

realtimepublishers.comtm

The Shortcut Guidetm To



Managing Certificate Lifecycles



Kevin Behr

Introduction to Realtimepublishers

by Sean Daily, Series Editor

The book you are about to enjoy represents an entirely new modality of publishing and a major first in the industry. The founding concept behind Realtimepublishers.com is the idea of providing readers with high-quality books about today's most critical technology topics—at no cost to the reader. Although this feat may sound difficult to achieve, it is made possible through the vision and generosity of a corporate sponsor who agrees to bear the book's production expenses and host the book on its Web site for the benefit of its Web site visitors.

It should be pointed out that the free nature of these publications does not in any way diminish their quality. Without reservation, I can tell you that the book that you're now reading is the equivalent of any similar printed book you might find at your local bookstore—with the notable exception that it won't cost you \$30 to \$80. The Realtimepublishers publishing model also provides other significant benefits. For example, the electronic nature of this book makes activities such as chapter updates and additions or the release of a new edition possible in a far shorter timeframe than is the case with conventional printed books. Because we publish our titles in “real-time”—that is, as chapters are written or revised by the author—you benefit from receiving the information immediately rather than having to wait months or years to receive a complete product.

Finally, I'd like to note that our books are by no means paid advertisements for the sponsor. Realtimepublishers is an independent publishing company and maintains, by written agreement with the sponsor, 100 percent editorial control over the content of our titles. It is my opinion that this system of content delivery not only is of immeasurable value to readers but also will hold a significant place in the future of publishing.

As the founder of Realtimepublishers, my *raison d'être* is to create “dream team” projects—that is, to locate and work only with the industry's leading authors and sponsors, and publish books that help readers do their everyday jobs. To that end, I encourage and welcome your feedback on this or any other book in the Realtimepublishers.com series. If you would like to submit a comment, question, or suggestion, please send an email to feedback@realtimepublishers.com, leave feedback on our Web site at <http://www.realtimepublishers.com>, or call us at 800-509-0532 ext. 110.

Thanks for reading, and enjoy!

Sean Daily
Founder & Series Editor
Realtimepublishers.com, Inc.

Introduction to Realtimedpublishers..... i

Chapter 1: The What and Why of PKI.....1

It’s a Matter of Trust3

PKIs.....4

Digital Certificates5

 What Does the Certificate Contain?5

CAs6

 Are All CAs the Same?.....8

Privacy9

 Encryption.....10

 How Encryption Is Used to Protect Privacy on the Web?12

Authentication.....13

Data Integrity14

Non-Repudiation.....15

Putting It All Together17

Summary18

Copyright Statement

© 2006 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.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 Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.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.

Realtimepublishers.com and the Realtimepublishers 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 Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.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.

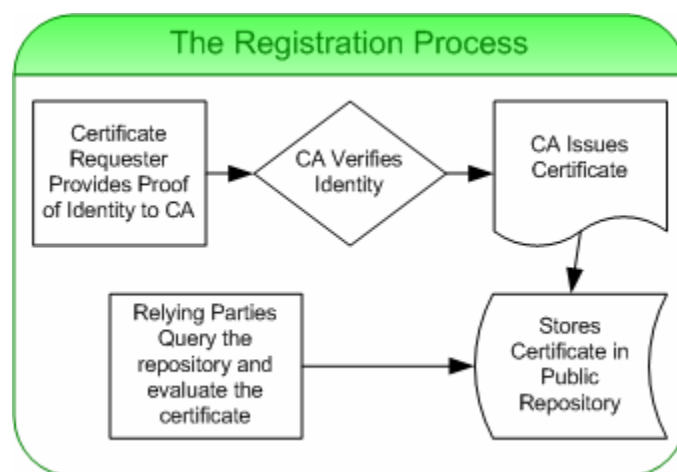


Figure 1.2: The digital certificate registration process.

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.

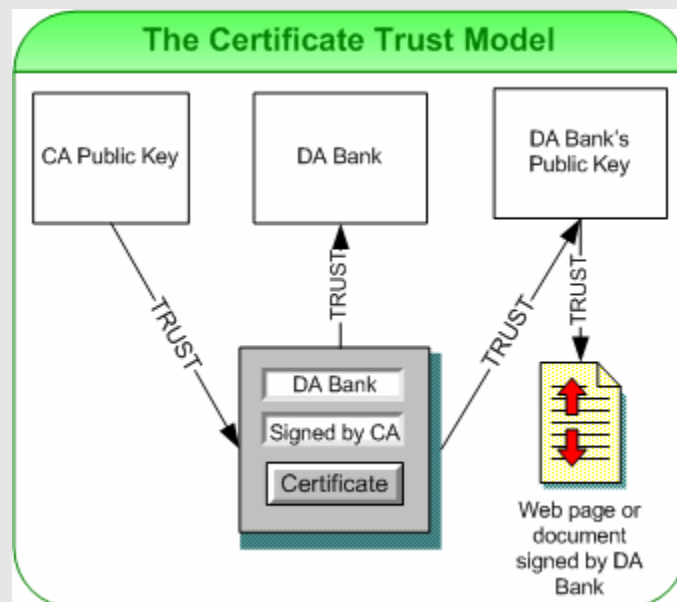


Figure 1.3: The digital certificate trust model.

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.

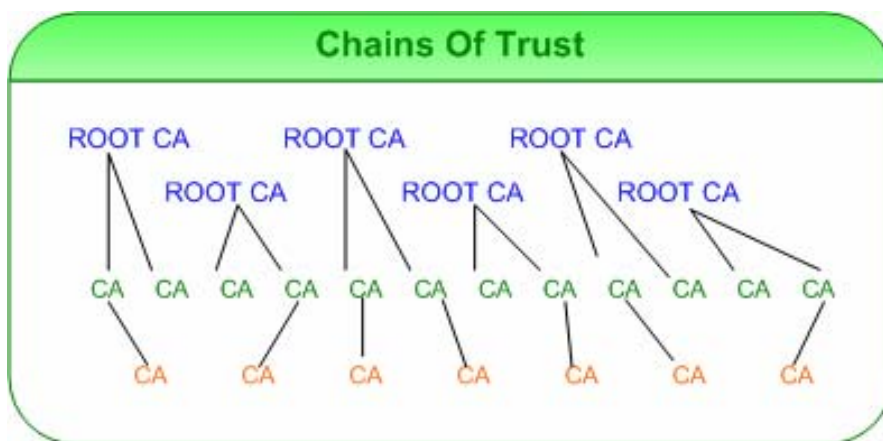



Figure 1.4: The chains of trust between CAs.

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).

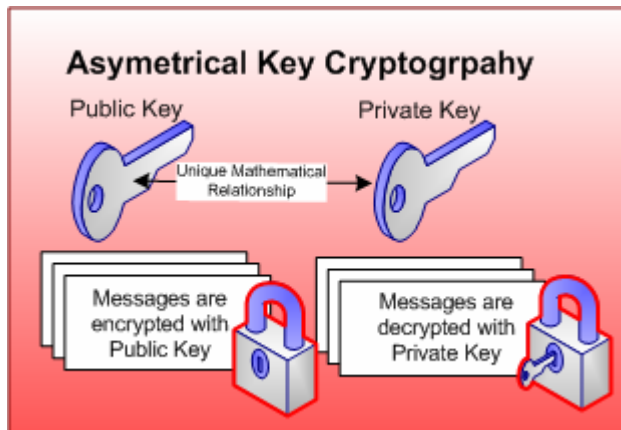


Figure 1.5: Public keys are used to encrypt data and private keys are used to decrypt it.

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).

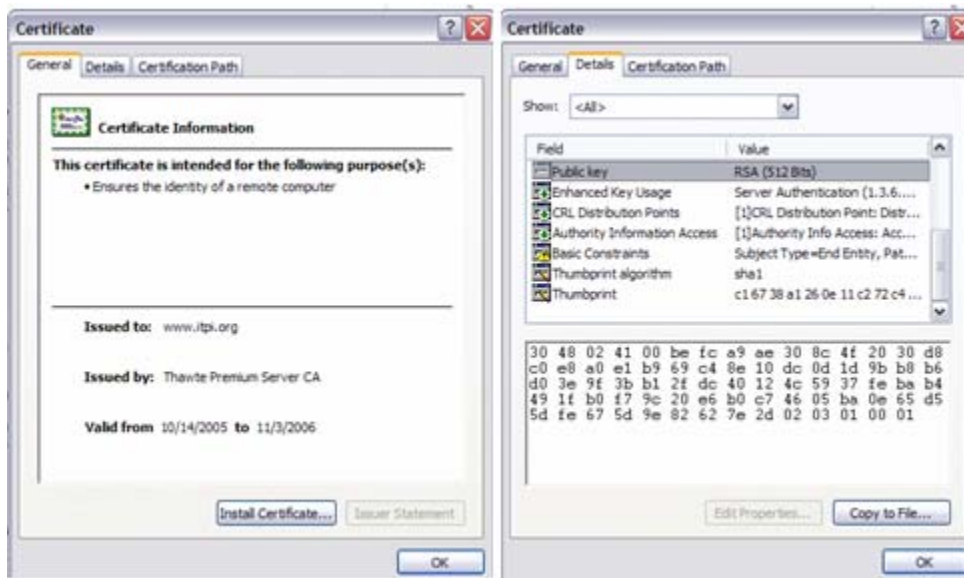


Figure 1.7: By double-clicking on the gold lock icon in your Web browser you can view the contents of the digital certificate on a secure site.

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).

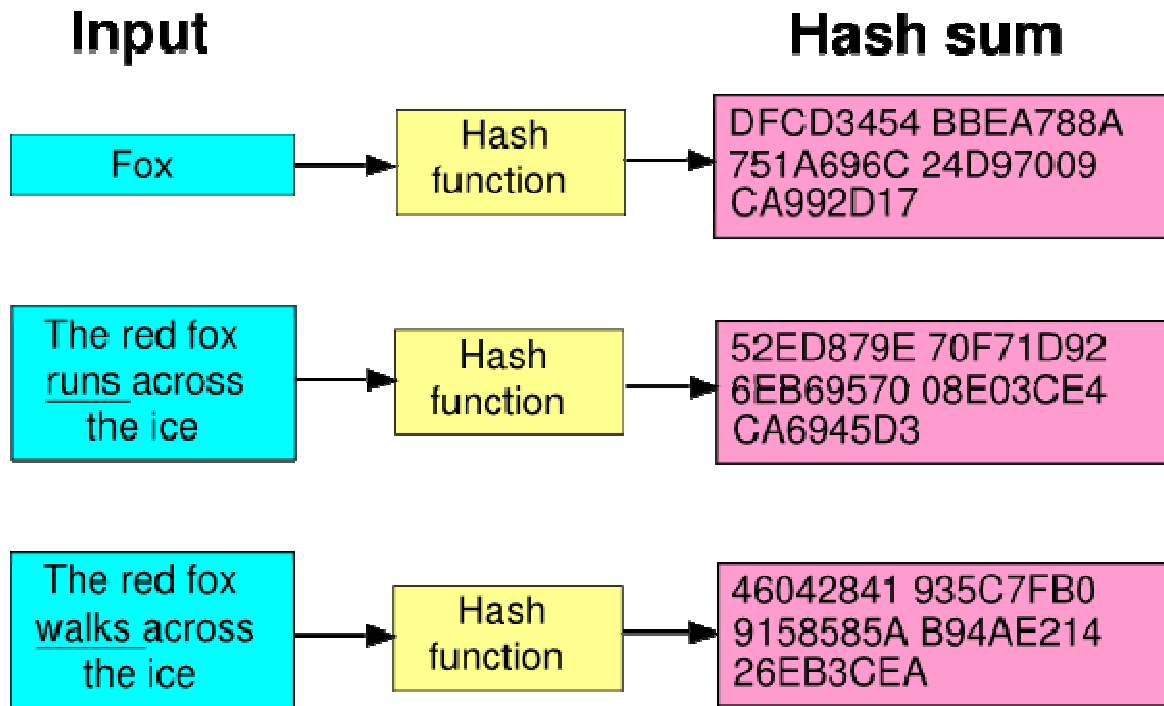


Figure 1.8: A slightly different text would yield a different hash sum or digest.

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.

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>.