process.

Encryption

Today, there are a variety of encryption methods to secure data. There are two primary groups in use. The first group is called symmetrical encryption because it uses symmetric keys. Symmetric keys are essentially identical keys or keys that use a shared secret. The most popular symmetric encryption algorithm is called Data Encryption Standard (DES). It was chosen as a standard for encrypting data in 1976 as a Federal Information Processing Standard (FIPS) for the United States government.

In fact, much of the work around encryption—and even PKI—started in the mid-seventies as the then-government and academic Internet was in its early developmental stages. The development of the DES was questioned by many academics as the National Security Agency (NSA) and other government agencies were involved. Some thought that it had “backdoors” that would allow government agencies to easily crack the messages and eavesdrop on any electronic conversations at will.

The original DES specifications were fairly weak. Academics were able to crack DES encoded messages in as little as 24 hours. With the emergence of Triple DES (3DES) and higher cipher bit depths that are commonly 128 bits, it is much more difficult to crack the ciphers used by commercial applications. There are at least four different methods to achieve 3DES. The most popular uses two encryption keys and the data is encrypted with the first key, decrypted with the second key, then encrypted again with the first key.

Confused? It’s easy to see why it is difficult to crack these ciphers. In fact, almost all the cracking done is in the realm of theory and is not practical, given the state of processing power at this time. As processing power increases, it will be very important to increase the complexity of the cipher methods. However, remember, there is also a fine balance between security and the amount of computing overhead introduced by encrypting streams of data. The computer must not only process the entire data stream but also decipher it in a timely fashion. The stronger the cipher, the more CPU time required to decipher the messages. No one wants to wait hours to get an urgent message or minutes to process an e-commerce transaction while they wait for their computer to decrypt a data stream.

DES and its modern variants such as 3DES are still in wide use, although it is important to mention that the United States government as of 2002 is now using a newer encryption standard called Advanced Encryption Standard (AES). This new cipher was invented by two gentlemen from Belgium—[Joan Daemen](#) and [Vincent Rijmen](#). AES is still not in wide use at this point, but many expect it to come into wider acceptance over the next several years due to its superior strength and lighter load on CPUs for encoding and decoding.

Symmetric key encryption is vulnerable to interception of the shared secret or key much more so than other methods such as asymmetric key encryption. Because of this vulnerability, distribution of keys is a major problem especially when one considers just how many keys would need to be distributed in the course of a day. It is imperative to protect the secrecy of the keys, thus, a different method must be used for PKI.

Asymmetric encryption is the type of encryption used for PKI, and it utilizes two keys—one public for encrypting the data and the other private to decrypt the data (see Figure 1.5).

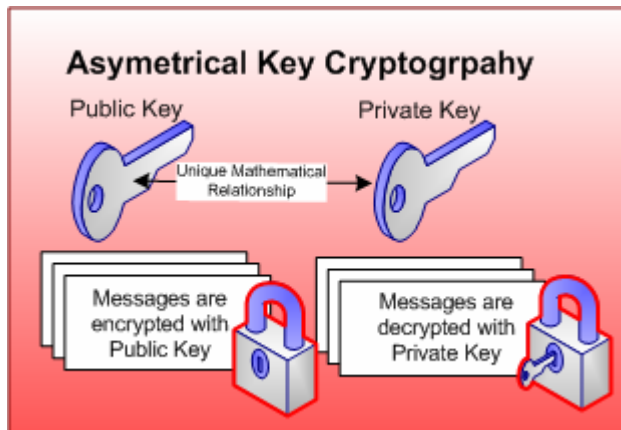


Figure 1.5: Public keys are used to encrypt data and private keys are used to decrypt it.

The advantage to asymmetric key encryption for public use lies in the ability to distribute a public key to everyone. This key allows anyone to encrypt data, such as an email message, and know that the only person who can decrypt the message is the holder of the private key. Although it is technically possible to determine the private key through factorization of the public key, it is extremely unlikely, given the amount of computational power required to do so in a reasonable period of time.

The first asymmetric key encryption algorithm was invented by Clifford Cocks in the early 1970s. Clifford was an employee of the United Kingdom intelligence agency GCHQ, and the invention was held in secret until 1997. There was an entire asymmetric key encryption system specification released by Whitfield Diffie and Martin Hellman. This publication outlined a public key exchange system complete with a public key agreement.

Since the advent of asymmetric key encryption in the 1970s, there has been a considerable amount of work done to further public cryptography. So much has been done that some researchers have even duplicated each others' inventions without knowledge of the original work. This happened in 1977 when Rivest, Shamir, and Adleman published a method identical to the work of Clifford Cocks. This would not be known for some time as his work was held in secrecy by the United Kingdom intelligence community until 1997. Although the modern implementation of this type of encryption is known for the combination of Rivest, Shamir and Adleman or as RSA encryption.

Encryption at a Glance

- **DES**—DES was developed by IBM. It uses a 56-bit key. DES is a symmetric key algorithm.
- **3DES**—The 3DES algorithm is a variant of the 56-bit DES. 3DES operates similarly to DES, in that data is broken into 64-bit blocks. 3DES then processes each block three times, each time with an independent 56-bit key. 3DES is also a symmetric key algorithm.
- **AES**—The National Institute of Standards and Technology (NIST) recently adopted AES to replace DES. AES provides stronger security than DES and is more CPU efficient than either DES or 3DES. AES offers three different key strengths: 128-, 192-, and 256-bit keys. AES is a symmetric key algorithm.
- **RSA**—RSA encryption uses asymmetric keys for encryption and decryption. Each end, local and remote, generates two encryption keys—a private key and a public key. To send an encrypted message to the remote end, the local end encrypts the message using the remote's public key and the RSA encryption algorithm. The result is an unreadable text. The remote end uses its private key and the RSA algorithm to decrypt the cipher text. The result is the original message. This RSA encryption technique is used for digital certificates.

How Encryption Is Used to Protect Privacy on the Web?

You may have noticed, if you visit an e-commerce site or a banking Web site, your Web browser displays a picture of a golden key in the bottom right corner. You may also notice that in the browser toolbar, the site's URL starts with an `https://` rather than `http://`. The `https://` denotes that the site is using SSL to encrypt the data stream between your computer and the Web server to which you are connected (see Figure 1.6).

If you click the picture of a lock in the lower right corner of your Web browser, you will be presented with a dialog box that contains a digital certificate. In this case, your Web browser has a list of root CAs built-in and has already verified that the certificate presented by the Web server you connected to is registered and valid with this list of verification authorities. This process is called authentication.

Your browser automatically inspects the certificate to verify:

- That the certificate is valid
- That the certificate is signed by a recognized CA—compared with roots already installed in the browser
- Ensures that the certificate is installed on the domain specified in the certificate details

This means that the certificate owner has proven that they are the appropriate party to receive your confidential information. Once your browser has verified that the certificate is valid, the browser then negotiates an encrypted session with the remote Web server.



Figure 1.6: What your browser will show you during an encrypted session.

Authentication

It is equally important to know who is on the other end of your secure data stream as it is to have the data encrypted. There is little point in scrambling data only to have it fall into the wrong hands. Authentication, most simply put, is the act of making sure the Web site; person, or institution you are dealing with is authentic. Authentication leverages the same PKI trust infrastructure to provide you with assurance via a CA that the person or institution that you are communicating with is indeed representing themselves correctly. This is done through the issuance and verification of a digital certificate. Remember that a digital certificate is issued by a CA and binds their identity to this electronic document. You can then check to see whether this document is valid and is in good standing with the issuing CA. You can even look at the standards for verification that the CA uses to decide how much you want to trust the entity with which you are working (see Figure 1.7).




Figure 1.7: By double-clicking on the gold lock icon in your Web browser you can view the contents of the digital certificate on a secure site.

Data Integrity

We have explored the importance of privacy and authentication; let's look at the role of integrity in the certificate lifecycle. Data integrity means that data received is exactly the same as data sent and that the message has not been modified in any way. This idea can also be applied to data values or messages that are transmitted over networks and stored on media such as tape backups or hard disks. This integrity must be able to be proved without regard to time or distance.

It is one thing to verify who you are doing business with, and another to make sure your communications are private. It is equally important to be able to provide a high level of assurance that the data you are looking at is exactly the same as the data sent to you. In the electronic world, this assurance can be accomplished via a variety of methods, with the most common method being a *checksum*.

A checksum is simply the result of a calculation performed on the basic components of the message, which is usually broken down into bytes. The bytes are then summed and the total included with the message to let the receiver know the exact message state the sender intended at transmission. The receiver can check the integrity of the message by performing the same calculation on the arrived message and comparing the actual value with the value included with the message. If they are the same, it is likely that the message has been untouched. If they are different, it is possible that the message has become corrupted during transmission or that someone or something has altered it.

This method, of course, has many weaknesses. It assumes there is no one with malicious intent that would have something to gain from modification of the message. It is very simple to change the position or value of part of the message and have it add up to the same total of bytes. That is why in commerce-grade applications, it is necessary to use a more sophisticated method that moves far beyond simple character or byte counts. You probably sensed that cryptography was going to be part of this increased sophistication and you were correct.

The need to provide higher levels of assurance far beyond the detection of accidental errors drove the development of many complicated cryptographic algorithms that not only total the amount of bytes but also document their position in the document. This algorithm or function is called a *hash algorithm*. A hash algorithm is a method for creating a summary or a digest of data such as the email example used earlier. The message would be passed through a hash function that would produce a fixed-length value as a result, this process is called *hashing*. This method is very similar to the simple checksum except that the hash algorithm assumes that there is an attacker that wants to modify the data. The hash algorithm does not merely add the numbers but applies a cryptographic table to the data and generates a unique fingerprint of the data in the form of a *digest*. This fingerprint is then attached to the message for comparison with the hashing fingerprint on the receiver's end. If they are the same, it is highly unlikely that the data has been modified (see Figure 1.8).




Figure 1.8: A slightly different text would yield a different hash sum or digest.

This level of integrity detection makes it much more difficult to modify data without detection. As is the case with most complexity, there is a cost. In this case, the complexity may provide a higher degree of assurance, but it also places a load on the processor in the same way that a complicated encryption cipher does.

Non-Repudiation

So far, we have explored setting up the security around the transaction and verifying the identity of the other party as well as talked about measures to protect the integrity of the data exchanged between the parties involved. The last piece of the puzzle is non-repudiation. What good is it if all the other pieces are in place, but after the transaction has occurred, one of the parties denies that they had anything to do with the transaction? This would certainly crush the appeal of online commerce if folks could just say they never bought the item in question and nobody could prove otherwise.

Ensuring that any transaction—which is, in essence, a legal contract between two parties—took place is just part of the challenge. According to the relevant ISO standards there are actually eight distinct areas governing the concept of non-repudiation:

- Approval—Non-repudiation of approval service provides proof of whom is responsible for approval of the content of a message
- Sending—Non-repudiation of sending service provides proof of who sent a message
- Origin—Non-repudiation of origin service is a combination of approval and sending services
- Submission—Non-repudiation of submission service provides proof that a delivery authority has accepted a message for transmission
- Transport—Non-repudiation of transport service provides proof for the message originator that a delivery authority has given the message to the intended recipient
- Receipt—Non-repudiation of receipt service provides proof that the recipient received a message
- Knowledge—Non-repudiation of knowledge service provides proof that the recipient recognized the content of a received message
- Delivery—Non-repudiation of delivery service is a combination of receipt and knowledge services as it provides proof that the recipient received and recognized the content of a message

There are two main components to non-repudiation that we are concerned with:

- Non-repudiation of origin—Non-repudiation of origin simply means that it can be proved that a particular party did, in fact, originate the transaction. By using authentication, it can be proved who is involved and their identity can be “signed” to the transaction. In the brick-and-mortar world, you are undoubtedly familiar with the concept of signing. Credit card companies and merchants must be able to prove that the correct individual initiated the transaction. Stores print a transaction receipt that has their merchant number and contact information on it. A description of the item sold and the cost with any additional taxes or surcharges is presented to the customer.
- Non-repudiation of delivery—On the receiving side, it is also possible to prove that the transaction was indeed received by similarly signing the end results. Typically, this is done by collecting the customer’s signature on the transaction receipt. This signature is then compared with the signature on the back of the credit card or other identification to prove that the customer is indeed authorized to make this transaction. If at any point in the future it becomes necessary to prove that a particular transaction was made involving either of the two parties, the transaction complete with signatures can be examined.

Putting It All Together

At this point, we have covered many of the foundational technologies and talked about how they work. Let's look at an entire example transaction to make sure you can apply these concepts in the real world.

It's December 17th, one week until the holiday season and you need to buy one more gift for someone special. You have exhausted the local stores and have found nothing suitable. Your search takes you to an online retailer—or e-tailer, as they have become known. Let's call the retailer The Underwater Acme Discount Superstore. You pull up your trusty Web browser and enter <http://www.theunderwateracmediscountsuperstore.com> and begin looking through their massive selection of wares. On the third-page, you score a direct hit and find an amazing marzipan fruitcake (made underwater) that your friend is sure to love. Almost too excited for words, you click the buy-it-now button and put the item in to your shopping cart. In your surprise and amazement at the huge selection Acme has, you missed the small golden lock that appeared in the lower right corner of your Web browser. It appeared the minute you added the item to your cart. Quietly, your Web browser received a certificate from the Web server when you navigated to the online store. In the background, your browser validated the certificate with the root authority that issued it. It checked to make sure the dates were valid and that it was not on the certificate revocation list of the CA. Once the match was made, the lock appeared and you did not need to worry about dealing with an imposter. Had the certificate not checked out, your Web browser would have presented you with a dialog box describing what the issues may have been with the certificate. Providing that you spelled the domain name correctly, you could be assured that you were on the authentic Underwater Acme Discount Superstore Web site.

From this point on, all your communications with the site were encrypted with the public key from Acme's digital certificate. The only people that could read your data stream are the folks at the Underwater Acme Discount Superstore—the holders of the corresponding private key. Once you have completed entering the shipping information and filling out a gift card for your friend, the transaction is turned into an invoice. This invoice has your payment information and the totals for the item plus shipping and any taxes. When you click to agree to the invoice, it is run through a hash and a digest is created. This digest allows the Underwater Acme Discount Superstore to attest to the authenticity and correctness of the invoice at any time in the future. They also can sign the invoice with their private key to further attest to prevent any repudiation later.

Summary

The role of digital certificates is central to providing critical assurances around commerce on the Internet. PKI truly creates a trust infrastructure that is enabling a multi-billion dollar e-commerce wave. With consumers having more concerns than ever about the safety and privacy associated with online transactions and credit card use, digital certificates play a key role in inspiring consumer confidence.

This chapter has covered the definitions of many key terms such as PKI, digital certificates, and root CAs. It also called out the difference between chained and root CAs and talked about some of the differences in the validation and verification processes between CAs. All of these discrete technologies work together to comprise a larger PKI that delivers:

- **Privacy**—All the data associated with a transaction or communication takes place in such a way that only the necessary parties have access to the data. This is satisfied by encrypting all data between the customer's Web browser and the merchant's Web server using SSL. For SSL to work in a way that inspires confidence from the user, an SSL certificate issued by a CA is necessary. Remember that this certificate should be issued by a root CA for maximum browser and application compatibility.
- **Authentication**—How does a consumer know that the Web site is authentic and operated by the company he or she thinks they have surfed to? The answer is authentication, and digital certificates provide the consumer with the assurance that they are on the right site and that they are dealing with the right folks. The secure lock in the lower right corner of their Web browser shows that the certificate presented by the site is valid and that a third-party is vouching for the authenticity of the Web site.
- **Integrity**—The critical elements of privacy and authentication have been established. Now, a method to guarantee that all data transmitted and received is exactly as it was at origin without any changes from a malicious party is needed. By utilizing a hash algorithm, the messages can be signed and sent with a digest that can be used to verify the integrity of the message contents.
- **Non-repudiation**—By using digital certificates to sign both the original and the delivered data, it is possible to prove that both parties were involved in the transaction. Coupled with integrity, this can be proven the same today or at any point in time—just as one could consult the merchant credit card receipt to prove that a consumer was present for and completed a sales transaction with their own signature.

PKI isn't the perfect solution to all the issues faced in an e-commerce powered world, but it does provide the basic foundation for trust and sufficiently lowers the risk of eavesdropping, malicious intent, and fraud on the Internet without being overly burdensome to implement. As encryption algorithms and PKI improve, you can expect increasing levels of assurance to make online commerce safer. With the amount of money at stake, you can rest assured that the security envelope will continually be pushed in the right direction. Many companies seeking to establish a brand awareness around the security of their products are innovating new more advanced security technologies. Some online banks and online securities trading companies are now issuing smart cards to their customers. These new security devices integrate with existing PKI technologies to form a greater assurance of identity and authenticity. Look for new applications of digital certificates, even in the consumer world. It is becoming more popular for corporations to issue certificates for all their employees in order to have two-way PKI in which both the site and the user are vouched for by a CA. The next chapter will dive deeper into certificate management with a look at root management.

Chapter 2: Root Management

This chapter will continue to build on the concepts outlined in Chapter 1 and extend them to common real-world scenarios. It will show you how to cut through the marketing hype and get to the bottom of the services provided by commercial CAs by introducing two documents that all commercial CAs produce. The chapter will then give you the inside track on methods you can use to manage your own certificate lifecycles by examining how commercial CAs manage their own

- Root keys
- Policy and procedures governing the keys
- End user agreements

Finally, the chapter will examine how the policies and procedures used by commercial CAs can map into your overall PKI management strategy, answering questions such as:

- How can you know what a CA does to secure their PKI?
- Who does the PKI serve?
- Is there a standard method you can use to compare different CAs?
- How are CAs audited?

This chapter will answer these questions as well as explore


- Access restriction,
- Backup of keys
- Auditing the environment to ensure operational integrity and reliability

Certificate Policy and Certificate Practice Statements

The first chapter mentioned that all commercial CAs are not the same as they have differing policies and procedures in place governing the lifecycle of the certificates that they issue and manage. Many CAs also have a range of certificate products that offer varying degrees of assurance and authentication based on the type of usage scenario required by the customer.

How do you best compare the products and services offered by various CAs? The answer is to look at what they actually do as opposed to what they market. The best way to understand the actual policies and procedures in use by a CA is to look at the two documents authored by CAs that describe how they manage their particular PKI. These documents are called the Certificate Policy and the Certificate Practice Statement (CPS). The Internet Engineering Task Force (IETF) PKIX working group, the group that is responsible for developing PKI standards, has created a standard document framework that CAs can use to make assertions about the practices and policies they use to issue and manage certificates.

These statements aren't just valuable for the prospective end subscriber of a CA; they are equally valuable to a relying party. In PKI, *relying party* refers to the person, company, or institution that relies on the information presented to them in the digital certificate. If you were to visit an e-commerce site to purchase something, you would be the party relying on the PKI to provide you with assurances of privacy via encryption and assurance of the authenticity of the Web site you are visiting via the information verified in the certificate presented to your Web browser.

 Remember, you can double-click the gold lock icon in the lower-right corner of your Web browser to see the contents of the digital certificate any time you are on a site for which the URL starts with `https://`.

The degree with which the relying party can depend on the accuracy of a digital certificate and its corresponding key is directly related to the particular policies and procedures of the issuing CA and the certificate holder (subject) of the certificate. The IETF RFC 3647 document authored by the PKIX working group serves as a framework for policies and practices documentation around the management of certificates. The Certificate Policy and the CPS and their suggested structures hold very valuable information that will come in handy when it is time to compare the practices of various CAs. The goal of the documents is to give interested parties a standard framework with which CA policies and practices can be compared.

 The IETF RFC 3647 is available at <http://www.ietf.org/rfc/rfc3647.txt> and the PKIX working group page is available at <http://www.ietf.org/html.charters/pkix-charter.html>.

According to the RFC, the Certificate Policy document is designed to act as a named collection of rules indicating the applicability of a digital certificate to a particular community of users. It can also refer to a specific application with common security requirements, such as which levels of verification are required by the CA for each certificate product it offers. Often these common requirements are tied to the monetary value of the transactions that will be authenticated by the certificate. The common requirements are not limited to commerce and can extend in to the types of end users or communities they support. These communities could be business units or academic research groups spanning multiple universities around the world. Whether evaluating a certificate created by an end subscriber or a certificate issued by a commercial CA, it is important to evaluate the assertions made by the CA to make sure that they match the application for which the certificate was presented. It is also very important that the procedures for validation and verification of the certificate subject are equal to the levels of assurance required by your particular application. If you are selling expensive items, you need to provide your customers with the highest level of assurance that you are authentic by purchasing the highest assurance certificate product from the appropriate CA.

When examining the various policy documents presented by a CA, you might find that there may in fact be several policies depending on the intended use scenario of the various certificate products offered. These differing policies reflect the needs of end-subscriber communities or corporate divisions. Inside of an enterprise banking network, for example, a one-size-policy-fits-all approach would not be the ideal approach. It would be wise to delineate the policy governing the management of certificates along the lines of risk and build separate policies for groups with differing risk profiles. For example, the bank's marketing personnel would certainly have different certificate needs than someone in banking operations that is responsible for transactional data. The bank's digital certificates could be managed with less transactional risk in marketing; therefore, the policies and procedures governing the certificate lifecycle can be much more lenient and geared to the speed of execution required by the business. This setup would be especially evident in the policies around issuance and verification of the certificate subject.

In the bank's operations group, the certificate lifecycle would be governed by stricter sets of controls around verification and issuance due to the increased risks associated with the processing of checks and payments.

The documentation framework outlined in RFC 3647 can also be used for creating documents besides the Certificate Policy and CPS. The format works well to construct relying party, subscriber, and other agreements outlined in a Certificate Policy or CPS (see Figure 2.1).

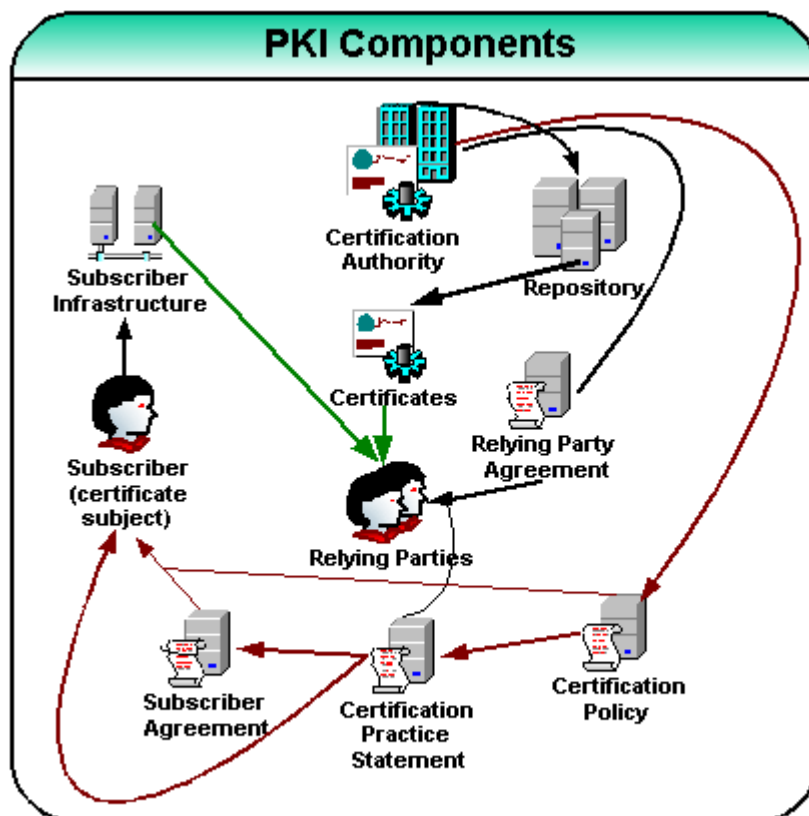



Figure 2.1: PKI components depicted with information flow. The green lines represent the flow of assurance from the certificate subscriber to the relying party. The red lines represent the flow of trust to the subscriber. The black lines represent the flow of trust from the CA to the relying party.

Think of the Certificate Policy as the basis for trust or accreditation in the issuing CA. When acting in the role of the relying party, it is wise to examine the certification policy associated with the certificate to make sure it matches the application for which it is being offered. For instance, when looking at the certificate type, it should match the type of transaction you are performing. If you are spending large sums or handling sensitive information, you want to be offered the highest level of assurance by a Class-3 certificate. If you are using a Web email service and just want basic security and privacy, a Level-1 or Level-2 certificate should be fine.

The CPS is an assertion by the CA covering the practices used by the CA to issue certificates. It is important to note that although the format of the document is outlined by an RFC, the actual content of the documents differ from CA to CA. The importance of understanding the very underpinnings of any network of trust cannot be overstated. The practices governing the issuance of certificates constitute the very cornerstone in the network of trust. If you are not comfortable that the CA can do what it purports to in its CPS, you need to consider another CA that offers services that match your needs more appropriately.

 A very sharp colleague recently shared such a scenario regarding the suspicious claims of a particular CA, which purported to automatically verify that the requester of a certificate was authorized to make changes to the domain name on their certificate application. The colleague pointed out that many domain name registrations commonly hide the names and information of the domain managers. So how could this information be verified automatically? The advantage of this automation was supposedly the speed of certificate issuance. Fast is not always the best—especially when dealing with security controls such as authentication. For high assurance applications of PKI, it is better to choose a CA that uses thorough checks, which might mean that they are performed manually by an actual person. It might take considerably longer to get your certificate and you might even have to mail or fax documents to the CA.

Key Terms in a Certificate Policy and a CPS

Activation data—Specific data values beyond the keys themselves that may be needed to activate cryptographic modules or access-control infrastructure. This data needs to be guarded, as it represents a security control in place to protect sensitive keys. The strength of the activation data and the strength of the access-control methods or devices must be commensurate with the sensitivity of the protected information.

CA certificate—A certificate for a particular CA's public key that has been issued by another CA.

Certification path—Certificates based on other certificates. These must follow an ordered, often hierarchical, path of certificates together with an initial certificate in the path that can be processed to validate the final certificate in the path (see Figure 2.2).

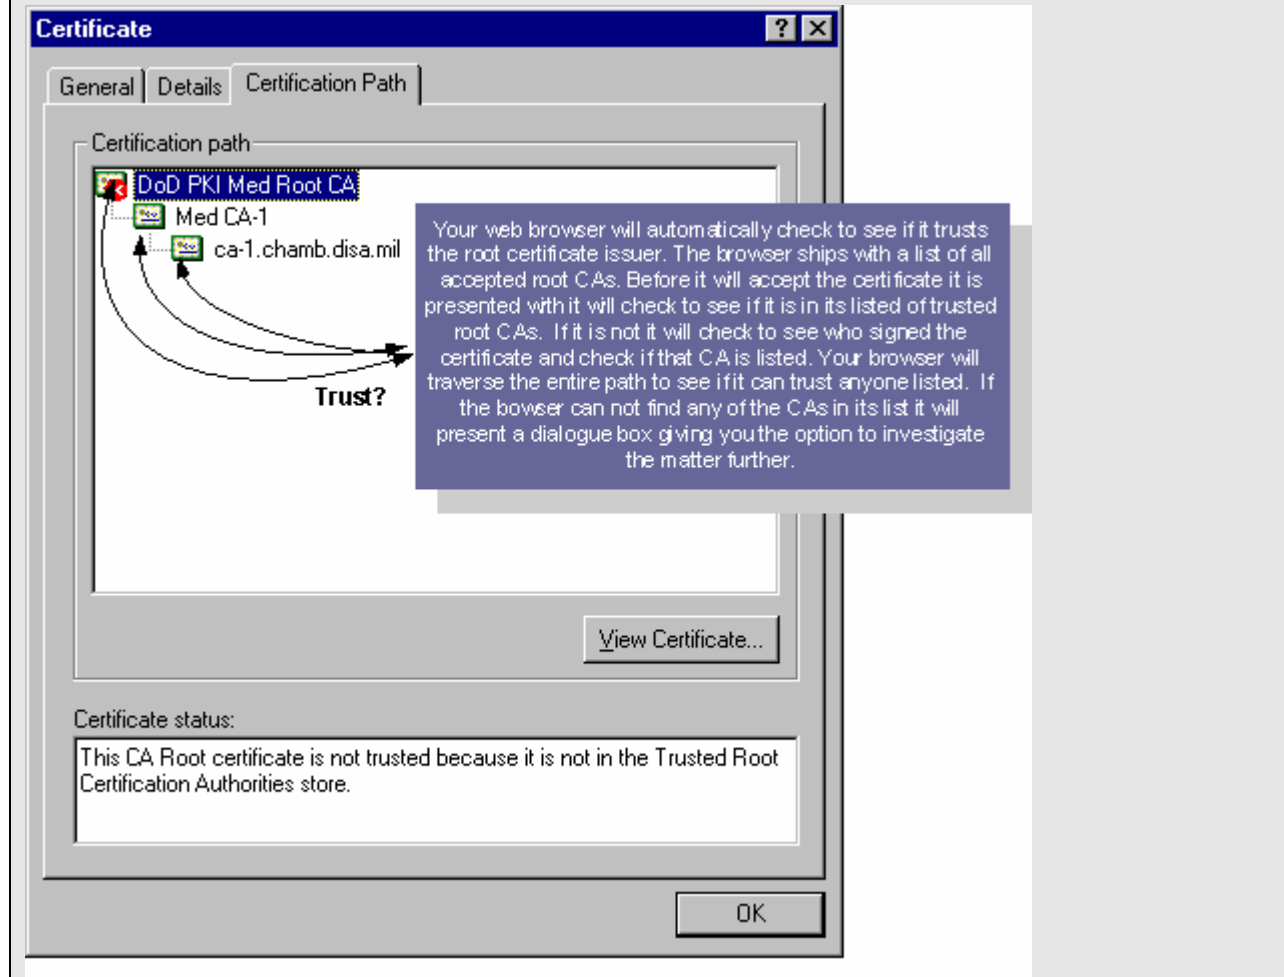
Certificate subject—The entity or end user that is requesting the certificate.

Issuing CA—The CA that issued the certificate, also referred to as the subject CA.

Policy qualifier—Information that accompanies an X.509 certificate in the form of a Certificate Policy identifier. The policy identifier is one of several data fields present in a digital certificate. This information might point to the URL at which a CA makes its CPS or Certificate Policy available.

Registration Authority—The entity that identifies and authenticates certificate subjects. This entity does not sign or issue certificates; the responsibility is usually delegated by a CA to perform these tasks on behalf of the CA.

Relying party—The party that receives the certificate and acts in reliance to the certificate directly or by other signatures verified by the certificate. This is also referred to as a being a certificate user.



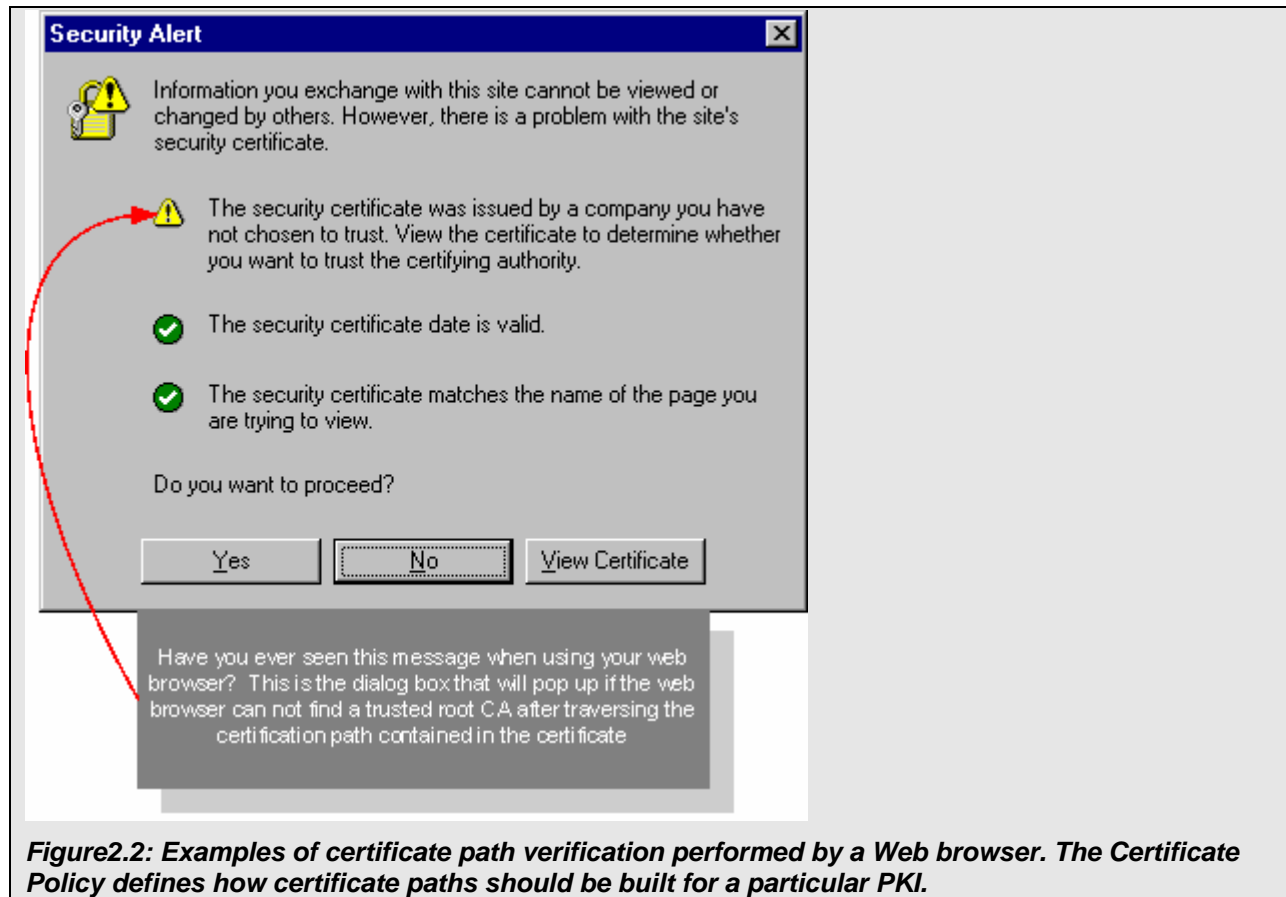


Figure 2.2: Examples of certificate path verification performed by a Web browser. The Certificate Policy defines how certificate paths should be built for a particular PKI.

The Components of Certificate Policies and CPSs


To find a Certificate Policy or a CPS from a CA, a good place to start is by visiting the CA's Web site. Most commercial and many private and community-specific CAs have links to their policy and CPS directly off their homepage. If you can't find the documents with ease, try using a search engine and query for the CA's name followed by CPS or Certification Practice Statement.

Once you have found the correct document, be sure to look at the date and version. The Web page on which CAs list their Certificate Policy documents and CPS usually have the latest document revisions at the top of the page in descending order with the older versions below. The CPS of many CAs is constantly evolving, so there are often several revisions.

The IETF RFC 3647 provides a basic framework or outline for the structure of these and other documents provided by CAs. The framework outlines nine basic elements. Not all elements are relevant to all CA applications or communities, so, in the course of reading documents authored by certain CAs, you might notice that some items are followed by the term *no-stipulation* for a specific element in the framework. In most cases, this simply highlights a provision that might not be applicable for that particular product or CA. This framework was designed to be used for all CA policy and procedure documents.

The RFC 3647 framework:

- Introduction
- Publication and Repository
- Identification and Authentication
- Certificate Lifecycle Operational Requirements
- Facilities, Management, and Operational Controls
- Technical Security Controls
- Certificate, CRL, and OCSP Profile
- Compliance Audit
- Other Business and Legal Matters

 More information about these nine elements is provided later in this chapter.

The main difference between the Certificate Policy and the CPS two documents has to do with their intended use. The Certification Policy document is intended to define the PKI and its community or application from a requirements standpoint. The point of this document is to instruct the PKI community members as to what they can or can't do within the confines of the PKI. Also the Certification Policy sets out the requirements and standards imposed by the CA on the PKI community. In contrast, the CPS focuses on how the CA and its participants must implement policy, procedures, and controls to meet the requirements that are spelled out in the Certificate Policy.

Other Documents

Although the Certification Policy and CPS are the central documents used to enumerate the policy and practice requirements of CAs, there are often several other documents that are very important. Not all Certification Policies or CPS are presented as legal contracts or even legally binding statements. Often contracts are handled in standalone documents that may be referenced in a CPS or Certification Policy. Other documents that may be referenced or even included within the Certificate Policy or CPS include:

- Subscriber agreements—These are the agreements that govern the purchase of certificates
- Relying party agreements—The agreements that stipulate what assurance guaranty is being provided to the person trusting a certificate by the issuing CA
- Security policies—The security policy documents for a particular CA
- Operational or training manuals—Documents that describe how to use the PKI of a particular CA
- Other standard documents from the IETF, ISO, other CAs, or a corporate entity
- Policies governing the use of email for subscribers and relying parties
- Key management planning documents—Templates that provide guidance and describe lifecycle considerations for certificate subscribers
- Employee manuals and guides—The internal Human Resources documents used by CAs
- A CPS abstract—An abbreviated version of a CPS that may omit sensitive or confidential information

The American Bar Association's Digital Signatures Guidelines Tutorial

The American Bar Association's section of Science and Technology Law has developed a free tutorial that can be accessed on its Web site at <http://www.abanet.org/scitech/ec/isc/dsg-tutorial.html>. The tutorial covers the basics of digital signatures and PKI as well as takes a broad look at the various legal implications of these technologies. Background into the nature of electronic commerce and a solid method for evaluating certificates are just two reasons to visit this great site. In addition, the guidebook can be downloaded in several document formats.

Do You Need to Develop a Certificate Policy and/or CPS?

If your organization is going to depend on digital certificates to deliver goods or services or issue them to employees or other business partners, you will need to develop a Certificate Policy and/or CPS. The RFC 3647 framework is a great place to start your efforts. As mentioned earlier, the RFC outlines nine sections that can be addressed depending on the scope of your PKI implementation. Are you running an e-commerce site and just need a certificate to provide assurance to your customers or business partners? If so, your documentation requirements will not be as elaborate and need not address certificate issuance, for example. If you plan to use PKI to issue certificates to all your employees, you will need to address the full certificate lifecycle in your documents:

- Introduction—This section is used to describe the company and how it intends to use PKI
- Publication and Repository—This section will answer the questions: How will you publish your public keys and how will you store the certificates you issue?
- Identification and Authentication—This section will answer the questions: What are the processes and controls you will use to authenticate certificate requesters?
- Certificate Lifecycle Operational Requirements—This section will answer the questions: How will you manage the infrastructure, policies, procedures, and controls around your PKI? Who will be responsible for what?
- Facilities, Management, and Operational Controls—This section describes what systems of controls you will build around physical and logical access to the facilities and infrastructure that house your PKI.
- Technical Security Controls—Use this section to describe the preventive, detective, and corrective controls that keep your certificates safe from others. For example, describe how you separate your production PKI network from your test and development networks and limit access to it. Also describe your defense-in-depth strategy for protecting the information assets in your PKI.
- Certificate, CRL, and OCSP Profile—Use this section to describe how you will publish a list of expired or revoked certificates to your PKI users.
- Compliance Audit—Describe your audit partner and frequency. Offer a brief description of the audit plan and objectives used as well as how the user can obtain a copy of the audit results.
- Other Business and Legal Matters—This section is used for general comments and/or items you might want to cover outside of the other sections.

Who Should Be Involved?

Assemble a working team comprised of the stakeholders in your organization that are involved with and affected by the use of PKI. Often the team will be lead by your security organization, as PKI is often part of a larger security policy or roadmap:

- IT operations—The ops staff can speak to technical and infrastructure questions that may arise.
- Sales/marketing—What is the business value of this technology and what business processes depend on PKI?
- Liaison to business partners—Often IT organizations have someone designated to address the concerns of business with IT and vice-versa.
- Legal council—They are crucial for offering decisions about what risks are acceptable in a general business context.
- Finance—If you are going to put dollar and cent values on the business processes that PKI enables, you will want their buy-in.
- Relevant division heads—Which business groups will depend on PKI? Whether they know it or not they need to have a voice and be part of the process.
- IT security—At most companies, PKI is a subset of a broader security policy framework. Security must be present to make sure that the PKI policy and procedures are integrated into this framework and that the provisions outlined in the documentation are consistent with the broader security roadmap. IT security may act as the PKI team lead.
- Physical security—It is very important to involve those in charge of physical building and area access. Commercial CAs use multiple levels of physical security to provide assurance that only authorized staff have access to systems used to issue or modify certificates. This is an important control consideration for enterprise use as well.
- Internal audit—It never hurts to bring your most knowledgeable control experts in if they have the time.
- Human Resources—If your PKI is used to authenticate employees or contractors, you will want HR at the table to make sure that the policies and procedures are consistent with employee manuals.
- Corporate compliance (HIPPA, GLBA, CFR11, FISMA, PCI, SOX) team leader—If your company is regulated, they will need to integrate any relevant policy and procedures into their control documentation.
- Business continuity/disaster recovery team lead—If PKI is going to affect your company's ability to maintain business continuity, you will need to integrate your work with the overarching business continuity and disaster recovery plans.

This group is also an effective team to determine the operational and corporate risks involving the use of PKI and certificates. The documents you create will most likely work with existing documents such as Human Resource employee handbooks, IT security policy, and IT change management or other operational run books. This is also a great team to classify the methods used to store and backup the certificates. Are the certificates so sensitive and mission critical that they need special protection? Ask questions to determine whether the certificates will be required to ensure business continuity. If third-party business partners utilize your certificates, you might want to consider authoring a CPS abstract that is an abbreviated CPS that contains only information that is relevant to business partners but does not disclose sensitive information regarding your security or HR policies.

It is also wise to establish an update plan that will address changing needs and scenarios and incorporate them into the documentation. A good rule of thumb is to make the evaluation of your PKI policy and procedure documents a part of an annual security policy review. If your PKI implementation will be used to protect email or enable secure e-commerce activities, it will be wise to construct an appropriate privacy statement for your employees, business partners, and customers; this statement is best handled in the context of your overall security policy framework.

As policy and procedure are meant to be living documents that are changed to reflect actual evolving practices, you will do well to consider revisions, distribution, and notification out of the gate:

- How often will this team or a smaller subset revisit the documentation to determine whether it still meets the needs of the business?
- How will your business partners or employees be notified regarding changes and updates to the various documents your company publishes?
- How will they know whether they are looking at the latest version?
- Where will these documents be stored?

These questions represent the belief that your PKI implementation will evolve as your business needs and partnerships change.

Rather than force everything into one huge documentation set you might want to consider each PKI use case as a separate implementation and create separate documentation. It does create a larger document base to maintain, but, if your organization is extending its PKI implementation beyond the walls of your company and into other companies, this will prove to be important. By the very nature of the Certification Policy and CPS documents, you might opt to protect sensitive references to internal security policy, authentication, and escalation methods from the eyes of anyone outside of your company. Abstracts of briefs can be prepared for each community or use case and should be represented by an internal-master Certification Policy that documents any and all certification policies and practice statements and the entities to which they pertain. This way, a regular review will cover both internal and external versions and all the documentation will receive the periodic scrutiny necessary to make sure the documents are still relevant.

Access Restriction

Another topic covered in most CPS documents is the policy of limiting access to PKI production systems, such as servers and network infrastructure that store and serve certificates to users. Many IT organizations have become very comfortable with pervasive server access by IT support personnel. This is often provided with customer service in mind. The rationale is that the less hurdles a systems administrator has to go through to solve a problem, the better. When it comes to access to live production systems containing digital signatures, this logic does not hold up. There have been several studies showing that the largest single cause of system outages is human error. By reducing the field of suspects, it is much easier to eliminate change or human error as a causal factor if something goes awry. If your PKI implementation is critical to the flow of business, you will want to consider removing persistent account access to the PKI servers and implementing a limited functional access only via approval. When the approved work is complete, access is then removed. Typically, during an effective IT audit, the auditors (internal or external) will want to see a list of user accounts on mission-critical servers. Make sure that you can account for every access account and answer why each account is enabled. Also consider the roles and responsibilities of your IT engineers. Do any of their official responsibilities make it more difficult for them to follow procedure and process?

Separation of Duties

CAs make sure that the roles and responsibilities of its employees do not compromise their ability to adhere to policy and procedure. A classic example in finance is the old rule about being able to add vendors to the accounting system and being able to print checks. This check and balance was instilled to prevent the AP clerk from adding a new “mystery” vendor and then printing the checks only to steal the funds for themselves. Checks and balances can also include signing limits or dual signatures for checks over \$500.00.

Preventing Unilateral Changes and Transactions

One of the largest causal factors behind IT service outages is change. Eighty percent of the time it takes to resolve an outage is spent determining just what changed! Many security breaches happen as a result of firefighting problems, as access controls are often bypassed to expedite troubleshooting efforts. Commercial CAs know that accountability for changes is key to preventing disasters and security incidents. For these reasons, commercial CAs have developed elaborate systems of control to prevent a single person from compromising the entire PKI.

CAs utilize security controls around the signing and storage of certificates and private keys. As CAs must protect their private keys at all cost, a formal signing ceremony of sorts is usually designed to reduce the risk of collusion or unilateral decisions that could prove to be fatal to a PKI. The next chapter will address the roles and responsibilities of those managing a PKI infrastructure, as the chapter looks at the certificate lifecycle components and what must be managed at each stage.

Strong Change Management Integration

Any proposed changes involving certificates or PKI should most likely involve security personnel. As the effects of PKI can be felt everywhere, on many corporate networks it is wise to vet all proposed modifications through a rigorous change management process. At a minimum, the change manager must obtain the appropriate approvals from the potentially affected business customers. If PKI is used by your sales and marketing team to sell on the Web, this is even more important, as any certificate modification or installation has the ability to affect the viability of e-commerce sites or employee network access.

Backing Up

If you have read some of the CPSs issued by CAs, you might have noticed that there is often a section that mentions that all critical keys are backed up in case of hardware failure and to address disaster recovery requirements. The documents state that they only back up their own certificates not those that belong to end subscribers. If your certificates are important to the everyday operation of your company, a plan to provide continuity of service should be put in place immediately. It is wise to inventory all the software and systems that depend on certificates to operate and make a list. If your organization has a business continuity and disaster recovery plan, find out how these systems and software are addressed. It is a great idea to look for underpinnings or new dependencies that might have cropped up due to changing requirements. A meeting with the person in charge of the overall disaster recovery/business continuity plan to discuss the role that certificates play in business operations would be a great place to start. Make sure that they understand that digital certificates may take days to provision and make sure that any documents used in the validation process with your CA are available in the event of a disaster so that if it is necessary you can re-validate your organization.

The most common form of backup in many organizations is the magnetic tape. Whether your organization relies on tape or other media for backup, be sure to understand who has access to this media. I have heard stories from more than one security professional that has casually lifted a backup tape out of the tape library and found certificates complete with private keys on the tape with absolutely no security. If someone is able to obtain your private key, they can issue certificates signed by your organization! This can prove disastrous, which is why backups of digital certificates must be subject to a system of controls that is designed to meet the assurance requirements of your company.

It is crucial to build a policy for the backup and safe storage of your certificates. Special provisions may need to be made to safeguard your high assurance for Class-3 certificates. When building this policy, make a concerted effort to quantify the damage possible if your certificate were to be compromised. The American Bar Association offers a wealth of information in a free PKI tutorial on their Web site that will prove helpful in identifying your risks (<http://www.abanet.org/scitech/ec/isc/dsg-tutorial.html>). Would there be a dollar loss or would the certificate give someone the magic keys to sensitive corporate data? Once you have a feel for the risks, it is much easier to design a protection plan that is commensurate with risk. For high-value certificates, you might want to consider some or all of the following controls:

- Consider the use of a FIPS-approved Hardware Security Module (HSM) to store your most sensitive keys (rather than a general-purpose server). It is a great idea to purchase a backup HSM in addition to your primary HSM. As HSMs typically cost thousands and can cost tens of thousands of dollars, they are most appropriate for high-risk scenarios or to protect very sensitive keys (see Figure 2.3).
- Use an offline backup server to store root or sensitive certificates. Control access to this server.
- Use change detection software such as Tripwire or Solidcore to monitor the integrity of your certificates. These types of software can create a fingerprint of your certificates and alert you if they are changed in any way.
- Encrypt your certificates and keys if they must be stored on a shared server.
- If certificates must be stored on tape, utilize tape encryption and control the physical storage of the tapes. The tapes should be kept in safe deposit boxes or controlled access scenario. Make sure that retention policies for longevity of storage and destruction are well thought out for these valuable items.

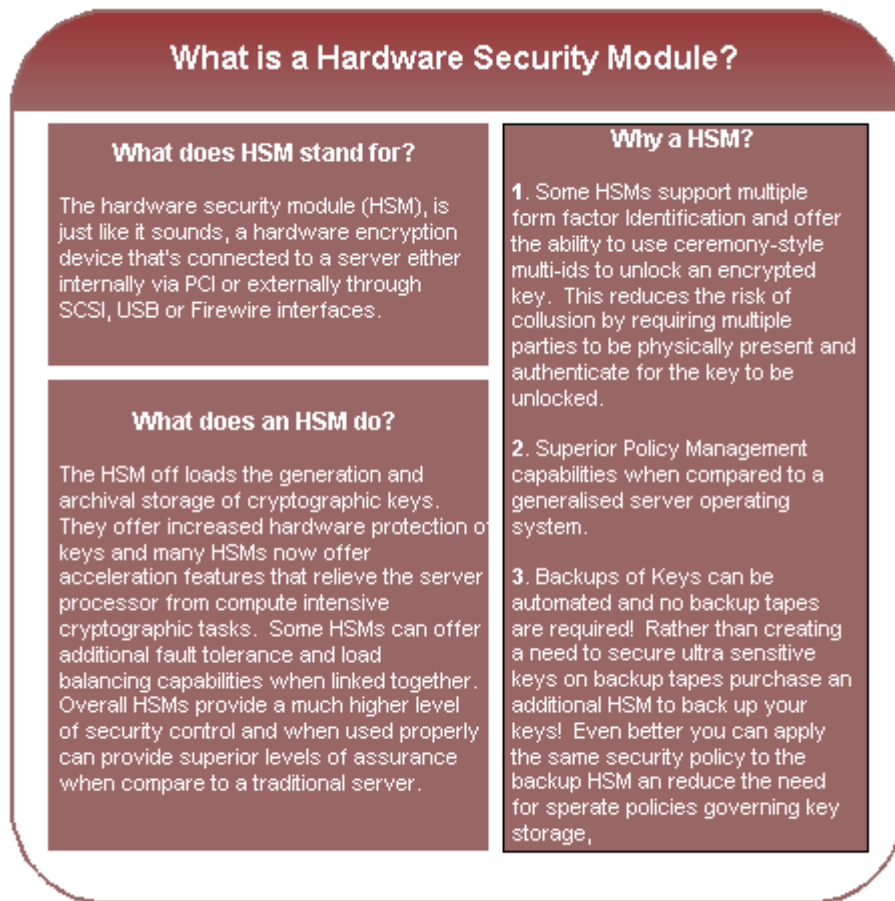


Figure 2.3: The definition and key reasons to use an HSM.

What is FIPS Certification?

The National Institute of Standards (NIST) has developed guidelines for security hardware manufactures and FIPS 140-2 (<http://csrc.nist.gov/cryptval/140-2.htm>) is the current specification for security requirements for HSM. Manufacturers must comply with the guidelines in order to sell their respective products to the United States federal government.

The specifications are dovetailed with a standardized testing methodology and an approved testing lab list. Once a vendor's hardware has been successfully tested by one of the approved labs, it is added to a list of approved solutions at the NIST Web site (<http://csrc.nist.gov/cryptval/>). This list is a helpful way to start evaluating competing HSM solutions.

It makes sense that commercial CAs go to great lengths to protect their private keys. If a commercial CA's private key were compromised, its entire PKI trust hierarchy could be destroyed—not to mention all the end subscribers validated by that particular CA (see Figure 2.4). Most commercial CAs go to very elaborate lengths to make sure that their private keys stay private. This is accomplished through the judicious implementation of security controls.

Controls common to commercial CAs:

- HSMs are used—Many CAs use HSMs to create and store their critical keys. These modules are capable of encrypting the keys and require an authorization process to be completed before the keys can be unencrypted and transferred to another device or server. In many cases, these HSMs are not statically connected to a network or, if they are, they are connected to an isolated production-only network with access granted only to those with an approved need.
- Backups are controlled—Backing up critical keys is a very serious business at CAs, and these backups are subject to rigorous controls and regular auditing to guaranty that this information is available only to authorized personnel. Limits are placed on retention and backups are audited regularly.
- Ceremony—Requiring multiple parties to be present with their part of the overall secret is called a ceremony. Many CAs use this same control method to guard the generation and signing of certificates and keys. This method makes it more difficult to compromise keys and certificates due to an increased number of people involved (reducing the risk for collusion).

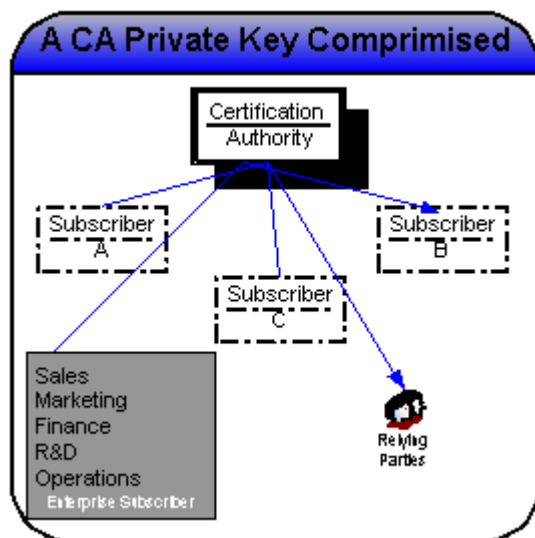


Figure 2.4: Everyone below the root CA in the PKI hierarchy is exposed if the CA's private key is compromised. In cases in which the compromised CA is a root CA, other CAs could even be compromised without their knowledge.

Audit

One of the most misunderstood terms among IT security professionals is audit. Many believe it means simply to check up on things or to investigate. Even worse, I have talked to many IT staffers that believe controls belong to the auditors. Controls are owned by the business. To be effective over the long-haul, high-performing IT organizations design controls to act as sensors that indicate whether current performance is furthering the goals of the business (keeping risk in check) or impeding them. In other words, controls are management. Audit is not a management function; rather, it provides an ongoing commentary on the effectiveness of controls in place and a list of controls that are missing (see Figure 2.5).

IT audit can be a critical part of managing adherence to critical controls such as policy and procedure. However, don't leave the development of policy, procedure, and controls to auditors—that is not their job! That is not to say that audits shouldn't review documentation. There are few people in the business world with a better grasp on process and controls than auditors, so leverage their expertise. Although you are gleaning from their knowledge, do yourself a tremendous favor and read up on control theory and IT audit frameworks such as Control Objectives for IT (COBIT—<http://www.isaca.org/Template.cfm?Section=COBIT6&Template=/TaggedPage/TaggedPageDisplay.cfm&TPLID=55&ContentID=7981>), which is published by the Information Systems Audit and Control Association (ISACA). Audit is best performed by trained auditors that have adequate independence from day-to-day operations and projects. In fact, the auditor ideally should report to a completely separate management chain. Independence gives the auditor an unbiased viewpoint from which to judge adherence with an organization's stated policies, procedures, and controls.

Controls 101

Preventive—Preventive controls focus on what can be done before an undesired action or condition occurs. A classic example of a preventive control is a lock on a door. In the case of systems control, functional access based on approved work is a preventive control. When compared with static access, functional access can drastically reduce the amount of unauthorized changes. If you can't access the server, it is much more difficult to implement a change on it. In PKI, a preventive control is the validation of a certificate applicant's identity preventing the applicant from falsely representing themselves.

Detective—Detective controls monitor an environment or a process and look for a certain set of conditions to occur, then generate a notification or alert that the condition indeed occurred. A great example of a detective control is a fire alarm. When a certain threshold of smoke and heat are met, the alarm emits a piercing sound to warn the building's occupants that there is a fire. In PKI, a detective control could be used to protect the integrity of a key by monitoring it against a baseline to detect any modifications to the certificate.

Corrective—Somehow an undesired behavior occurred and now you need to remediate the situation so that you can get back to performing as if nothing ever happened. A classic example of a corrective control is a steering wheel. Data backups are an example of a corrective control. If a certificate is corrupted or lost due to a hardware failure, the certificate can be restored from a backup.

Commercial CAs that are licensed are required to have their operations audited. The audits in the United States and Canada are performed under an AICPA/Chartered Accounts of Canada coauthored document called the Webtrust Program for Certification Authorities (this document is available at http://ftp.webtrust.org/webtrust_public/tpafile7-8-03forthefweb.doc). The basic principles serve as an excellent mechanism with which any organization can evaluate its performance against its intentions or policy assertions. Often, service providers and CAs will also undergo a Statement on Auditing Standards for Service Providers (SAS70) audit to validate the broader system of controls in place around their IT environments. The SAS70 was designed to make a public statement about a service provider's controls and what the auditor found to be true or untrue about those public assertions. Although the SAS70 and Webtrust audit programs do not serve as an effective method to compare one CA to another, as they will most likely have different assertions regarding their own system of internal controls, it is important that they successfully pass these audits. The main difference between the two audits is that the Webtrust audit has a defined framework and control set that is audited and the SAS70 that audits the service provider's control assertions and has no predefined framework. The SAS70 audit focuses on the control assertions that the service provider has decided to make public.

The main categories of Webtrust audit evaluation are:

- CA Business Practice Disclosure
 - Identify the Certificate Policy and CPS under which the CA issues certificates
 - What are the CAs certificate lifecycle practices?
 - Does the supplied documentation disclose all this information to subscribers and relying parties?
- Service Integrity
 - Does the authentication method employed by the CA effectively screen applicants in a repeatable fashion?
 - Does the CA effectively protect the keys and certificates it is responsible for managing throughout their lifecycle?
 - Does the CA publish revocation lists in a timely and accurate manner?
- Environmental Controls
 - Subscriber Personally Identifiable Information (PII) is kept confidential and is used in a way that is consistent with the promises and warranties outlined in the CA's privacy statement.
 - Operational integrity and continuity are sufficiently maintained
 - Systems are maintained only by authorized personnel performing work that has been authorized and vetted by appropriate control processes such as change management
 - Any software or systems developed are built using repeatable and authorized processes




Figure 2.5: Audit represents a critical component in a successful PKI implementation. Audit verifies that the foundation of trust promised by a CA is actually delivered.

In addition to the Webtrust (<http://www.ansi.org/>) checklist, there are a variety of sources for audit and control objective guidance. The ANSI X.9.79:2000 standard codifies the control objectives highlighted in the Webtrust CA audit into its PKI Practices and Policy Framework. There is no shortage of audit guidance from organizations such as the Information Systems Audit and Control Association (ISACA—<http://www.isaca.org>). There is much guidance available on their Web site free of charge. Simply use the search bar and enter PKI to get a great list of articles and guidance to peruse through. ISACA members also have access to a wide variety of audit checklists and specific audit guidance around security and PKI.

If your organization does not have a full-time internal audit staff, you might want to consider hiring or contracting with a third-party organization to audit your PKI implementation annually to make sure the policy, procedure, and controls you have invested in are in fact being followed every day.

Summary

This chapter has covered quite a bit of ground building on the PKI concepts covered in the first chapter, moving beyond basic certificate and digital signature technologies to how certificates are actually provided and looking at what needs to be done once a certificate is purchased. It is easy to get overwhelmed by the wealth of information and seemingly steep learning curve associated with the purchase of a digital certificate. The real work begins when it is time to depend on PKI. Building a dependable PKI solution requires serious thought and effort around policy, procedure, and controls. Implementation of digital certificates is not where the heavy lifting is done. Building out an effective system of controls to support the goals of the business requires involvement from many stakeholders and often represents the bulk of the work. Making sure that the digital certificate lifecycle is understood and effectively managed doesn't require a technical certification or advanced Computer Science degree but rather a focus on the goal of the business and how digital certificates can diminish constraints around commerce and security and open up a whole new avenues of business.

By choosing to see digital certificates in the context of a lifecycle, you are now free to focus on management, metrics, and controls in an enabling manner. Rather than view certificates and PKI as something you buy, implement, and forget about until a disaster occurs and strips your organization of its ability to derive e-commerce revenue through theft, fraud, or reputation loss, why not tackle policy, procedure, and audit when it is least expensive, both in monetary and emotional capital, early in the lifecycle. Nothing hurts worse than failures, outages, or miscues shortly after the fanfare of a new rollout. The common misconception of business and IT executives alike is that the steepest curve in any new IT project is getting the new technology to work. Seasoned IT operations engineers know that most common incidents could and should be prevented with simple controls designed into the project before it is ever starts. Nearly 80 percent of all IT outages are caused by people and process issues, according to research performed by Gartner.

Using the standardized policy and procedure document framework (RFC 3647) built by the IETF's PKIX working group, it is possible to compare critical components of the Certificate Policies and CPSs issued by commercial CAs. This chapter covered that a CA can actually issue several policies and practice statements depending on the PKI community they are addressing. As with the commercial CAs, it is important for you to consider whether abridged or specialized versions of your documentation should be developed for use with business partners or customers so as no to reveal confidential security practices and controls beyond the confines of your trusted company staff.

With a better understanding of the makeup of the Certificate Policy and CPS, it is also easier to make more informed decisions as a relying party. Especially when you can slice and dice the issuing CA's CPS to verify that the certificate you are presented with is the correct type for the application in use. Knowing how to strip all the marketing hype away and get to the bottom of the validation and verification methods employed by a particular CA not only helps you decide whether to trust a certificate but also helps you decide whether a CA's practices line up with the requirements of your application. By using the CPS as your guide to what a CA really does in the way of validation and verification, you can steer clear of misunderstandings and get the level of assurances you really need.

The next chapter will leverage these concepts and explore the different phases of the certificate lifecycle. It will explore how CAs manage these phases differently—some CAs provide adequate verification policies and practices but lack in other areas such as reissue or revocation.

Chapter 3: The Certificate Lifecycle

The first two chapters focused on defining and exploring key elements of PKI and how the root certificates and infrastructure need to be managed. This chapter will cover the entire functional lifecycle of a digital certificate. The certificate lifecycle consists of five distinct phases: issuance, re-issuance, expiry, renewal, and revocation. This chapter is full of helpful hints, best practices, and resources that will save you time, help you avoid embarrassing and expensive site outages, and steer you towards the correct certificate product for your application. Let's start at the very beginning of the certificate lifecycle with issuance.

Issuance

This guide has defined digital certificates as the binding of a vetted identity (company, role, or person) to a pair of digital asymmetrical keys. Whether purchasing a certificate from a commercial CA or issuing your own certificates within your company or enterprise, the concept is the same. Because a digital certificate is a form of credential used to identify a person, company, or a role, it is crucial that a validation and verification process is sufficiently rigorous to indemnify the level of assurance provided by that credential. For example, if you were a relying party shopping on an e-commerce Web site for a new laptop computer that costs \$1500, you would want to know that the digital certificate presented to you by the Web site guaranteed that you were not being spoofed. You would hope that the CA had stringent standards to make the merchant prove it was who it purported to be. The first chapter discussed the differing levels of assurance offered by most commercial CAs. The levels of validation and the verification methods are more rigorous as the level of assurance provided by the certificate product escalates. Keeping with the focus on delivering useful shortcuts, let's break these types into functional categories that focus on their common usage rather than brand names and hype.

There are myriad uses for PKI and digital certificates; the following list highlights the most common uses:

- Securing communications with encryption to provide privacy and integrity
 - Secure Sockets Layer (SSL) encryption for externally facing Web applications— You want any personally identifiable information kept private, so you will need to encrypt the traffic between the Web server and the end user computer with asymmetrical keys. In this application, the focus is on providing assurance of privacy and data integrity to your customers and business partners.
 - SSL encryption for intranets—In this application, the focus is on intra-company use. Typical applications include intranets or Web-based email systems for remote use. The main focus is to keep your data secure from eavesdropping via encryption. This type of application is sometimes called *domain authorization* because the CA needs to verify the certificate requester's right to administer the domain and make the request on behalf of the company. Your users know who you are and already trust the company, so there is little need for stringent validation and authentication of the company itself.
 - Email—Most email traverses the Internet much like a postcard. The message payload is typically not protected and potentially could be viewed by malicious eyes. By using PKI, you can safeguard the contents of your sensitive email by encrypting it with public/private keys and providing a high level of assurance that only the intended party can read the message in clear text.

- Providing authentication and integrity
 - Network and application authentication—Digital certificates can be used to prove identity, save users from having to memorize multiple passwords, and keep remote-login procedures simple. Certificates can provide critical assurance that the remote user is indeed who they say they are and reduce the need for multiple forms of credentials to just a few simple methods such as smart cards or secure biometric devices. Certificates can also be used as part of authentication for a virtual private network. VPNs allow use of the public Internet as a private network by encrypting all traffic that passes between the end user and the corporate network. Many VPN applications allow the user to function as if he or she were on the local corporate network.
 - Web site or company authentication—This application focuses on providing assurance to your clients that when they pull up your Web site, they are dealing with your legitimate site and not an imposter. The client's Web browser automatically verifies that the certificate presented by the Web server is valid and is the correct certificate for the server in question. If this information is correct and the certificate is valid, the lock symbol will appear on the bottom of the Web browser indicating the site is legitimate.
 - Software signing—For developers of commercial software products, signing the files that make up a software package with their digital certificate enables you, the end user, to be assured that the program you think you are installing is exactly that and nothing more. By leveraging the authentication and integrity functions of PKI, you can be assured that the software comes from the correct source because the CA validated and verified the company, and is the correct set of files and only the files the publisher intended for you to have (and therefore contains no viruses or malware added by others) because a digital checksum or hash is used.
- Non-Repudiation
 - Transactions—As Chapter 1 discussed, PKI provides merchants with the ability to prove that a particular user engaged in a transaction in the merchant's store. If customers could select items and have them shipped, then later claim that the transaction never occurred, business on the Web would be problematic. PKI allows the merchant to prove that a certain user transmitted a certain transaction. Additionally, the user can be assured that the transaction took place with the merchant that the user intended to do business with.
 - Software—Ensures that files signed by the private key of a software vendor did in fact come from that vendor.

Most of the common PKI applications hinge on effective vetting by a third party, as the issuance of faulty or fraudulent certificates could have disastrous effects. If successful, the malicious party, for example, could impersonate a reputable vendor on the Web or even have signed viruses, Trojan-horse software, or other malware as having originated from a trusted software company. Many users will accept software updates or patches if they appear to be from legitimate sources.

 For more information about fraudulent digital certificates, see <http://www.computerworld.com/softwaretopics/software/story/0,10801,58857,00.html>.

To further illustrate the controls used during the validation of a certificate issuance request, let's look at two very popular certificate applications and contrast the information that maybe required by a CA to complete them.

Contrasting Validation and Verification Procedures

Many companies have leveraged intranets to act as a valuable information and resource repository for their employees. The intranet becomes such a part of daily operations that when users travel, they depend on remote access to the intranet to stay productive. A common approach is to make the intranet available to employees by creating VPNs through the public Internet. Certificates can help to alleviate many security concerns that arise when considering the prospect of making this information globally accessible to remote users. Namely, PKI can provide a credible means of proving that users attempting to access company resources are indeed employees of the company and are authorized to view or interact with the systems made available through the intranet. In addition, by offering SSL encryption, concerns of confidentiality and data integrity can be addressed.

Many CAs offer a commercial certificate product for intranet use. These products often feature a streamlined application process focused on quick turnaround. As a result of the reduced requirement for company authentication, as employees already trust their own company, the process instead focuses on command and control of the domain in question. The CA wants to be able to verify that the certificate requester is named by the company in a public record as having control over the domain for which the certificate is requested. This verification can often be performed with an automated search of a domain registrar's database.

It is important to note that if your organization has opted to use the private or anonymous registration features offered by some domain registrars, you will not be able to be automatically verified. This information must be publicly visible to be automatically checked. In addition, check to ensure that your CA differentiates intranet certificates that are visible via the intranet from internal-only intranets. Doing so will allow the CA to add another important control in the enrollment process. The additional check should be to verify that the host name or IP address of the host listed in the CSR designated for internal intranet use is not visible or publicly accessible from the Internet. A best practice of a CA is to check the IP address associated with the internal intranet host to make sure it falls within the specifications outline by the Internet Engineering Task Force (IETF) RFC 1918.


What Is the RFC 1918 Address Space?


The IETF is a standards body that creates proposed Internet standards through several working groups. The PKI standards were created by the PKIX working group at the IETF. RFC 1918 was created to allocate a range of IP addresses for internal organizational use. The addresses cannot be reached via the Web, as the routers on the Internet backbone know that they are not publicly routable. These addresses were reserved to encourage companies to conserve public IP address space by using a private IP address scheme behind their respective firewalls. The network allocations as outlined by the RFC 1918 include:


[10.0.0.0](#) - 10.255.255.255 (10/8 prefix)

[172.16.0.0](#) - 172.31.255.255 (172.16/12 prefix)

[192.168.0.0](#) - 192.168.255.255 (192.168/16 prefix)


 The complete text of RFC 1918 is available at <http://rfc.net/rfc1918.html>.

 The IETF can be found at <http://www.ietf.org/>.






 The PKIX working group at the IETF can be found at <http://www.ietf.org/html.charters/pkix-charter.html>

To request a commercial intranet (domain) authenticated certificate, you will need to create an email address or an email alias, which is an address that forwards all mail it receives to an existing email address that will act as an authorizing contact for certificate matters at your organization. Most CAs have established a predetermined list of authorizing contact email addresses that they want you to use. All future correspondence regarding your domain will need to use this address. If you are unclear as to what email address or alias to create, search the CA's site for a step-by-step registration or enrollment tutorial. If you can't find a document explaining the requirements for an authorized contact email address, you will need to call the CA to have your questions answered directly. Many CAs now offer online chat operators to answer these types of questions. Examples of predetermined email aliases include [webmaster@\(your-domain-name-here\).com](mailto:webmaster@(your-domain-name-here).com), [hostmaster@\(your-domain-name-here\).com](mailto:hostmaster@(your-domain-name-here).com), and [info@\(your-domain-name-here\).com](mailto:info@(your-domain-name-here).com). The key for many CAs is that the predetermined email address match an email address of a contact specified on the publicly visible domain registration.

You will need to generate a public and private key pair on the host on which you want to install your new certificate. To do so, it will be important to know which Web server software you are using. At the time of writing, Microsoft Internet Information Server (IIS) and the Apache Web server are the two most common Web servers in terms of actual visible installations on the Web. Microsoft's TechNet offers step-by-step instructions on how to generate key pairs and request certificates from commercial CAs as well as tutorials covering self-signed certificates for internal usage. The Apache project also has excellent online documentation for its HTTPD server that answers many of the common questions associated with setting up SSL and generation of key pairs.

 Both the public and private keys used to enable SSL functionality on your Web site will need to be created by your Web server software. The Web server software will also create a Certificate Signing Request (CSR), which you will present to the CA for them to sign with their own private key.

The private key generated by this process should be backed up immediately through a secure storage method such as a HSM or a protected disk or secure tape backup. Remember that if this key is ever compromised, so is the assurance associated with the site it protects!

-  More information about Microsoft IIS can be found at <http://www.microsoft.com/WindowsServer2003/iis/default.aspx>.
-  More information about the Apache Web server project can be found at <http://httpd.apache.org/>.
-  A report by Netcraft on the most common Web servers on the Web can be found at <http://survey.netcraft.com/Reports/0603/>.
-  Microsoft IIS TechNet articles are available at <http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/559bb9d5-0515-4397-83e0-c403c5ed86fe.mspx>.
-  Information about setting Apache HTTPD Version 2.2 servers for SSL can be found at http://httpd.apache.org/docs/2.2/ssl/ssl_faq.html#keyscerts.

Now you will need to generate a CSR on the Web server on which you will be installing the certificate. This process is also performed by your Web server software (see Figure 3.1). The CSR process will use the key pair generated in the last step. If you have questions about exactly how to configure your server to generate the CSR, check the references listed in the previous resource box for help with the Apache HTTPD and Microsoft IIS Web server software.

```

root@atlantis:~
[root@atlantis root]# openssl req -new -key www.mydomain.com.key -out www.mydomain.com.csr
Using configuration from /usr/share/ssl/openssl.cnf
Enter PEM pass phrase:
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [GB]:US
State or Province Name (full name) [Berkshire]:New York
Locality Name (eg, city) [Newbury]:New York
Organization Name (eg, company) [My Company Ltd]:My organization
Organizational Unit Name (eg, section) []:IS
Common Name (eg, your name or your server's hostname) []:www.mydomain.com
Email Address []:

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
[root@atlantis root]#

```

Figure 3.1: The CSR generation process on a host running the Apache HTTPD Web server.

As an alternative to wading through the potentially large amounts of technical documentation that came with your Web server software, you can look for help in the certificate enrollment guide from your CA. These guides tend to be targeted both at the specific software packages and the activities associated with the issuance process. If after consulting all these resources, you still need help, try pulling up your favorite search engine in your Web browser and entering the phrase “how to generate a CSR.” You will want to look for results from the domain name of your preferred CA first, as any CA specific instructions would be covered in their own documentation.

In most cases, the output of the CSR generation process will produce a text file of which the contents will look something like the example in the following sidebar (your characters will be different as the result of random key generation and your unique host information gathered during the CSR process).

What Does the Output of a CSR Look Like?

When your Web server software generates a CSR, it will generate the private/public key pairs that will be signed by the CA you have selected. To present this information to your CA, the Web server software will gather information such as the organization name and several other details that it will combine with the public key; it will then encrypt it all with the newly generated private key. This process produces an encrypted text blob that can be pasted directly into the CA's enrollment form. The CA then uses the public key in the blob to decrypt the contents of the CSR, and in doing so, verifies that you are in possession of the private key, as you used it to encrypt the whole CSR. The following example shows CSR output:

```
-----BEGIN CERTIFICATE REQUEST-----
MIIBzDCCATUCAQAwYsxDHDAaBgNVBAMTE3d3dy50aGlzaXNhdGVzdC5jb20xCzAJ
BgNVBAYTAipBMRkwFwYDVQQIExBXZXR0ZXJlIFByb3ZpbmNIMRlwEAYDVQQHEWID
YXBIIFRvd24xEjAQBgNVBAoTCVRlc3QgQ29ycDEbMBkGA1UECxMSVGVzZdGluZyBE
ZXBhcnRtZW50MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDvzfmv7vJ9bOyQ
dxMLIgtDIEFz7MWsOUoZOPTq3qsTTXPW61q01jY8eQfs96I5xPjxALPeT4m74cce
UtYxldG7pLJiB3SGU94yvyvHDiyV+6mV/e++KWT2ql0Jv1emmobmAGdUxdx2pW9C
Epr0DmcVny6VGWAI36bG0NdYrNix4QIDAQABoAAwDQYJKoZIhvcNAQEEBQADgYEA
BfSHgDr9Vc460YG+IAiWuVWEife8B4QOojiV8oUxJJdqBA2CEEmXLWfa7/mfUtd5
EQd6voLDT8axpXPbOrmwa3kzEZvQZhg+QvKEylfncqdWbDUK71tO0fVafBKwRQfE
73J/THmVABZuz9T6X3+KWGxGDiYw0sY3bE7OjBCwr14=
-----END CERTIFICATE REQUEST-----
```

When unencrypted by the CA, this example CSR contains the following information:

Server domain name:	www.thisisatest.com
Country code:	ZA
State / province:	Western Province
Town / city:	Cape Town
Organization:	Test Corp
Organizational unit:	Testing Department

The CSR process uses the private key to encrypt and sign the information collected during the CSR generation process, thus binding the key pair to the host. This request is what the CA will sign with its private key and incorporate into the certificate the CA issues to you following successful verification and validation of your credentials. This process is very similar to the process used by organizations that issue credentials such as a state driver's license and a passport. In this scenario, you must complete an application and supply credentials to prove your identity. Once the agency has been satisfied, it takes the information you supplied on the registration form and incorporates it into an official license and signs it with the credentials of the particular agency to, in essence, vouch for the information you have provided.

A very important and often-overlooked part of the CSR process is the designation of what is referred to in PKI as a common name. During the process of actually generating the CSR, you will be asked to fill in several information fields as (see Figure 3.2):

- Organization (O)
- Organizational Unit (OU)
- Country (C)
- State (S) (do not abbreviate)
- Locality (L)
- Common Name (CN)

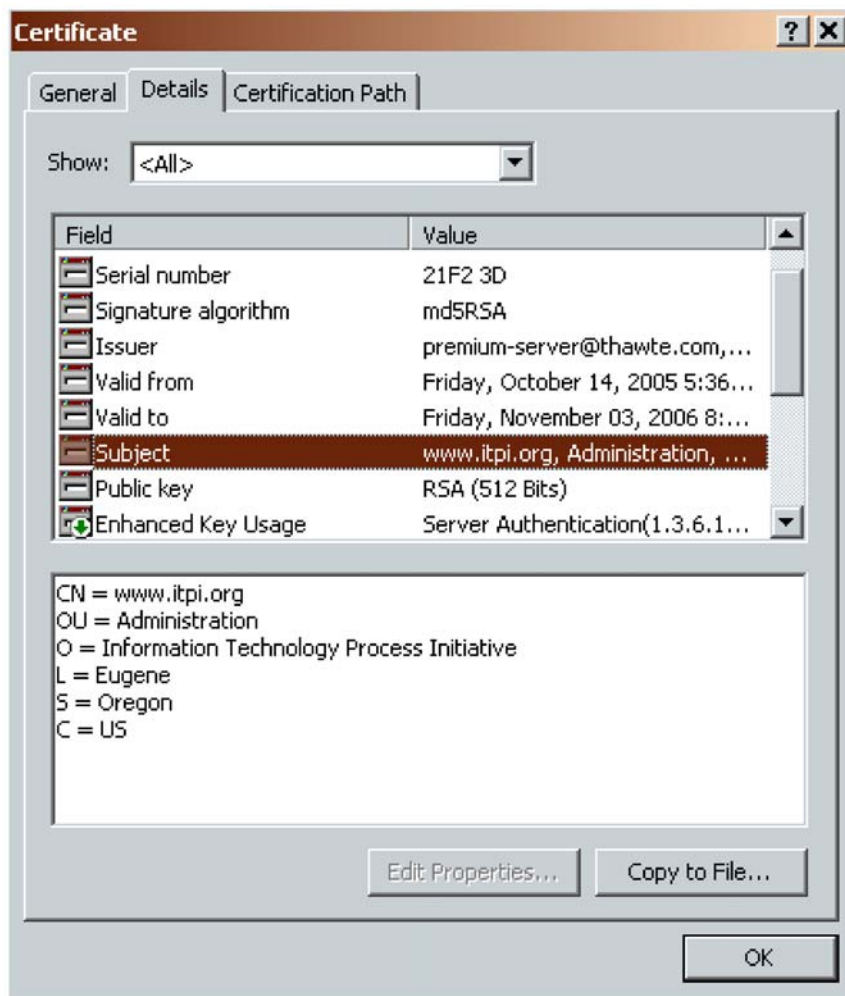


Figure 3.2: When on a site using SSL, double-click the lock symbol at the bottom of your Web browser and the certificate properties window will appear. By selecting the Details tab, you will see a list of fields present in the certificate. Select the subject field, which stores the value you entered during the process of generating your CSR.

The Common Name (CN) specification can be a bit tricky if you are new to generating a CSR. Typically, the common name is comprised of the host name that the CSR was generated on appended to the domain name `yourhostname.(your-domain-name-here).com` (see Figure 3.3). If the name you enter as the common name is just your domain name, your site visitors will likely receive a pop-up window alerting them that the actual secure site hostname you are logged onto is different than the name specified on the digital certificate. This message is hardly comforting, and users will often shy away from using a site that causes their browser to display warning messages. When specifying a common name for an internal-only facing intranet site, simply specify the host name by itself as one word. When it is external facing, you will need to use the hostname plus domain name format mentioned earlier. This format is often called a Fully Qualified Domain Name (FQDN), as it contains all the information an application would need to reach this host from the Internet (providing your DNS is set up correctly).




Figure 3.3: The Microsoft IIS Certificate Wizard prompts for the entry of a common name. Notice that the domain name `mydomain.com` is combined with a host name `www`.

Another piece of information that you will be prompted for is the location or *Geographic location* as it is sometimes called. This parameter refers to the state in which the organization resides. It is very important to spell out the state name completely rather than abbreviate it (see Figure 3.4). If you abbreviate the name of the state, your resulting certificate may not work.




Figure 3.4: The Microsoft IIS Certificate Wizard prompting for location information during the generation of a CSR. Notice that the State New York is not abbreviated but spelled out completely. If you abbreviate your state name the resulting certificate may not function correctly.

Next, find the certificate enrollment or registration page for your chosen CA. Before you actually begin the enrollment process, check to determine whether the CA has an online CSR test page. A CSR test page will allow you to paste the contents of your CSR into a text box. You can then click a submit button and the page will display the contents of your CSR in readable text so that you can verify that the CSR contains all the correct information before you submit it.

Once you have generated your CSR and feel comfortable that it contains accurate information about your host and company, complete the online certificate application by completing the relevant contact information, designating an authorizing contact, selecting the certificate product you want, specifying the type of Web server software you are using on the host, and providing the number of host licenses you will need. Unless you are using a load balancer that is spreading server traffic across several Web servers, you will need only one host license. The final step in the enrollment process is to provide payment, which is usually done with a credit card.

Once the certificate has been processed by the CA, it will typically notify the authorizing contact via email that they can download the certificate. Typically, this download is done by providing a URL and a PIN code to be used as a login credential. Once the authorizing contact has logged in to the CA's site, the contact will be able to download the signed certificate.


Once the certificate is in the possession of the authorizing contact, they will need to follow the steps outlined in the documentation for the particular Web server software being used to install the certificate for the site. This point is a good time to verify that the methods used to back up and protect the corresponding private key have been effective.

Higher-Assurance Certificates

For domain authenticated certificates, most CAs verify that the person generating the CSR is authorized by the company listed as the owner of the domain to administer the domain and request certificates on behalf of the domain. For company-authenticated certificates, the CA needs to authenticate the right to administer the domain and the right to request and manage certificates on its behalf—but it needs to go much further. Make sure that your CA validates at least two and preferably three additional points.

For a high-level assurance certificate used for e-commerce transactions or code signing, read your CA's Certification Practice Statement (CPS) to ensure your CA checks the following points:

- Company name validation—Does the company actually exist and is it licensed to do business? Sometimes the CA will request copies of the articles of incorporation or other legal documents that vouch for the legitimacy of the company.
- Does the company have the right to use the domain name?
- Is the telephone number on the application linkable to a published company phone number? Some CAs will call a published number or send a letter through the postal system to verify that the certificate requester is a legitimate employee at the company and that they are authorized to request the certificate on behalf of the company.
- The CA must prove that the subscriber has a legitimate copy of the private key associated with the public key included in the CSR. In most cases, the CA will spell out this method in its CPS. Often, for security reasons, the description of exactly how a CA does this may be a bit vague and open-ended.
- Make sure that your CA splits the authentication and verification duties in their company by examining the CPS section that spells out separation of duties policy. A good application of this process is to have one CA staff member authenticate the domain and the right to control the domain by the authorizing contacts listed on the certificate application. This task can be performed by pulling up the domain information using a service called Whois at the Internic (see Figure 3.5). Verification would then be handled by another CA staff member and should entail calling the organization listed in the certificate application and verifying that the authorizing contact is both employed by the organization listed and that the listed contact has the right to request and manage certificates on behalf of the organization. If both the authentication and verification processes determine that everything checks out, the certificate application will be approved as long as the CSR is valid.

 The Internic Whois domain query tool can be found at <http://www.internic.com/whois.html>.





Figure 3.5: The InterNIC Whois domain query tool.

Managing Multiple Certificates

Many CAs offer solutions for the organization that needs to manage multiple certificates. Even relatively small organizations that start out with a single Web server and just need a single certificate find themselves in the position of having multiple servers and thus the need to manage multiple certificates in short order. If your venture is successful, it makes sense that you will need to scale your infrastructure.


Many Internet Service Providers (ISPs) and Managed Service Providers (MSPs) manage large server and network infrastructures that require management of digital certificates. In these types of arrangements, service providers are looking to add value to their existing lines of business by assuming the management of the certificate lifecycle for their client base.

 The next chapter will take a close look at the decision process around outsourcing PKI lifecycle management and whether managing certificates in-house makes sense for your organization.

If you find that you need to manage multiple certificates, most commercial CAs offer a multi-certificate management service. This service can be useful for many reasons:

- If you want a consolidated view of all certificates in your possession—complete with expiry and host information.
- If you want to receive the best volume pricing on all your certificate purchases.
- If you want the ability to request certificates on your clients' behalf.
- If you want expedited turnaround time from the CA on all certificate requests.
- If your clients want you to manage everything for them including the CSR and enrollment process.

A key benefit with these programs is that they require validation and verification up front, which once completed, permits authorized account users to request certificates with significantly reduced issuance times based on the initial authentication. It is also usually possible to add domains to your account at any stage, with the CA performing the necessary authentication procedures in each instance. Be sure to carefully read the subscriber agreement for these types of product offerings. It is a good idea to develop a permissions hierarchy for all staff involved in the process that dictates exactly whom will be responsible for managing requests and approvals for certificates. In some cases, extra staff might be required to create a workable separation-of-duty schema. Processes such as change and configuration management will need to be reviewed to make sure they adequately address the increased risk associated with handling sensitive customer information such as private keys. A look at your organization's privacy policy may be in order due to brokering contact- and security-related information on behalf of your clients—not to mention the special infrastructure you may need for the backup and storage of private keys for your clients. Physical security may also be a concern for this additional infrastructure.

 It is a best practice to have any agreements that legally bind you or your organization analyzed by your legal counsel to make sure you are not biting off more than you are willing to chew contractually.

Re-Issuance

Once you have secured and installed your certificate, there are several scenarios that might occur that cause the need for a re-issuance of the certificate. One of the most common is the loss of the private key. Make sure that your CA will re-authenticate your request from scratch in the event of a private key loss. Otherwise, it will be much easier for people with malicious intent to hijack your certificates. The extra steps involved in re-authentication and verification are trivial compared with the damage that can come from an unauthorized party controlling your private key.

Other situations that might necessitate a re-issuance are changes to the version of your Web-software, contact information, or the hostname of the server housing the certificate. Any changes required to the information in the CSR will most likely trigger re-authentication and verification processes as well. It is important to plan for the time that these processes may consume when considering changes to any of these parameters. The CA will not operate on your timetable for verification. Most likely, the re-issuance process will go smoothly, but make sure you have carefully planned for these types of changes and have allowed for potential time delays if your request will require re-authentication and verification. When considering a CA, be sure that you understand what triggers the need for re-authentication. It pays to inquire up front and have a full understanding of your CAs practices. This information is available in the CPS published by all commercial CAs. If it is difficult for you to understand the rules and practices your CA has outlined in the CPS, by all means, engage a sales support person directly and get answers to your specific questions. If after talking directly with the CA you still have confusion, you might want to look at the policies of another CA to determine if they are more inline with your expectations.

Expiry

With all the work required to properly authenticate and verify digital certificates, some people may ask “Why even have an expiration date?” Certificates expire in order to allow companies to manage risk over time. By forcing periodic re-authentication and verification, the PKI is continually cleansing itself of outdated or outmoded information that could be used to increase risk of key theft or key revocation due to inaccurate information in the certificate. The same logic used for state- or country-issued credentials such as drivers licenses and passports apply. By periodically forcing all certificate subjects to re-prove their identity and having that information verified, both the CA and the certificate are subject to reduced risk of fraud. It is very important to understand your CA’s policy with regard to expiry in order to minimize the risk associated with unexpected termination of your certificate and to understand what the timeline of a potential certificate renewal may look like.

Most CAs offer both a 1-year and 2-year certificate product. The 2-year pricing is usually more attractive and should be considered. The main reason is less administrative overhead on your team’s part. When you factor in that the 2-year option is cheaper and requires less work, it is usually the best route to take.

Renewal

When submitting your first certificate request, renewal is a distant thought. But to prevent costly outages and prevent your clients from getting scary messages about your certificate being expired, schedule your renewal the instant you complete the issuance process.

I wish I had a dollar for every horror story I have heard from systems administrators about unknowingly letting domains and certificates expire. Even with the auto-notification emails used by most CAs and registrars, domains and certificate renewal often get put off to the last minute. It is both embarrassing and costly in many cases to let a certificate expire. When your users pull up your secure Web site and are about to pay for their goods, they will see an ugly message from their Web browser telling them that your certificate has a problem and that the date is no longer valid. The user will be presented with the option of choosing to trust the certificate in spite of its date. Many users will not choose to trust an outdated certificate. This situation sends out the message that the site is not well managed—not to mention that it may have been compromised!

Many IT organizations depend on a change management process to protect their critical services from interruption and security breaches from risky unauthorized changes. Typically, the change management process handles the approval and scheduling of IT infrastructure changes. This scheduling is often committed to a master calendar, which serves many purposes. Primarily, it acts as a control to deter change-related collisions and to notify users of changes that may affect them.

One effective way of making sure that certificate renewal does not get lost in the operational shuffle is to put in a change request for the certificate renewal and schedule the work to be done 1 or 2 months prior to the certificate expiry date. This method eliminates the need to put the reminder for renewal in an individual's schedule or task list.

Another effective method is to leverage a Configuration Management Database (CMDB) as spelled out by the standard for IT service management, the IT Infrastructure Library (ITIL). The CMDB is designed to contain all relevant information about the IT infrastructure and its relationships and dependencies. This database is not limited to hardware and software; it also may contain documentation, policy and procedure information, and depictions of network and security architecture.

 The ITIL is a registered trademark of the United Kingdom Office of Government and Commerce. More information can be found at <http://www.itilpeople.com/>.

In my tenure as CTO and CIO at several organizations, I have used a CMDB to keep track of important information about domain and certificate expiry. A simple report run on a daily basis can illuminate which domains and certificates need to be renewed long before a crisis is looming.

Integration of renewal with change management and configuration management processes can also alert authorizing contacts to changes that might need to be made to the certificate itself. It is very important to understand what types of changes your CA will allow to your renewal. A good CA will have firm boundaries around exactly what can and cannot be changed under the renewal process. For example, you should not be able to change any details contained in the CSR. Details such as contact information, the email address of the authorizing contact, or the business contact can be changed without the need for re-issuance. These changes may trigger validation and verification efforts to provide assurance that your certificate is not being hijacked by an unauthorized party.

If any of the details in the CSR must change, re-issuance rather than renewal will be required by your CA. Most likely, your CA will then need to re-authenticate and re-validate your information. This can add time delay to the process—especially if there is inaccurate information in your original certificate request.

Renewal is also a perfect time to consider the quality of service your CA has provided you during your subscription period. Have you been satisfied with their performance? Were they available to answer your questions in a reasonable timeframe? In today's global economy, you might find that you are dealing with a company an ocean away. Do they have support hours that match your region or time zone and native language? It is important to look at their offering in light of new products and pricing from other vendors. Are their prices and service quality in line? If not, consider switching to another CA. I highly recommend checking for changes and updates to both the subscriber agreement and CPS to determine whether you have gained or lost any rights or assurances that might be important to you.

Revocation

Why do you need revocation? If your private key is somehow compromised, you would want it revoked. If your organization is found to be engaged in illegal activities, your key will be revoked for you. Despite all the best practices and best efforts of your staff and the fully vetted, background-checked dedicated team at your CA, mistakes can happen. When they do, it is good to know that there is a way to stop a bad certificate from deceiving non-suspecting relying parties. The revocation concept is not new. Most state and federal law enforcement agencies have databases of credentials such as state issued driver's licenses. They can run the numbers against the database to make sure the license hasn't been altered and is in good standing. Credit card companies use the same principle to prevent stolen or missing credit cards from being used once their disappearance has been reported.

In PKI, these lists are called Certificate Revocation Lists. A CRL is a signed, time-stamped blacklist of revoked but unexpired certificates and is issued by a CA periodically (usually daily). CRLs are a commonly recognized standard (as spelled out in the X.509 RFC) and are deemed by many organizations as acceptable for use in non-online or low-value commercial applications. The largest issue for security-sensitive organizations is that the data might not be current enough if the CA publishes the list only once a day. The other main issue with CRLs is that they can become impressively large and unwieldy over time. To address the size issue, the Delta CRL mechanism, which only contains the changes since the last CRL was issued, was developed.

The CRL Distribution Point (CDP) mechanism was established to tackle the issue of large list sizes by partitioning the CRL into relevant PKI communities. CDP partitions have their own designators or pointers that are embedded in the certificate so as to point the query to the correct CRL partition.

In the financial services world in which large transactions such as funds transfers require an online check of the certificate's status as of that exact moment, the Online Certificate Status Protocol (OCSP) was developed. OSCP has since become the successor to CRL; although many Web browsers still use CRL, it is commonly disabled by default. As new releases of Web browsers become available with OSCP enabled, it will completely replace CRL. OCSP was enabled to allow for nearly instantaneous checking of a certificate's real-time status with a digitally stamped and signed result making it suitable for high-value/high-assurance transactional scenarios. One downside to OCSP is its performance. Due to the fact the requests are happening in real time and the results must be stamped and signed by the CA, the performance can fluctuate greatly and can be problematic for Online Analytical Processing environments or high-volume commerce. When choosing a PKI provider, ask questions about their investment in the OSCP infrastructure. Have they built their plans around OCSP becoming the standard or have they built their infrastructure to support the older standard and have adopted a wait-and-see approach with OSCP? Like CRL and CDP, OCSP was also spelled out by the PKIX working group at the IETF and is part of the X.509 standard.

Summary

This chapter has covered a lot of ground, starting with issuance and working to revocation. Along the way, it covered best practices that will speed you along the certificate application process and help you obtain a certificate that will function seamlessly for your end users. Answering the question "Just how do you generate a CSR?" allowed for a look at the CSR and what it contains. The chapter revealed that you as the certificate subscriber are responsible for generating the public/private key pair that the CA will ultimately verify, validate, and sign with the CA's private key.

Tools such as online multi-certificate management programs from CAs were presented as a potentially effective solution for managing certificates for multiple clients or for a single-user multi-certificate environment. This chapter made sure to tip you as to what changes could be made in the re-issue and renewal processes that wouldn't require you to be re-validated and re-authenticated to avoid introducing additional delays in the process.

The next chapter will look at whether it makes sense to outsource PKI functions to PKI service providers. It will also examine the enterprise CA model to determine whether there is any benefit to becoming your own CA and issuing your own certificates for your employees and internal applications. With all the flexibility that being your own CA offers, there is a major tradeoff in complexity. But not to fear, the chapter will cover the processes controls and key practices you will need to know to make an informed decision.

Chapter 4: Managing PKI Infrastructures

The previous chapters have discussed the components used in a Public Key Infrastructure (PKI). This chapter will address the key decisions that IT professionals need to make in planning and implementing the infrastructure that will best support their security needs—while addressing the demands of their budgets. The chapter will focus on

- Defining the requirements and challenges of the infrastructure
- Identifying the key components of the infrastructure
- Choosing the best means to provide the infrastructure—specifically, comparing outsourcing with building an in-house system

The Requirements for an Effective PKI

Why does an organization need a PKI? Although this topic has been examined in previous chapters, a brief review will help set the stage for the decisions you need to make in designing and building your own infrastructure.

PKIs help people secure information. Specifically, PKIs enable the exchange of information between entities in electronic form. There are a variety of specific functions for which PKIs are used.

One use is to encode a data stream. This could be all data exchanged on a TCP channel, such as used by the secure hypertext transfer protocol (HTTPS) or secure File Transfer Protocol (SFTP). It might be encoding the text within the message, transmitting other portions of the message in plain text so that other components, such as mail servers, can effectively deliver the encrypted payload without being aware of the encryption that protects the sensitive data. By encrypting the data, it can be sent securely through public channels such as the Internet. This enables electronic commerce and exchange of confidential information to be extended to nearly everyone, displacing the expense and added maintenance of private networks and Value-Added Networks.

Encryption keys can be used to secure data stored in databases and on file systems. These keys protect long-term assets and should be stored securely. If the keys are lost, the encrypted data becomes unusable. If the keys are compromised, the data is no longer secure. A PKI can implement technology and procedures that secure the keys that protect sensitive organizational data.

Another use is to confirm the identity of the sender. This is the process of using digital signatures. The data itself may not contain data that is sensitive. If electronic networks are to be used to conduct legal commerce, some form of non-repudiation of the parties involved in a transaction is required. If the parties met, they would sign paper documents, and their signatures would stand as evidence of their willingness to participation in the transaction. Digital certificates can be used as legally binding digital signatures.

More frequently, the need of identity confirmation is much more mundane. For instance, if remote access to the corporate network is required, the virtual private network (VPN) server might require a digital signature to validate the user. A smart card can contain such a signature and improve the security of VPN access.

PKI Requirements

These needs help define four key requirements for a PKI:

- The PKI must be able to deliver a public key that the client can use to encrypt data. The system must be able to validate the authenticity of this key.
- The system must be able to authenticate the validity of a key, so authorized administrators must be able to revoke a key. Keys need to be revoked for a variety of reasons: the private key is compromised, the issuing entity is retired or no longer valid, or a new encryption algorithm or key size is implemented. Any of these reasons may require an existing certificate to be revoked.



For the purposes of this discussion, a certificate is a digital document that contains a public key and metadata concerning the issuer and the Certification Authority (CA) that validates it (for instance, an X.509 certificate).

- If the system is to count on the certificate to provide proof of identity, it must provide non-repudiation. If a digital certificate is provided to a consumer, the PKI must reasonably guarantee that it was issued by the certificate holder. Thus, similar to a written signature on a paper document, it can serve as evidence that the communication was sent from the certificate's owner.
- For the PKI to be utilized by automated systems, there must be a centralized means of defining and enforcing policy within the system. For instance, defining encryption algorithms, enforcing key lengths, revoking keys issued by terminated employees, expiring keys after a defined period of time, and so on.

The Threats to Your PKI

To implement this type of system, there are two key attacks that must be defended against. First, the private key must be jealously guarded. With the private key, any encryption provided by the system is laid bare. Further, if the key is lost, all data encrypted using the public version of the key is rendered virtually useless. Thus, the key must be backed up securely.

The second threat comes from the “man-in-the-middle” attack. If a malefactor can intercept a legitimate certificate and substitute his or her own certificate in its place, the attacker now controls the private key that can decrypt the payload. The attacker can substitute his or herself as the legitimate recipient of the encrypted documents. To protect against this attack, the PKI must provide a mechanism to identify the source of the certificate.

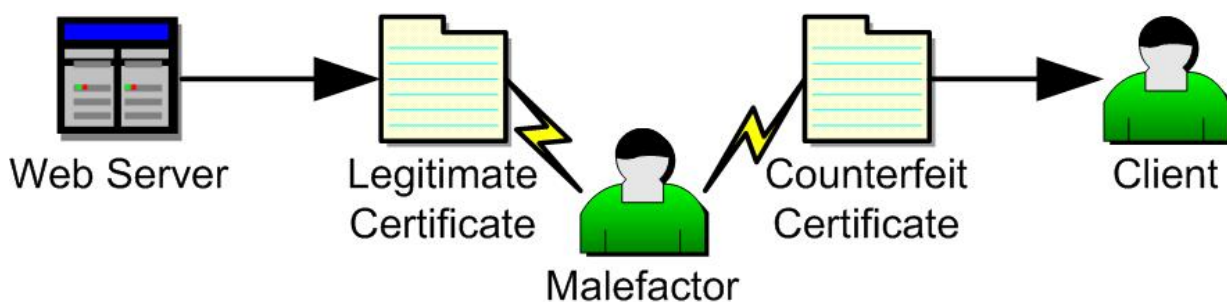


Figure 4.1: A man-in-the-middle attack.

Building a Network of Trust

Central to any PKI is developing a mechanism by which the legitimacy of a key can be confirmed. The simplest means is by direct trust. If a person or entity you know or can independently verify provides you with a public key, you trust that key directly. This could occur when a person hands you a CD with a public key file on it, or your employer provides you with a smart card.

The question becomes, how well does this system of direct trust meet the requirements of your need for a PKI? When you have independently authenticated the source of the key, that requirement is met. Of course, in the anonymity of the Internet, this could be much more difficult to confirm. Certificate revocation can also be a problem. The issuer can ask you to stop using the key, but they are wholly dependent on your compliance. If you issue keys with direct trust, and need to upgrade, changing the encryption algorithm or key length, you must contact each person who has a key and get them to exchange them. Non-repudiation also becomes more of an issue. As the issuer may not have tight control over the distribution system or control of the keys passed to clients, they may not stand by the issuance of the key. Also, policy becomes manual and cumbersome.

With all these disadvantages, one might think this system is seldom used. But, with a limited number of users and the right circumstances, it can be appropriate. For instance, sharing keys between the servers of business partners to share sensitive information is often accomplished through direct trust.

The Role of Certification Authorities

To better automate and manage the system of distributing and managing keys, a more formalized system has been developed. In this system, there is a centralized source—a Certification Authority (CA)—whose primary function is to validate certificates. The CA has its own signature that can be independently validated. When a user creates a certificate with his or her private key, the CA validates the public key and adds the CA's digital signature to the certificate. Thus, the CA stands as the validator of the certificate.




Figure 4.2: CAs validate certificates.

The CA acts as a trusted agency, so its function is typically fulfilled by a third-party organization. This setup promotes trust, particularly between independent organizations with divergent goals. A third-party CA best serves by remaining a neutral source, trusted by all parties.

The test of the CA system then becomes whether they can meet the requirements for your PKI. First, the CA can authenticate the certificates they sign. The CA signature on the certificate can contain a hash that can help determine any alterations to the certificate.

Certificate revocation is also readily handled. The certificate holder can inform the CA that the certificate is no longer in service. This can be an automatic expiration or an overt change or retirement of the certificate. When a user goes to validate the certificate, the user can be notified by the CA of the change in the status.

The certificate should stand as a non-repudiable identification of the certificate holder, similar to a signature. The CA validates the identity of that holder when they sign the certificate. The holder then cannot deny that it is his or her certificate.

The matter of control is also addressed. Because use of the certificate can be validated with the CA and the status can be controlled by the holder, working with the CA, the policies relating to how the certificate is used are more easily managed.

CA Systems

It would be simple if there were but one CA. Everyone would register their keys with that one authority. The CA could guarantee uniqueness of names and easily track how certificates are used and the policy that is applied to them. They would have a common signature and a simplified system of authentication.

In practice, there are many competing CAs. Each one offers differing levels of security and service, as mentioned in the previous chapters. Such is the situation on the World Wide Web. This model is often referred to as the web of trust. In the web of trust, a user can receive certificates from a number of CAs. The user may also directly trust a certificate issued from another user who is known to the user, such as a business colleague. This is really an outgrowth of the direct trust system. It is used particularly within an organization.

The major drawback is that policy is difficult to enforce. If a person has certificates with multiple vendors that do not report with one another, that person must carefully manage the certificates they have with those vendors. Also, much of the burden of identifying the legitimacy of the certificate source falls onto the user.




Figure 4.3: Web of trust.

It can be advantageous for a CA to create subordinate authorities. In this case, a root CA holds the credentials to validate multiple subordinate CAs. Each of the subordinate CAs can issue and validate individual certificates, and the subordinate CA can be validated against the root CA. For instance, an organization might contain subsidiaries. Each subsidiary might need to create and manage its own CA. The root CA for the organization can validate the certificates of the subordinate CAs and stand guard against someone creating a counterfeit CA within the hierarchy.

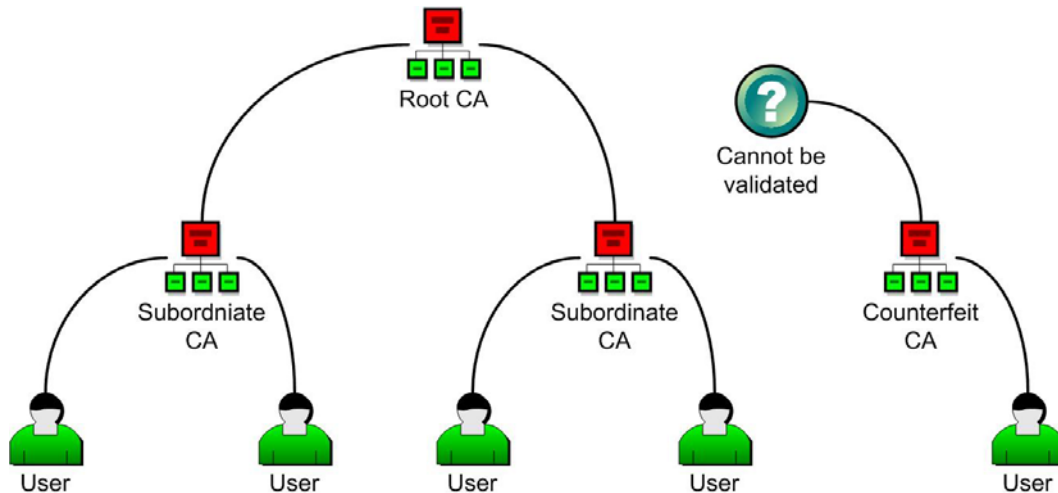


Figure 4.4: CA hierarchy.

Digital Certificates

Digital certificates are used to help standardize and automate the tasks of exchanging these public keys and authenticating signatures. A digital certificate contains a public encryption key and metadata that is used to identify, validate, and control use of the key. For instance, a typical certificate might contain:

- Owner's name
- Name of the CA that validates the certificate
- The public key used by the CA
- Owner's public key
- Period for which the key is valid
- Digital signature that can validate the key

There are several standards for controlling how certificates are constructed and authenticated:

- X.509 is the oldest and most commonly used standard in commercial applications. It provides for direct authentication from a single CA, web of trust, and hierarchical systems of trust. To accommodate these varieties, the certificate itself is extensible and more complex to use.
- OpenPGP (Pretty Good Privacy) was adopted from encryption software originally targeted for encrypting emails. It is lighter weight than the X.509 standard. It relies on direct authentication or the web of trust.
- Simple Public Key Infrastructure (SPKI) was developed to greatly simplify certificates for which only a single CA is required. These are ideal for use in intranet applications where validation of the source needs to be less rigorous. The infrastructure used to implement SPKI is referred to as Simple Distributed Security Infrastructure (SDSI). It is based on the work of Ron Rivest and is being developed by the Internet Engineering Task Force (IETF).

As has been previously mentioned, the X.509 standard is the most common standard for exchanging public keys. This standard is maintained by the IETF. They created a working group for maintaining this standard called PKIX. The charter for this group can be found on the Internet at <http://www.ietf.org/html.charters/pkix-charter.html>.

Elements in a PKI

As PKIs have grown, there are several common services and structures that have grown to meet their needs. Understanding these key elements can help one plan and administrate a PKI.

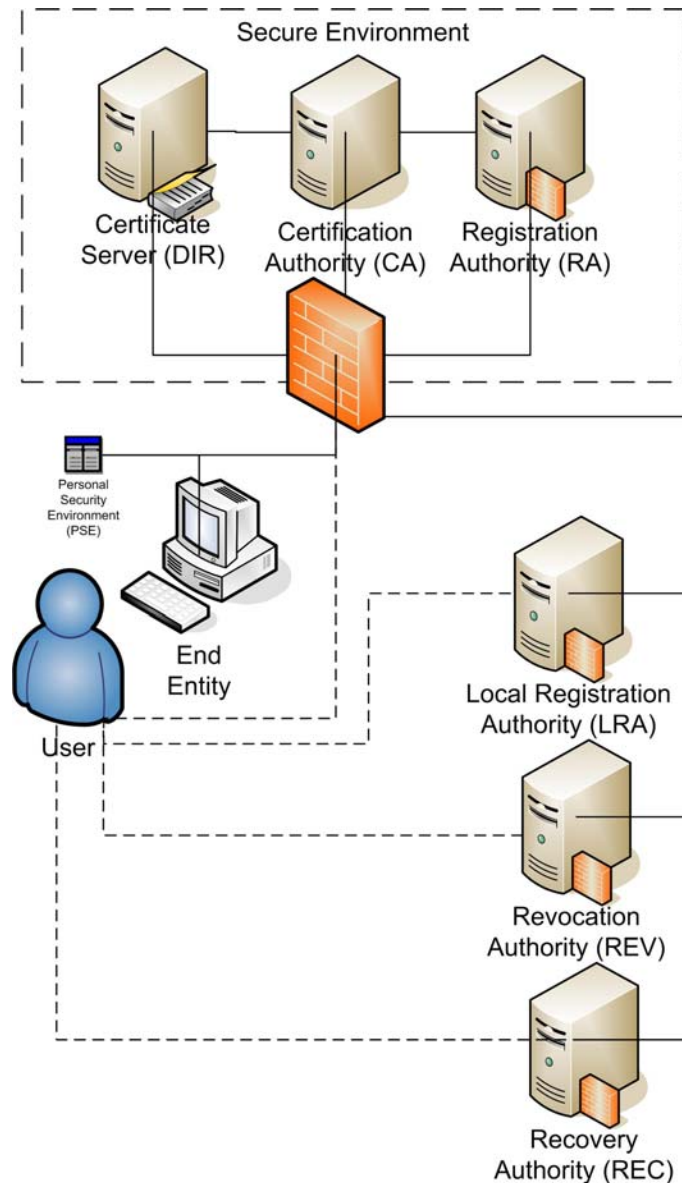


Figure 4.5: Components of a PKI.

CA

The CA component is responsible for generating digital certificates. A computer that wants to receive encrypted data generates a public key and submits that key to the CA computer, along with other required metadata, such as the duration of the certificate and the distinguished name of the owner. The CA computer compiles this information, adds its own public key, and signs the freshly minted digital certificate.

The CA system must be closely guarded. If its private key were lost, it could no longer create or validate new digital certificates. If the private key were compromised, a malefactor could forge certificates and violate the primary trust provided by the CA system.

Because of the need to protect this computer, it is typically kept in a high-security environment. This would include physical security, such as securely locked rooms, as well as regulations that enforce at least two people being present when the computer is directly accessed. The CA computer typically is not connected to the Internet directly.

Registration Authority

The registration authority (RA) allows users to apply for digital certificates. It collects the public key and metadata related to the certificate and typically handles corresponding with the applicant. Most RAs are public facing and require Internet access. Care must be exercised in the communications between the RA and the CA so that the CA security is never compromised. An RA may be a physical server or just an administrative service application.

RAs may be distributed outside the control of the CA. For instance, if an organization creates a PKI for their intranet, they may choose to distribute RAs in each satellite location. These distributed servers are referred to as local RAs.

Certificate Server

When the certificates are created, they need to be stored for future retrieval. This storage must be secure. The certificate server (or Directory—DIR) is the component that stores the certificates in a secure manner for future retrieval. Certificate servers are also used in support of the certificate revocation process.

Time Stamping Service

The time stamping service (TSS) is an optional but significant service. In order to provide non-repudiation, transaction must secure the time at which the certificate was used. Because the certificate holder maintains their policy, including the period of validity of the certificate, within the CA system, having a secure time stamp with the certificate helps to validate that the certificate holder was honoring the certificate at the time of use. If the trusted authority provides the time stamp (typically as a value-added service), it serves to indicate that the certificate was valid when in use.

Revocation Authority

There are times when a key must be revoked. This can be done as part of the RA or can be a service in and of itself. The revocation authority (REV) informs the CA that the certificate is no longer valid. The CA will then add the certificate to the revocation list and make the necessary changes with the DIR service so that the certificate is no longer issued.

Recovery Authority

If the PKI centralizes the secure storage of keys, the recovery authority (REC) provides a mechanism to obtain a copy of the key. Not all PKIs centralize the storage of keys. It is common for this service to be supplied in conjunction with one of the other aforementioned services, such as the RA or REV.

Personal Security Environment

The certificate holder needs a place to store their digital certificates, public keys, and private keys. This is known as the personal security environment (PSE). All the components that handle keys and certificates within the system also require a PSE for their secure information. A PSE can store data on a hardware security module, smart cards, or other security device. Typically, however, it is stored as a file on a computer. It is important to ensure that this information is well protected. As previously stated, loss of a private key invalidates the data encrypted with the accompanying public key. If the private key falls into nefarious hands, data encrypted with the public key is no longer secure and the trust is violated.

End Entities

Finally, the key contained in the certificate must be used to encrypt data. The PKI software applications that consume the digital certificate, find the public key, and use it to establish a secure connection with the server are termed end entities.

Certificate Management

Certificate management is the crucial function served by your PKI. The IETF provides guidelines for managing keys in RFC 2510 Certificate Management Protocols. It structures a number of critical elements that must be addressed by the PKI that bear consideration.

Initialization

Several elements of your PKI must be initialized before the system can be put to work. First, the CA must be initialized. A private key must be generated under controlled circumstances so that the key is secured and not copied or stolen. The process must be monitored by at least two administrators to ensure the integrity of the procedure.

Certificates must be created. Upon initialization of a new CA, the process begins as new requests are made. Over time, certificates will be revoked and replaced. The CA must store the new certificates within the DIR and prepare the publication and revocation lists used to validate the certificates.

The certificates will ultimately need to be exported to end entities so that they can use the public key therein contained to encrypt data. The CMP provides definition of the protocols used to regulate this process.

Publication of Certificates and Revocation Lists

Although the CA creates and manages the certificates, they are stored in the DIR. The CA must provide the DIR with copies of new certificates to store. If a certificate is revoked, the CA adds it to the revocation list, which is also stored by the DIR. The CMP lists several protocols for implementing this process.

Key Recovery

The CA can be used to create the private key as well as the public key used with a certificate. The advantage of this system is centralized management of the key pair. It requires the CA to maintain a secure database where the keys can be kept. If data is stored in an encrypted manner, recovery of the key can be crucial to restoration of the data. But there is a strong caveat to consider: If the database is compromised, all the data secured with the keys stored therein is at risk. The benefits should be carefully considered against the risks. If your organization chooses to centrally manage the key pairs, consider the following:

- The database is effectively an extended PSE. It must be kept in a secure environment and carefully guarded.
- There must be a clear, carefully followed protocol for recovering keys.
- Only keys used to store data that is kept in long-term storage needs to be secured. Keys used to temporarily encrypt emails or to provide digital signatures do not need this level of security.

Revocation

Once the certificates are in circulation, there are a variety of reasons that they may need to be revoked:

- If the private key is lost or compromised, the certificate should no longer be used to encrypt data. Either the data will not be decrypted (loss of private key) or it will not be secure (theft of private key).
- If the private key of a CA is compromised, the malefactor can coin counterfeit certificates. In this unlikely but devastating event, all the certificates managed and validated by the CA must be revoked and replaced with new certificates that are based on a new private key.
- If the metadata on a certificate changes, the hash values on the certificate change and thus the old certificates must be revoked and replaced with updated certificates. For instance, X.509 certificates contain the distinguished name of the server in which they are installed. If the distinguished name of the server is changed, the certificate must be updated. Doing so will result in the invalidation of the old certificates.
- If an organization needs to suspend a certificate for a short period of time—for example, to combat a Denial of Service (DoS) attack—it can be temporarily revoked and later reinstated.
- If a server or entity is retired or is no longer associated with the originating entity, his/her/its certificates are revoked. If an employee has a certificate for VPN access, that certificate is revoked when the employee terminates employment. If a server is replaced through a server consolidation, its certificates are no longer used and should be revoked.
- For security purposes, certificates come with expiration. Once they expire, they are revoked and, typically, a replacement certificate is made available.

Certificate revocation is performed by the CA. The CA maintains a revocation list that is stored in the DIR. If the certificate holder drives the revocation—because private keys have been lost or compromised—the certificate metadata changes or the certificate is to be retired; the holder must inform the CA. This can be done via the REV service.

Certificate Issuance

When an entity requests that a CA issue a certificate, the process is known as *enrollment*. The enrollment process is governed by the guideline established in the CMP. There are a variety of ways to implement enrollment.

An entity can generate a private and public key pair on his/her computer. The private key is secured in the PSE. The public key is packaged in some portable form, such as a CD, floppy disk, or USB key. It is transported in this form to the RA as the entity applies for the certificate. Validation of the entity is part and parcel with the application process and must adhere to the CA's Certification Practice Statement (CPS) for validating the user. The RA collects the rest of the requisite data (name of the certificate holder, expiration date, and so on), and submits the data to the CA. This process is governed by standards, such as PKCS#10 (see RFC 2511 for details). The CA creates a properly formatted certificate. A copy of the certificate is placed in the DIR for safekeeping and a copy is given back to the RA. The RA issues the copy of the certificate back to the requesting entity.

As most people do not venture in person to their CA, the process is often performed online. Similar to the previous process, the applicant fills out an application online and submits their public key (the generation of key pair and the process of submitting the key were detailed in previous chapters) to the RA. The RA uses the process outlined in its CPS to validate the applicant, then submits the key and metadata to the CA. A digital certificate is created, stored in the DIR, and returned through the RA to the applicant.

If the key pairs are centrally managed by the CA, the applicant begins by filling out an application. Once the application is validated, the RA forwards the request to the CA, which generates both the private and public key. The private key is secured and copied to a PSE (often in the form of a smart card) and a digital certificate is created. Both the PSE and digital certificate are returned to the applicant by the RA.



A very common use of certificates is to encrypt data shared by network routers. Cisco Systems, a leading manufacturer of network routing equipment, has developed a protocol for enrolling and changing certificates, dubbed the Certificate Enrollment Protocol. CEP provides for decentralized key generation and online installation. This protocol is not compliant with the CMP protocols established by the IETF.

Certificate Servers

Although the CA creates and manages the certificates, the certificate server stores those certificates and makes them available. Because of the central role that the DIR component plays, it deserves special attention.

The job of the certificate server is to provide access to certificates and revocation lists, so it serves much the same function as a directory server. Directory servers provide information concerning the entities within an organization. They can provide authentication information for security, email address lists, and other information. Because this function is very similar to that required of the DIR service, certificate services are often combined with the corporate directory services.



The PKI does not actually require or specify an organizational directory service. As this service is commonly used, many directory servers include certificate services as part of their functionality.

Directory Services

A directory service provides a database that can be accessed to identify the entities contained within an organization. Each entity gets an entry in the directory service database. The entity will have a set of attributes, such as its type (user, computer, printer, router, and so on), its name, and its network address. Each type of object in the directory will have a defined set of attributes. The definition of these attributes is called a *schema*.

To manage the objects, they are typically organized into a hierarchical structure commonly referred to as a Directory Information Tree. The DIT consists of container objects and leaf objects. Container objects are used to create departments, groups, or other organizational units (OU). The leaf objects represent actual entities, such as users, computers, routers, and so on. To help define objects within the tree and facilitate navigating to those objects, a namespace can be defined. The namespace holds the specification for the common objects that the DIT contains.




Figure 4.6: Example of a DIT.

X.500 Directory Specification

The ITU-T published a specification for directory services called X.500. It is a complex standard designed to cover a very wide range of uses. It has been organized into nine subsets of the standard. X.509, discussed ubiquitously through this guide, is one of those sub-standards. You can discover more about these standards at <http://www.itu.int/home/index.html>.

One of the key concepts is to develop a globally unique name of an object. By structuring the namespace correctly, this goal can be achieved. As one of the goals of the PKI is to identify a certificate holder, providing a unique name for the certificate holder is important. To provide a unique name so that a specific entity can be positively identified, the X.500 specification presents the concept of a DN.

A DN consists of a name that is broken down into components. The X.500 specification identifies five components that can be used to comprise a DN:

- Country (C)—The country in which the name is initiated
- Locality (L)—A localization, such as a city, region, and/or state to which the entity is affiliated
- Organization (O)—The organization within that locality to which the entity is affiliated
- OU—A group used to organize entities within the organization, such as Employees or Servers
- Common Name (CN)—A name for the entity that is unique within the OU in which it resides

Thus, if John Doe is in the Engineering department of MyOrganization located in Anytown, USA, his DN within the directory could be expressed as

C=USA, L=Anytown, O=MyOrganization, OU=Engineering, CN=John Doe



This is a navigational path. John Doe could easily be a member of an additional OU, such as Network Administrators. Thus, although there is only one John Doe, he can have many DNs that correctly identify him.

The X.509 specification has been implemented as several directory protocols. The following list identifies a few of the most common.

- Directory Access Protocol (DAP) is a very full-featured protocol. It allows the DIT to be navigated and objects within the tree to be modified. It can also include security protocols. Like many full-featured protocols, it is difficult to implement and use. DAP servers are often extended to build X.509 certificates and distribute them as required.
- Lightweight Directory Access Protocol (LDAP) was developed by the IETF to provide a directory protocol that was focused on the needs of Internet users. It is described in RFC 2251 through RFC 2267. The protocol is a simplified version of the full-scale DAP. LDAP has been implemented by several vendors to act as a directory service for their TCP/IP-based networks, providing authentication, location, and security services. LDAP servers are among the most popular of directory servers and are very frequently used to distribute certificates and post revocation lists.



When accessing a specific RFC by number, you can visit <http://www.ietf.org/rfc/rfcXXXX.txt> where XXXX is the number of the RFC. Thus, you can access RFC 2251 by browsing to <http://www.ietf.org/rfc/rfc2251.txt>.

- Domain Name System (DNS) is implemented as the standard means of mapping common text-based names to IP addresses. In the public Internet, it is used to ensure the uniqueness of common Internet names through the services of domain registrars. Because DNS is carefully monitored to prevent duplicate names, it becomes a standard source for developing a truly globally unique name space in which entities can be defined. Most TCP/IP networks implement DNS servers to resolve names within the network, so they are a very common network appliance.

The IETF proposed a service extension to DNS that allows it to be used as a certificate server, termed DNSSEC (see RFC 2532 and RFC 2931). The protocol allows a digital certificate to be distributed within a DNS message. This concept was not considered during the formation of the X.509 protocol specifications, so the process itself is not compatible with those protocols. This incompatibility has resulted in sparse adoption of the specification.

- Active Directory (AD) is Microsoft's proprietary directory service. It provides an LDAP-compatible interface and similar functionality to an LDAP service. Microsoft provides an LDAP-compatible interface to its directory to allow LDAP clients to access its information store. The primary limitation is that it requires the use of Microsoft Windows servers to implement and administrate. Microsoft Windows servers include a certificate service as part of their standard offering.
- Netware Directory Service (NDS) is the proprietary directory service developed by Novell. It provides similar functionality to that offered by AD. In conjunction with its implementation on Linux and open source platforms, Novell now offers eDirectory as an LDAP directory implementation.

Why All the Attention to Directory Services?

The X.500 specification is the basis for building digital certificates. The X.509 specification has always included a DN field used to identify the certificate holder. And because the primary functions of a DIR service are quite similar to those of a directory server, many directory server implementations include a DIR service as part of their standard offering.

Certificate Servers and Revocation Lists

One of the key functions of a PKI is management of a list of revoked certificates. The certificate server is an obvious choice to fulfill this function. When a certificate is revoked, the CA notifies the DIR of the revocation and the certificate is added to the revocation list. Each certificate has a serial number that serves as the unique identifier of the certificate. When asked to validate a certificate, the certificate server can quickly search the list. If it is located, a revocation notification can be issued. Accessing the list can be done in a variety of ways.

Checking Revocation Status Online

If the end entity can connect to the certificate server online, it can query the server to discover whether the certificate has been revoked. The certificate server searches the revocation list. It compiles a response, signs it, and returns it to the requestor.




Figure 4.7: Online certificate status checking.

Although this system is conceptually simple, it includes costs that should be carefully considered. The user must have an available connection to the certificate server. As long as the certificate server is local and available, this is likely to be the case. But if the user is occasionally connected, or the DIR is separated through the Internet or a WAN line, this can be costly. Also, the activities performed by the DIR server to look up the specific certificate and then compose and sign a message consume server resources. During times of peak use, this can impact performance or even dictate the addition of servers. The chief advantage is that the status of the certificate can be ascertained with up-to-the-minute results.

The IETF developed the Online Certificate Status Protocol (OCSP) to facilitate online validation (see RFC 2560). When the server receives the request, it ascertains the status of the certificate, the reason for revocation (if any), and compiles and signs a response. The server does not give any information about itself beyond the signature. This led to the development of a more expansive protocol that provided more information.

The revised protocol is dubbed, ironically, the Simple Certificate Validation Protocol. The expanded protocol provides the following features.

- The end entity can submit the request to a server that queries the CA that issued the certificate. It can then follow the hierarchy chain to the root CA to determine whether the CA can be trusted. It can then return the result to the end entity.
- If the end entity does not trust the server, it can request the CA certificates be returned so that the end entity can validate the trustworthiness of the certificate validation.
- The protocol can be used to download a list of revoked certificates.

Using Revocation Lists

As checking, signing, and rechecking the same certificate over and over can create a great deal of overhead, it is common for the CA to create a periodic list of revoked certificates and sign the entire list at once. This Certification Revocation List (CRL) can be downloaded and will reduce the load on the CA server. An end entity can download the list and know which certificates are revoked.

Of course, if the end entity must download the entire list from every CA with which it deals, that would create an unmanageable system. To help control the deluge of information, the list can be partitioned to reduce the amount of data required. At first blush, this may not seem to reduce the amount of data downloading at all. The X.509 specification allows for the identification of a CRL Distribution Point. The CDP can be used to partition the CRL by a common factor, such as the URL of a distribution list. Thus, an end entity need only download the CRLs of entities for which it holds certificates. Once the CRL is downloaded, it can be used to identify revoked certificates from the point in the partition. The end entity can download the list as frequently as required, but does not need to connect to the server each time the certificate is used. Still, this results in a lot of information being downloaded to obtain the results of a single certificate.

The solution lies in creating a certificate revocation tree. The tree is a hash of the values of revoked certificates and their CAs. By combining them so that each possible combination appears in only a small hash, the amount of data can be significantly reduced for a single certificate query. The hash values allow only small amounts of data to be downloaded to provide confirmation of whether a given certificate has been revoked.

The X.509 specification provides detailed protocols for CRLs. The entries in the list are required to contain certain fields:

- The version number of the X.509 specification used to define the CRL
- The Object ID of the algorithm used to sign the CRL
- The X.500 DN of the CRL issuer
- The date the CRL was issued
- The date that the next CRL will be issued
- The serial number of the revoked certificates

When the X.509 specification was expanded to accommodate the X.509v3 specification, it incorporated the definition and use of additional fields. Standard extensions in this version of the CRL specification include:

- Authority key identifiers that provide a unique identifier for the issuer's key
- Issuer additional name fields, such as an IP address or email address
- CRL number is a unique identifier for this CRL

There are also additional fields created for each certificate

- Reason code for the revocation of the certificate
- Certificate issuer name
- Reason for a suspension—used if the certificate is suspended rather than revoked
- Invalidity date states the date after which the certificate is considered invalid

The Design and Implementation of a PKI

As you consider a PKI, you should organize a formal IT project to produce the results you require. This guide has introduced the wide variety of questions, opportunities, and risks involved in building a PKI. This section aspires to help you put that knowledge to practical use.

Requirements Analysis

The first step in designing your PKI system is to determine what you expect of that system. One of the first considerations is the applications that are supported by your PKI. For instance, if all you intend to support is SSL encryption of Internet traffic, centralized key generation and storage is probably not required. However, if you intend to use keys to encrypt data stored on a file system—data that you might need to recover—centralized key generation and storage becomes more significant. Carefully consider the applications that will leverage your PKI:

- Email encryption to protect sensitive information—Consider whether the email will be persisted in an encrypted manner or decoded and stored in plain text after transmission
- Secure Socket Layer communications through hypertext (HTTPS), file transfer (SFTP), or other protocols—Because of the ubiquitous use of HTTP-based protocols to help work around corporate firewalls, the use of encryption has spread and become a critical application for many Web-oriented services, including Web services—based applications.
- VPNs allow people to connect securely through the public Internet while encrypting the data passed through the TCP/IP channel.
- Protection of ERP systems, such as PeopleSoft, JD Edwards, and SAP—These systems provide encrypted connections to secure communications between systems.
- Single Sign-On (SSO) applications allow users to authenticate to a single network service and access all the services within that network—These services use certificates to validate users to the resources on the network.
- Digital signature for signing documents or signing applications and code to mark it as safe for use.
- Secure electronic transactions through credit cards—Credit card companies are distributing cards with smart chips. These chips contain digital certificates and are used to identify and authenticate the user.

What Is the Primary Purpose of the PKI for Your Organization?

Most organizations will want to establish a PKI to implement security for their employees, customers, and business partners. This may be as simple as building trust with consumers on the Internet and as complicated as meeting the complex requirements of governmental regulations. The driving requirements for this security will help answer many of the questions you will ask as you select systems and establish protocols. For this application, the cost of PKI is overhead, balanced against the cost of security breaches and meeting regulatory compliance.

Some organization may offer PKI as a value-added service. For instance, an Internet Service Provider (IPS) might want to offer SSL services to its clients who purchase Web hosting services. In this case, the balance of cost control and security for you direct customers must be carefully weighed.

Some organizations will want to offer PKI services for sale. In this case, PKI becomes a primary product and the focus of a profit center. Finding the unique blend of pricing and features to distinguish one's organization will serve to determine the services and protocols offered by the PKI.

Design

Once the requirements for the PKI have been established, a system can be designed. The first step at this point is to determine whether to develop the PKI in house or to outsource the PKI. Each side offer advantages and disadvantages.

In-House vs. Outsourced PKI

In-house development and deployment of the PKI can leverage existing resources. If staff already exists that can adequately manage the needs of the PKI and the physical security requirements are in place, it might be economical to develop and deploy with your standing resources. In-house systems can be more flexible and provide you with control of the operation. You are also not at the mercy of the security of an external organization. If your PKI outsource partner were to make changes, they could compel you to make costly modifications on their timeframe rather than your own.

However, many organizations do not have the expertise or additional resources to manage a secure PKI. If protocols are not carefully adhered to, the sanctity of the secured data can be at risk. If under-trained or inexperienced personnel are pressed into service, the results can be gaps in the PKI operation or security. Also, outsource firms are highly motivated to keep the PKI up-to-date and operating smoothly. Their livelihood is dependent on providing a secure operating environment for their customers. If you do not have the existing physical security and expertise, it is often less costly to seek the services of an outside firm.

Mixing and Matching Outsource and In-House

The trusted Root CA is used to validate the subordinate CAs under it. For many organizations, a third-party CA is used as the trusted Root CA for subordinate CAs created by the organization. That shifts the burden of creating a very costly environment to the CA.

If an organization can create their own subordinate CAs, they can be nimble and flexible in meeting their own security needs, while still being able to revoke certificates if a portion of the hierarchy is compromised. But remember, every subordinate CA you create must be secured in its own right. Failure to do so can create a major breach in security and permanently damage the trust you build with your infrastructure users.

Choosing a CA

If you provide externally facing certificates, such as those used by public Web servers, you will want to contract the services of a public CA. That authority will stand as an impartial third-party witness and attest your identity to your customers or business partners.

Previous chapters discussed how to evaluate a CA. Carefully consider their CPS and internal procedures. Query them concerning the services they offer and value-added services that you can leverage. Be certain to have your specific requirements in hand to use as a guideline for the services you need. Also, take time to contact existing customers and learn from their experiences the type of service that you are likely to enjoy. Consider the entire cost of the relationship, not just the price tag on a 2-year, level 3 certificate.

Choosing Components

Some portion of your PKI will be managed in-house. Work with your existing vendors to discover the services they offer that you might already own and can leverage for the build-out of your PKI. For instance, if you have an LDAP server operating within your network, it might have a certificate server extension. You might be able to use that service without adding a new physical server or learning a new operating system (OS). Many companies have Microsoft Windows servers. These servers can be used to create self-signed certificates that are quite suitable for internal use, such as development and test server, internally encrypted application, and the like. They are simple to use and incur little added expense.

Depending on your needs, you might want to investigate best-of-breed solutions. Companies that live or die based on their implementation of PKI solutions may have products that better serve your esoteric needs. Explore the Internet and check with others to learn which products supply the needs of organizations similar to yours. As with a public-facing CA, take time to read customer reviews or speak to customers concerning their individual experiences with a vendor. These conversations often lead to the best decisions.

Establish Practices

To keep your PKI secure, you need to establish proper practices among your operators. You need to develop your own policies and procedures. The recommendations and patterns established by the PKIX working group and other parties provide the framework for creating your own CPS. Those policies must be implemented as training materials and practices used by your operators every day to assure your PKI delivers as designed.

Pilots

When creating a new PKI, it is often best to go slow. Lessons learned as you begin to implement a solution can guide the next step. Even when enhancing or re-organizing a PKI, there is much to be said for taking small, metered steps that lead to the best solution.

Include training of personnel as part of your pilot. Ultimately, the ability to conduct the procedures defined by your PKI requirements is just as important as the servers and systems that operate it day-to-day.

A successful pilot has a set of measurable goals that define success. Implement these measures to discover whether the solution indeed provides the results you require. Take time to compile the lessons learned and use them to improve the process and the next phase of the PKI implementation.

Roll-Out and Monitoring

Once the design is set and the system is confirmed, the wide-spread deployment of the PKI is scheduled and implemented. Without question, regardless of the careful planning and testing, small hitches and obstacles will be encountered. Collecting these issues and documenting them, along with their solutions, will help you refine and improve your PKI system going forward.

Your system should include monitoring of critical systems. You need to audit use of the system and the activities of your personnel. You need to monitor the resources used by your internal systems. Some of these audits may be required for regulatory compliance. Others will help you predict the capacity of your system and proactively allow you to refine and enhance it.

Summary

This chapter has concluded this guide with a close look at the decision-making process involved when choosing between building an in-house PKI and third-party outsource management for your organization. This chapter covered the basics of a PKI and identified the key components of a PKI. It then built on this information by providing practical guidelines for PKI design and implementation.

I hope this guide has helped you to develop and improve your PKI. May your computing be safe and secure!

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.