Realtime
publishers

The Evolving Threat Landscape and New Best Practices for SSL

sponsored by

Symantec™

Dan Sullivan

Copyright Statement

Threats to Enterprise Information Systems and the Need for Ubiquitous SSL

Threats to information security are becoming more sophisticated, targeted, and persistent. As new protective measures are developed and deployed, attackers seek out alternative means of compromising your systems. It is imperative for businesses to implement best practices to protect valuable information assets. This history of cybercrime and the evolution of malware demonstrate there is no single way to attack a system—there are many. Not surprisingly, there are many defensive measures you can put into place. Anti-malware, email filtering, and vulnerability scanning are just three commonly used methods to address the risk of attack.

One measure is fundamental to many others: encryption. Encrypted data appears to be random data but of course it is not. When data is encrypted, it is essentially inaccessible to anyone who does not have the proper decryption key. The decryption key is an electronic asset, such as a file or string of characters, that is used with a decryption algorithm to transform apparently random text back into its intelligible, unencrypted form. Encryption is such an effective and widely applicable tool, malware developers have used encryption to help avoid detection.

Cryptography is the study of encryption and is a well-established field. There are a number of widely used encryption algorithms in use at any time. Some algorithms are distinguished as strong encryption algorithms because the cost of breaking them far exceeds the likely value of the data encrypted by them. The Advanced Encryption Standard (AES), for example, is a strong encryption algorithm specified by the U.S. National Institute of Standards and Technology (NIST).

The Need for Ubiquitous Encryption

Having encryption standards and employing encryption in high-risk applications is no longer protection enough. Any application, including fairly unsophisticated Web applications, can become a rung on the ladder of attack that ultimately leads to a data breach or compromised system. Businesses need to protect their data as it moves between servers and other devices on the Internet. Securing a backend database is certainly a reasonable and expected course of action, but we are now realizing that other components in the application stack must also be secured.

The need for broader use of encryption stems from three characteristics of threats to information security today:

- More sophisticated attacks
- More targeted attacks
- More persistent attacks

Attackers and cybercriminals deploy attacks that combine these characteristics. The ultimate goal of the attacker is to break through defenses and access data or devices. Consider a few scenarios that highlight how the sophistication, targeting, and persistence of attacks can lead to eventual success for an attacker.

More Sophisticated Attacks

There was a time when scanning for viruses and keeping your operating system (OS) up to date with patches was all that was expected to keep your data and devices secure. In addition, in the early days of malicious software, we were more concerned with viruses that disrupted operations than with data breaches. As attackers learned there was money to be made, they shifted their emphasis from vandalism to significant criminal activity.

Attacks became more sophisticated to avoid detection. Viruses would make slight changes to their code each time they copied themselves to reduce the risk of being caught by pattern-matching techniques. Anti-malware researchers responded by developing an entirely new form of malware detection based on programs' behavior rather than patterns in their code.

We are now at the point where vast numbers of servers, desktop computers, laptops, tablets, and phones are connected via the Internet. The rate of growth in connected devices will likely increase as instrumentation is introduced into appliances, cars, and other components of the Internet of Things (IoT), see Figure 1.1.

From an attacker's perspective, the increasing number of devices presents additional points of access into systems. If a determined attacker cannot breach the defenses around a manufacturer's database, the attacker might try instead to gain a foothold into the corporate network by sending a phishing email with a malicious attached document or by injecting malicious code into a sensor that sends data to an application running on the manufacturer's network.

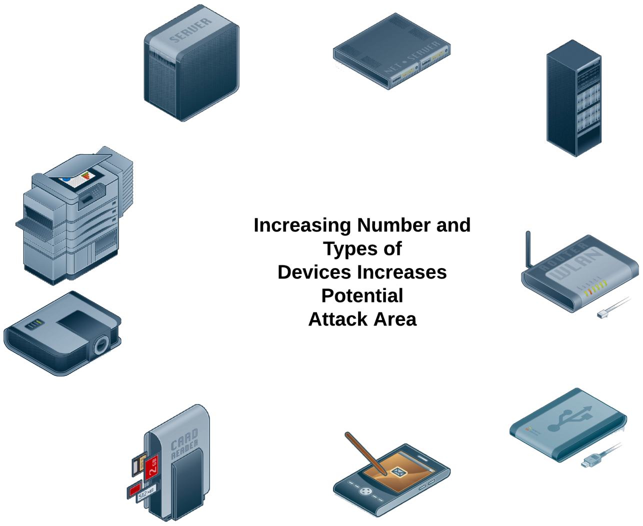

Figure 1.1: As the number of devices on the Internet increases, so does the potential attack area. New types of devices, such as IoT, introduce potential new avenues to compromise connected systems.

From an IT professional's perspective, these devices are additional components that could be used to compromise network, server, or application security. Imagine what goes through a systems administrator's mind as she considers the implications:

- How can we isolate corporate data on personal phones and tablets?
- How do we ensure disgruntled employees are not taking intellectual property with them when they leave?
- Should we perform some type of vulnerability scan on sensors sending data?
- How should we filter IoT data to ensure that it is not malicious?
- What happens when an employee does not encrypt data transfers between corporate servers and external networks?

Let's examine the last scenario in the list, using unencrypted communication channels, to see how it can lead to data breaches and compromised systems.

Data Breach Scenario: Too Little Encryption

Consider a simple example: A customer uses a free, public Wi-Fi service at an airport. The local Wi-Fi network is not encrypted. Fortunately, your developers have anticipated this problem and automatically encrypt login data using Transport Layer Security (TLS)/ Secure Sockets Layer (SSL), an industry standard for encrypting sessions between devices. What they did not anticipate, however, is a vulnerability in the way HTTP is commonly used. HTTP sessions do not keep information about the state of the session; instead, state information about the session is stored in a cookie on the client device. Cookies may include personally identifying information, such as a username and password. This setup is in place to allow the HTTP session to reconnect to backend systems as needed without constantly prompting the user for credentials.

The security risk with this scenario occurs when the cookie, containing personal information, is sent over the unencrypted Wi-Fi; it is transmitted in plain text (see Figure 1.2). Anyone using a scanning application, such as FireSheep, can discover and access the unencrypted information for the wireless network (unencrypted data is equally vulnerable to someone with physical access to a wired network).

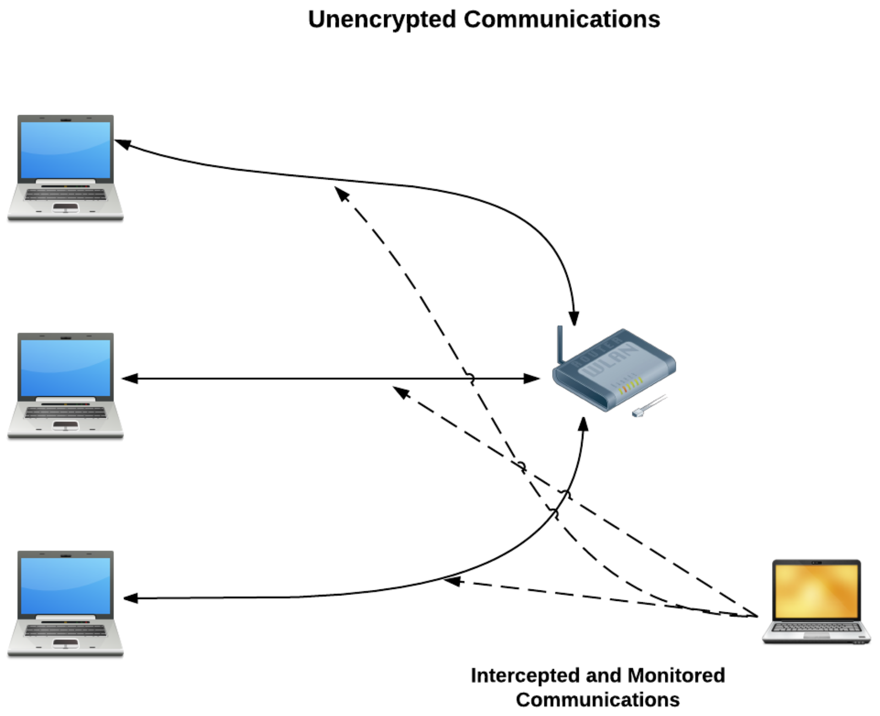
Unencrypted Communications
WLAN
Intercepted and Monitored Communications

Figure 1.2: Unencrypted communications are vulnerable to simple attacks that scan wireless frequencies or wired connections.

An attacker at a major public facility, such as an airport, sports stadium, or shopping mall, could collect large volumes of data from unencrypted wireless networks. Some of this data may contain identifying information, so it could be quite valuable to the right buyer.

More Targeted Attacks

Cybercriminals have adopted many of the efficiency practices one sees in legitimate businesses. These include a division of labor. For example, the person scanning Wi-Fi networks may be moderately knowledgeable about network protocols and filtering data files but does not have much use for the data itself. Instead, the attacker might contract with another cybercriminal to sell subsets of the data. There is likely a market for login information about anyone who works at one of the top 1,000 financial institutions in the United States. Similarly, any scanned data from engineers at electronic manufacturers may be of value to someone committing industrial espionage.

Some cybercrime market offerings are supplier driven. Selling personally identifying information collected from compromised networks is a good example. Other products are more demand driven. According to the RAND Corporation's recent study on cybercrime markets (L. Ablon, et. al. Markets for Cybercrime Tools and Data: A Hackers' Bizaar), there are providers of "stolen-to-order" data, which would be especially prevalent in industries with valuable intellectual property. Stolen-to-order is just one example of more targeted attacks that exist today.

Initial Entry: We Are Part of the Problem

Let's continue with the example scenario of stolen identifying information. The data collector decides to sell the information he collected to a black market data broker. The data broker then proceeds to sell it to someone interested in breaching a large U.S. bank. The bank attacker may collect financial institution data from a number of brokers. By collecting data on multiple financial institutions, the bank attacker reduces the risk that someone else will determine the ultimate target. By collecting data from multiple brokers, the attacker increases the chances of collecting a large amount of user identifying information from the target institution.

The bank attacker is interested in a particular system at the bank. There are many layers of defense around bank systems, so the attacker decides to focus on gaining access to the network. From there, the attacker can collect more information on everything from employees to the network architecture (see Figure 1.3).

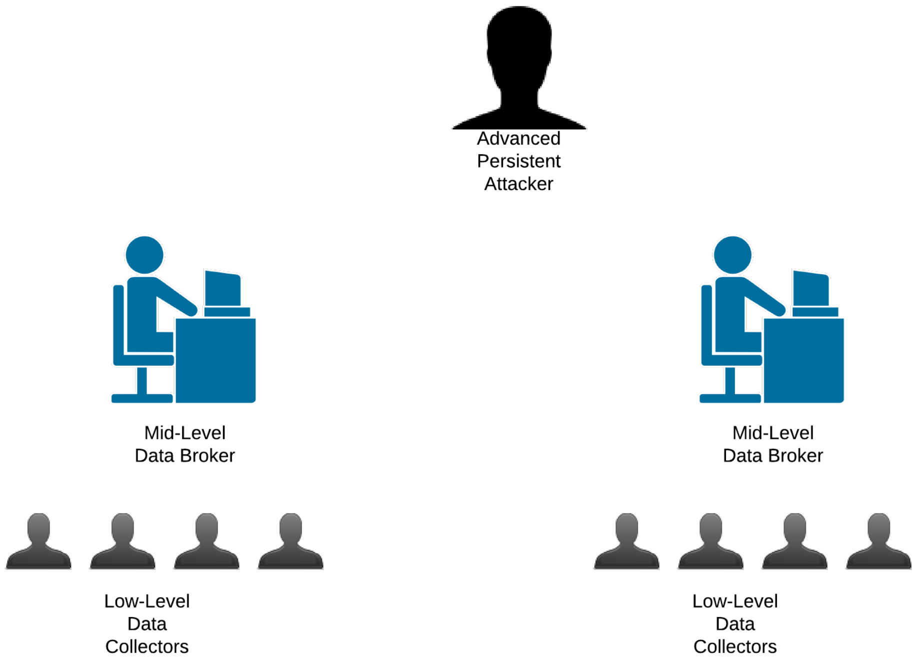

Figure 1.3: Cybercrime entails multiple levels and specialized services. In this case, low-level data collectors sell data sets to brokers who sell them to the final end user—a specialized attacker.

If lucky, the attacker may have culled a username and password from the stolen data purchased from data brokers. Assuming none of the username and password combinations are useful, the attacker turns to spear phishing.

This technique entails crafting emails targeted to a particular person. The bank data includes email addresses, so it is a simple matter to compile a list of potential targets. The bank attacker can filter the list based on the data collected from the unprotected Wi-Fi network. For example, other data from the victim's session could give information about the area of the bank the person works in or about collaborators in the bank. An enterprising attacker may use a social media source, such as LinkedIn, to collect additional data about possible targets. All of these efforts are designed to limit the list to high-valued targets and to collect enough information to craft a personalized spear phishing email that will actually hook the victim (see Figure 1.4).

A phishing email will include an attachment or a link to a compromised Web site. The goal of the text of the email is to lure the recipient into clicking on the link or open the attachment. At that point, malicious code would run that establishes a foothold for the attacker. For instance, opening a malicious document may trigger the download of a remote control program that allows the attacker to gain entry to the victim's computer and use it with the privileges of the victim.

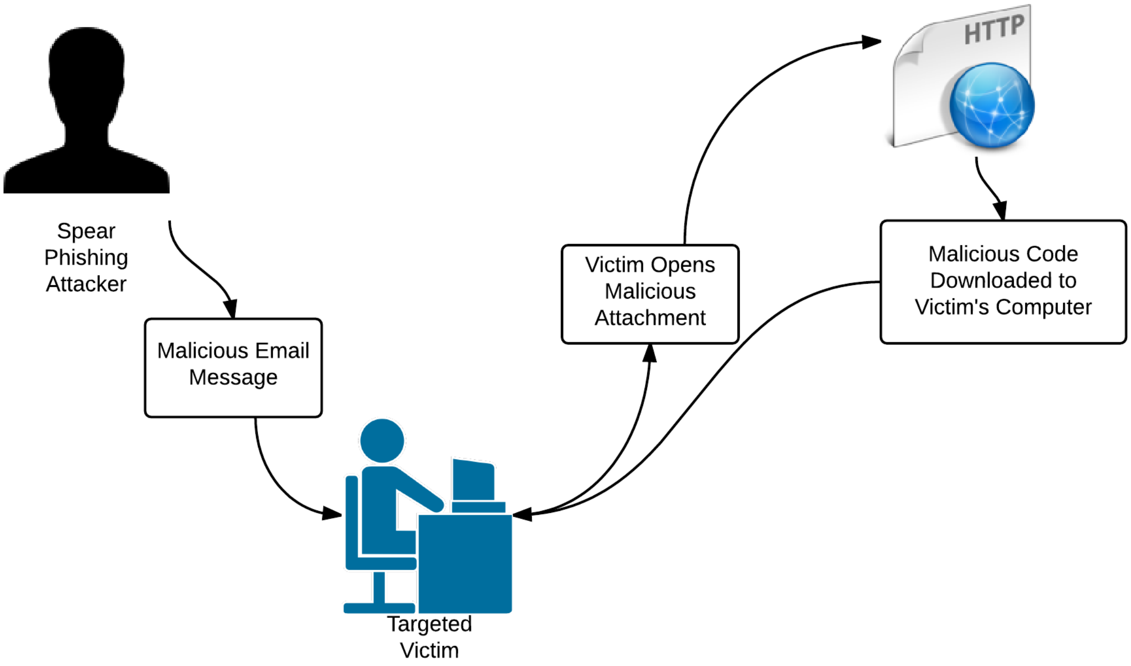
HTTP
Spear Phishing Attacker
Malicious Email Message
Victim Opens Malicious Attachment
Malicious Code Downloaded to Victim's Computer
Targeted Victim

Figure 1.4: Spear phishing targets particular individuals. The attackers goal is to lure the victim to open a malicious attachment or visit a compromised Web site so malicious software can be downloaded to the victim's computer.

Now that the attacker is in the network, a new phase of the attack can begin.

More Persistent Attacks

Attackers targeting a particular business or institution may have many reasons to choose their victim, including money, revenge, and politics. Targeted attacks may be persistent; that is, the attacker continues to probe, scan, and test for vulnerabilities over an extended period of time. Relatively small successes, such as compromising an intern's laptop, could be just the first step. Each success provides a new vantage point to collect information about the business. A persistent attack should be seen as a long-term effort by an attacker to gain access to systems or digital resources.

Persistent Attacks: Incremental Successes

Let's return to our bank attack scenario. The attacker has gained low-level access to a device on the corporate network. Probably, one of the first orders of business is to collect information on the structure of the network and the software running on it. The attacker may try a wide range of methods to gather information:

- Run vulnerability scanners on network devices to check for known vulnerabilities
- Collect contact information from the compromised employee's email to target a new set of potential victims
- Check for open ports used by common Web services, such as databases and Web servers
- Collect network packets; this information is useful if internal traffic is not encrypted

Each of these steps can help the attacker compromise additional systems. Imagine if a vulnerability scan finds an unpatched Web server. The vulnerability allows the attacker to run code with some modest set of privileges. The attacker quickly discovers the Web server is a frontend for a database application.

The attacker writes a script that executes on the database and returns metadata about the database structure and the type of data stored in it. The attacker decides this is a low-priority database. Instead of querying that database for a data dump, the attacker queries the database for a list of other databases. Database management systems provide methods for linking databases to facilitate sharing information. It is also a method for attackers to learn about other systems on the network, as Figure 1.5 shows.

When an attacker undertakes an advanced persistent threat, such as described in this scenario, there is no single tool or deadline that drives the endeavor. The goal is to attack until the system is breached and the data is collected, the operations are disrupted, or some other malicious outcome comes to fruition.

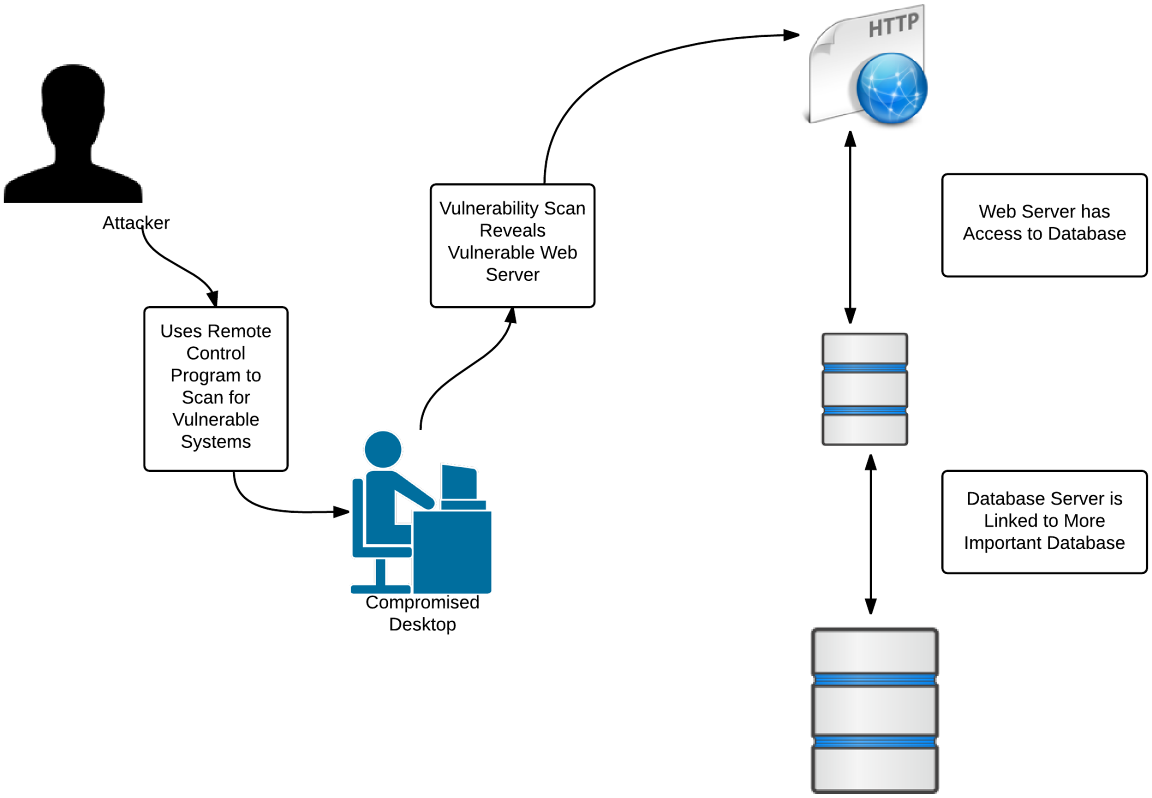
HTTP
Attacker
Uses Remote Control Program to Scan for Vulnerable Systems
Compromised Desktop
Vulnerability Scan Reveals Vulnerable Web Server
Web Server has Access to Database
Database Server is Linked to More Important Database

Figure 1.5: A persistent attacker will continue to look for vulnerabilities in the system, including liked databases.

Security Breaches at Large, Respected Companies

Sophisticated, targeted persistent attacks happen to a wide range of institutions. Many news headlines have included the names of well-known companies along with phrases like “breached” and “hacked.” These victim companies do not necessarily do things all that differently from other companies. The increasing economic incentives for cybercrime are driving all aspects of the cybercrime market, including increased attacks on businesses.

Security professionals will tell you there is no silver bullet. Attacks will continue. New defenses will be deployed and they will work well, at least for a while. When attackers cannot breach a defense, they try to go around it. If the value of the attack is large enough, attackers will invest the resources needed to achieve a successful attack. So, it is safe to conclude, you and your business are:

- Confronted by cybercriminals with extraordinary financial incentives to steal data from businesses,
- Facing a sophisticated market of cybercriminals that range from low-level participants, such as money mules, to highly skilled coders finding and exploiting vulnerabilities in widely used software,
- While contending with personal devices outside of your strict control,
- And fielding demands from lines of business to support innovative applications that often include collecting and sharing data across the Internet.

There is no question that the threats to businesses are real and substantial. There are ways, however, to mitigate the kinds of threats discussed here.

Addressing the Danger Posed by Advanced Persistent Threats

Addressing threats emerging from the existing cybercrime landscape is a constant, ongoing process. It is essential for IT professionals to adopt a strategy that is equally constant and ongoing. In particular, the strategy should incorporate four principles:

- There is no silver bullet
- Individual countermeasures fail
- Security in depth is essential
- Encryption is the first and last line of defense

Each of these principles helps guide the selection of specific countermeasures and procedures that constitute a response to current threats.

No Silver Bullet

As much as we would like one, there is no single tool, application, or procedure that we can implement that will protect our information assets and devices. There are several reasons for this reality.

Technology solutions, such as anti-malware and vulnerability scanning systems tend to focus on a single type of threat or problem. Anti-malware, for example, scans network traffic and downloaded files for malicious content. Such content may be detected by matching patterns of incoming data to known malicious software. In other cases, malware is detected by observing the behavior of a suspicious program and discerning malicious behavior. This measure blocks attackers trying to inject malicious software into your systems.

Vulnerability scanning systems, in contrast, examine devices to determine whether any of the software on those devices has known vulnerabilities that an attacker could exploit. In this case, the objective is to find weaknesses in your own systems that could be exploited by an attacker in the future.

As these examples demonstrate, there are multiple ways for attackers to compromise your systems. This certainty calls for multiple techniques to prevent such an event.

Individual Countermeasures Fail

Even with the best strategies, plans, and implementations, countermeasures can fail. Again there are several reasons for this reality. Human error may be the root cause. An employee may remove a malicious email from a spam folder believing it is a legitimate message. After opening the message, the employee clicks on a link to a compromised Web site that downloads malicious software. A network administrator might misconfigure a router that leaves it vulnerable to manipulation by an attacker. An application developer might improperly code a SQL query built using input from a user, leaving the query vulnerable to a SQL Injection attack. There are certainly other ways human error can lead to the failure of defensive measures, but this set of examples demonstrates the breadth of potential errors.

Software bugs are another reason countermeasures can fail. Well-publicized bugs such as the Heartbleed bug in the OpenSSL implementation of SSL/TLS show that even widely used software may contain vulnerabilities. The Heartbleed bug allowed attackers to read memory and capture security information, such as encryption keys.

Errors during the normal operations of some security applications can leave devices vulnerable to attack. For example, a systems administrator may automatically push an OS patch to the Windows 7 devices on the network. Although most of the patches apply successfully, some fail. Unless those failed patches are detected and corrected, the devices will still contain the vulnerabilities corrected by the patch. In other cases, a network failure can disrupt routine vulnerability scans. Devices that are unreachable during the network outage are not scanned and may harbor vulnerabilities that could have been detected by a successful scan.

The potential for failure in individual security components is one of the reasons security in depth has become a best practice.

Security In Depth Is Essential

Security in depth is the practice of using multiple defenses to protect your information assets. The goal of security in depth is to protect data and devices even in cases where one or more defensive measures fail. Security in depth includes redundant and complementary measures.

An example of a redundant countermeasure is to deploy anti-malware systems on the network and on individual devices. Traffic coming into or leaving the corporate network is scanned for malicious content. Concurrently, all devices on the network run local copies of anti-malware software. If the local anti-malware software is not up to date, for example, it may miss detecting the latest versions of a particular virus; however, the network scan could detect and block the malicious content.

Complementary measures address different kinds of vulnerabilities. For example, consider a situation in which malicious code with a SQL Injection attack went by the anti-malware systems. The malicious code is now in a position to execute against accessible databases. Routine vulnerability scans have identified applications vulnerable to SQL Injection attacks and they have been patched. In spite of the fact that malicious code is on the network, it is not in a position to utilize the code against the accessible database systems.

Security in depth employs a mix of technologies, and one of the more fundamental of these technologies is encryption.

Encryption: The First and Last Line of Defense

Encryption is such a fundamental operation in security processes that it is hard to imagine a reasonably secure environment without it. Encryption is the first line of defense when sharing information over the Internet or with mobile devices. If you are sending important strategy documents, private patient information, or confidential details of financial transactions, then it should be encrypted. As discussed earlier, an attacker with access to unencrypted Wi-Fi can easily monitor and capture data transmitted over that wireless network. If a person using that Wi-Fi network does not encrypt data before sending it over the Internet, it could be exposed to an attacker. If the attacker does collect data that had been encrypted on-device, then it would be of no use to the attacker. The captured data would look like random data. Without the encryption key, the attacker would have no way to decrypt and use the encrypted data.

Encryption is also your last line of defense. As noted earlier, security countermeasures can fail. Attackers with time and resources can launch prolonged, persistent attacks that slowly compromise or work around each of your defenses.

Consider a possible scenario: While working at home, an employee visits a compromised Web site, which results in malicious code being injected to that device. The next day, the employee connects to the corporate network. The compromised device detects a connection to a corporate network and downloads additional tools to explore and exploit the network. After scanning hundreds of machines, it finds a script that logs into a database. The script contains the username and password of an account used to perform basic data loads.

Next, the attacker uses those credentials to log in to the database and finds the account has several administrative privileges. Although the set of privileges should not be sufficient to compromise the database server, an unpatched vulnerability in the database allows the attacker to gain database administrator privileges to the server. At this point, all data in the database is available to the attacker who downloads the most promising looking data.

There is now little that can be done to prevent the unauthorized use of the data. It has been downloaded from an internal server and is no longer subject to the controls and security measures in the corporate network. The only way to prevent this type of unauthorized use is to store encrypted forms of data in the database. In that, the attacker would have spent significant effort to break into the network, scan for vulnerabilities, exploit multiple vulnerabilities, and finally identify and download potentially valuable data. Unfortunately for the attacker, the data is useless because it is encrypted.

In the event other security measures fail and data is stolen, if the data is encrypted, it will be of no use to the thieves. Encryption stands as the last line of defense even if others are compromised.

How SSL/TLS Certificates Protect Assets

Up to this point, the discussion has focused on the risks your business faces from cybercrime and the challenges you face in protecting your information resources. It is now time to turn to a method and set of practices to enable encryption-based protections for your data.

SSL and TLS are protocols for securing communications and authenticating users and services. The purpose of this guide is to discuss how to apply SSL/TLS technologies to protect data and your systems. The low-level details of how these protocols work are outside the scope of this guide, but it is important to understand basic features of the SSL/TLS protocols. First, TLS is a successor protocol to SSL. The two are functionally similar enough that for convenience, this guide will simply refer to both as “SSL.” The reader can assume that this reference indicates “SSL/TLS.”

SSL Certificates

SSL is an encryption protocol that uses digital certificates, or simply certificates. Certificates are often compared with passports because they are secured artifacts that vouch for the identity of the holder. They share additional characteristics with passports:

- A standardized form agreed upon by all issuers
- The trustworthiness of the artifact depends on the issuer
- The artifact identifies the holder
- The artifact is valid for a set period of time
- The artifact is tamper resistant

There are also parallels with how certificates and passports are used. A boarder patrol officer would not let someone into a country without checking for a valid passport. If it appears that someone tampered with the passport, the boarder control officer may turn the visitor away. Similarly, administrators may want to implement controls over what servers, desktops, laptops, and other devices communicate with corporate devices. Certificates act like passports for servers and other devices. Devices implement policies that determine when another device will be allowed to communicate with it. For example, a server policy may dictate that communication is allowed only with devices holding a valid certificate from a small list of trusted issuers. Any device with a certificate from one of those trusted issuers is allowed to establish communication channels with the server; all others are not.

Just as different countries issue certificates, different organizations can issues certificates. The SSL standard is an open standard that can be implemented by anyone with interest in doing so. As the need for SSL certificates is so common, OS, browsers, and other applications that depend on them include a list of trusted issuers.

Many of us depend on these lists by default. For example, we expect our browsers to validate certificates when we use sensitive services, such as online banking. An attacker who wants to set up a bogus Web site that appears to be a legitimate bank site may be able to do so. The attacker might be able to secure a URL that is close to a legitimate bank's URL so that customers that make a simple typing mistake may end up at the bogus site. The attacker may even create a certificate that states it is the legitimate bank. Fortunately, our browser would recognize that the certificate is not from a trusted vendor and block access to the site.

Later chapters will describe further details about how certificates are used, but first let's examine additional details about certificates and the standard body established standards and best practices. Certificates come in several forms:

- Domain validated
- Organization validated (OV)
- Extended validated (EV)
- Wildcard

These certificates all perform the basic passport functions described earlier. They are, however, suited for different needs. The first three require varying levels of validation and vetting. If concerns about customer trust are primary, an EV certificate is the best option; however, OV might meet less-demanding use cases. Domain certificates are appropriate in cases where minimal trust is required. Wildcard certificates, unlike the other types mentioned, are distinguished by the administrative convenience.

Domain validation certificates are used for Web sites and Web applications within a single domain. For example, if you own the domain "mycompany.com," you would be able to procure a valid domain certificate for that domain. That certificate, which is implemented as a file, would be stored on your servers and accessible to SSL programs that authenticate identities and encrypt data.

Domain validated certificates offer encryption, but only provide basic authentication because they only confirm that the person applying for the certificate has the right to use a specific domain name, not whether the organization is trustworthy (or even exists).

Domain validation requires some basic efforts on the part of the certificate issuer, such as verifying the owner of the domain name. These requirements are sufficient for some uses where demonstrating a high level of validation is not necessary. For major businesses, such as banks and online retailers, there is often an interest in implementing more extensive security controls.

OV certificates require all the validation measures of domain certificates as well as additional checks on the organization. OV certificates offer reliable authentication and encryption for the cloud because they validate that the organization claimed to be responsible for the domain or server actually exists, and that the person applying for the SSL certificate for that domain or server is an authenticated representative from that organization. This certificate type is commonly used by banks, retailers, insurance companies, and other organizations that have a direct financial stake in ensuring the confidentiality and trust of their client communications.

EV certificates require substantially more effort on the part of an issuer to verify that a legitimate organization or business owns a domain. Instead of just checking domain registration, for example, issuers validate physical addresses of an organization and check other public records. EV certificates are the best choice for server-to-browser connections because they offer the strongest level of authentication and the clearest validation that the connection is secure. With EV certificates, the legal, physical, and operational existence of the organization is verified, as is the right of that organization to use that domain. Using EV ensures that the organization's identity has been verified through official records maintained by an authorized third party, and that the person requesting the certificate is an authorized agent of the organization.

Wildcard certificates extend the functionality of a single server certificate to apply to multiple subdomains within a domain. This certificate type allows for subdomains such as "www," "support," and "admin."

The types of standards that are offered and the procedures implemented prior to issuing a certificate can vary. The CA/Browser Forum is a voluntary organization of certificate issuers, known as Certificate Authorities (CA) and browser vendors. The organization works to promote trust in SSL technologies by establishing standards. For example, in 2005, the CA/Browser Forum issued the Extended Validation Guidelines for EV certificates. The standards include details on the use of color and icons to help browser users recognize an EV-enabled site.

Summary

Businesses and organizations face constant, sophisticated threats to their information infrastructure and digital assets. Encryption is a fundamental technology that provides both the first and last lines of defense. SSL certificates are a crucial component for implementing SSL and other security controls. SSL is a set of protocols that enable encryption to protect data privacy and authentication of servers and applications to promote trust. The next chapters will examine methods for using SSL and choosing the right types of SSL certificates for your organization's needs.