



internet security

keeping up with security challenges

a backgrounder from hp

## how to avoid getting on the front pages for the wrong reasons

This backgrounder is intended to take you “from 0 to 60” on the topic of Internet and intranet security. It does not address specific server-based security issues, and it assumes that you are aware of the basic security requirements of authentication and authorization (such as passwords and access lists) on at least one operating system. After reading this paper, you may think that the Internet is a frightening neighborhood to play in. It is, indeed, and HP Services, a security-savvy professional services organization, can help your company stay safe.

Before the advent of PCs and the Internet, computer systems and the data on them were protected by physical means. The computer itself was locked behind strong walls, and access to it was granted through secured terminals, with a dedicated connection for each terminal. With the emergence of PCs and LANs, information sent to terminals is no longer protected, because it can be downloaded to a PC and taken out of the building on a floppy disk or similar storage medium. In fact, with the proper software (see *sniffer* in the glossary), you can monitor data on the LAN, including data destined for PCs other than your own.

The Internet is an internetwork that connects various smaller networks. No one owns or controls it. Only *TCP/IP* addresses, which allow host computers to be identified individually, are controlled by a central authority. This authority is the *Internet Corporation for Assigned Names and Numbers (ICANN)*.

Data that travels from one host to another does not follow a fixed path. It is impossible to predict which paths and through which hosts a data packet will flow. You must therefore assume that every packet sent over the Internet can be intercepted and/or altered. For this reason, all sensitive communications should be encrypted, perhaps by using public-key technologies.

Many different kinds of information travel over the Internet. Various protocols are used to transmit each type of information. The *Hypertext Transfer Protocol (HTTP)* is used for Web pages, the *File Transfer Protocol (FTP)* for data files, and the *Simple Mail Transfer Protocol (SMTP)* for e-mail.

Think back to some of your recent Web transactions, and consider how safe they were. Did you write anything in an e-mail message that you wouldn't want displayed on a billboard? Do you buy products over the Internet using a credit card? One risk in these practices is that individuals can set up hosts and use sniffer software to intercept and store all of the information that passes through them.

## encryption

To prevent someone from reading the information you are sending across the Internet, you can encrypt it. It isn't the *method* of encryption (or algorithm) that must be kept secret; it's the *key* used in the algorithm. No matter what kind of *encryption* you use, hardware encryption is generally more secure than software encryption, because in software encryption, the key is held in the computer's memory, making it accessible to hackers.

### symmetric encryption

With symmetric encryption, the same key is used to encrypt and decrypt the message. Therefore, the sender and recipient must agree on the key to use. If the key is compromised, the sender and recipient must agree on a new key. The problem is how to exchange keys without someone else intercepting the message. (In old spy movies, you may recall that characters kept a codebook in a briefcase shackled to their wrist.) The keys must be exchanged in a secure manner before messages can be encrypted using the keys.

### asymmetric encryption

With asymmetric encryption, or *public-key encryption (PKE)*, keys are never exchanged. The sender and recipient use a *key pair*. Either key in the pair can be used to encrypt a message, as long as the other one is used to decrypt the message.

Both Netscape and Microsoft® Internet Explorer can use PKE when encryption is required. A lock icon at the bottom of the Internet Explorer screen and a shackle icon on the Netscape screen indicate that the browser and server have established a secure connection. (See *Secure Sockets Layer* in the glossary.) The server initiates the switch to secure mode, and the browser responds to the switch. The user does not need to take any overt action.

The server software sends its public key to the client software (browser). The browser picks a key at random, encrypts it under the server's key, and sends it to the server, which decrypts it. The browser and server now have agreed on a symmetric key, which was exchanged in a secure manner.

Why don't the client and server continue to use PKE instead of exchanging and using a symmetric key? PKE is processor intensive, whereas symmetric encryption is not. Although encryption within a Web browser is transparent, data files and e-mail attachments usually must be encrypted explicitly before they are sent. Some systems, such as Lotus Notes, have transparent encryption built in.

Even if the text on a Web page says that your credit card information is being sent securely, it is not secure unless the lock or shackle icon appears on the browser screen. And even if the data is *transferred* securely, it may not be *held* securely on the server. This is how it was possible to steal and post credit card numbers from CDUniverse customers—an event that hit the front pages in January 2000.

Several companies, including HP, manufacture external devices and PCI cards that increase the processing speed of PKE while reducing processor overhead. Examples of these products include the HP AXL300 and HPAXL600L PCI Accelerator Card for ProLiant servers and the HP Atalla PayMaster, SignMaster, and WebSafe2 Internet security processors. Drivers for these PCI cards are included with many operating systems, such as current versions of the Microsoft Windows family.

## certificates, smart cards, and other authentication

Authentication is the act of verifying that things (or people) are what (or who) they claim to be. In the physical world, such verification can be accomplished by looking at the object, if it is familiar, or by looking at a picture identification card. We can assume that if an official agency, such as a state department of motor vehicles, issued the identification card, the person carrying it has proved his or her identity already. By contrast, passwords are easy to steal or guess, so they are not considered strong forms of authentication.

On the Web, individuals can use *certificates* from a *certification authority (CA)*, or they can generate their own for authentication. Assuming that the CA has exercised due diligence, a signed message can be taken as authentic. Some CAs issue two or more certificates to a user; one certificate can be used for signatures and another for encryption. For example, the user can generate a certificate for signatures and have a CA sign it to verify that it is authentic. The CA can generate a certificate for encryption and then retain it. This allows users to recover their encryption keys from the CA after identifying themselves if they lose *smart cards* or passwords that protect their private keys. This is important when data is being stored encrypted; otherwise, data can be permanently lost.

Some companies and entire countries are moving toward the use of smart cards for authentication. It is not difficult to replicate a credit card, including the magnetic stripe. If you have captured or guessed the PIN (or one isn't used), you can now create bogus transactions on the Web or over the phone. On the other hand, because the smart card is intelligent, the *public-key infrastructure (PKI)* can ensure that the original card is in use, rather than a copy of the card. This is because a smart card, unlike a magnetic stripe card, can use symmetric or asymmetric encryption on all communications. Information going to the card can be encrypted under its public key as recorded by the issuing authority (American Express, for example), and the card can decrypt the information using its private key. A copy of the card would not be able to decrypt the information, and the reader can disable the card so that it cannot be used again. Even if someone learns the user's PIN, the card cannot be copied because no one can access the private key locked inside the card.

Smart cards can also be verified offline. The user inserts the card in the reader and enters a PIN, which the card itself can verify. Again, unlike a magnetic stripe card, the smart card can answer yes or no without compromising the PIN presented to it (the PIN entered on the reader) or the PIN that unlocks it (the PIN stored inside the card).

Other types of authentication can be used with or without smart cards. Some methods are biometrics (retinal or facial scans and fingerprints), tokens (similar to a smart card but not requiring communication between the card and the server), and handwriting recognition.

## key management

No matter which *encryption method* is used, it's important to handle and store the keys securely. As you may remember from old spy movies, people are not secure. No one person should know all parts of nonpublic keys being used.

There are various ways to split keys into separate parts. One way is to have several people choose a number that has the same number of characters as the key that you are generating. All of the numbers are entered into a device that combines them in some way to generate the key. The device can then inject this key directly into the

encryption device. Because the key may need to be regenerated in the event that the encryption box is destroyed, the key parts are put into separate sealed envelopes and stored in a manner that protects them from harm and that needs more than one person to retrieve them.

The same key should not be used for more than one function. PIN verification, message encryption, and data storage encryption should all use different keys. Like separation of duties, separation of keys (or key variants) leads to better security.

## **data partitioning**

For performance reasons, you might opt to keep any materials intended for the public directly on the Web servers themselves. Normally, this includes only information that people could otherwise locate via other advertising channels (catalogs, images, marketing brochures, and so on). Any dynamically generated data (such as stored billing and shipping information) should be kept as far out of reach of the Internet as possible.

Furthermore, any data that your customers supply via Web-based forms should be removed immediately from the Web server through as many *firewalls* as needed to safely secure it. This is the most fundamental security precaution that you can take. Never store any private information on the Web server itself, because you cannot be sure that the server is always under your control. Should an attack occur, a few Web pages may be altered, but your important assets will remain secure.

## **addresses and name servers**

Each host on a TCP/IP network must have at least one address assigned to it. A host can have multiple addresses assigned to it—either many on one physical interface or one per many physical interfaces. Multiple addresses can be used to spread the load across multiple network interface cards, or some addresses can act as administrative addresses. A good example of this is the TSM port on an HP NonStop S-series server. This port should be put on a private LAN (perhaps with just the administrative console, or with other administrative consoles).

Addresses are stated as four sets of decimal *octets* (8 bits) separated by periods. An example is 192.244.38.6. A name such as “www.hp.com” is linked to the host’s IP address by a name server, a program that translates domain names into IP addresses. Name servers are arranged in a hierarchy somewhat like an inverted tree, with the *InterNIC* name server at the top. If a name server cannot locate a host, it asks the next one up the chain until it reaches the root InterNIC name server. If the InterNIC name server cannot locate the host either, it passes the request to the domain name server (DNS) to which the host belongs. For example, “inraweb.hp.com” is a host that belongs to the “hp.com” domain. When the domain was registered, the HP hostmaster notified InterNIC of the IP address of its primary name server. Therefore, the InterNIC root name server knows that the hp.com name server is required to know the addresses of all servers belonging to the hp.com domain. Whenever a request is made to connect to a host that ends in hp.com, this name server is asked to make the translation from name to IP address.

References to a host by name can be “hijacked” by convincing one or more name servers that a bogus host is the real host. One way to do this is to intercept a name server’s call to a more authoritative name server and return the address of the bogus

server to it. Future references to the host that uses this name server will then be routed to the bogus server. There is very little that can be done to prevent this form of hacking, but you should be aware that it can happen.

Unscrupulous businesspeople can use DNS names to confuse people and cause them to disclose more information than they normally would. For example, a server named "aolbilling.com" could be part of a scam to have people with AOL accounts divulge credit card numbers. The scammer would send e-mail messages to AOL users telling them that they need to reregister their credit cards. Because the server name sounds as if it could be part of AOL, many users will follow the directions.

Similarly, some apparently reasonable sounding DNS names may be owned by people or companies that have nothing to do with the root name—for example, "MichaelJFox.com." This is another way to deceive Internet users intentionally with DNS names. Entities that own DNS names based on brand or celebrity names that are not their own are called *cybersquatters*.

## FTP and HTTP

FTP is used to move files around a TCP/IP network, such as the Internet, whereas HTTP is used to render or display Web pages written in *Hypertext Markup Language*, or HTML (and, in the future, XML). HTML code looks much like the code for TeX, a document formatting system written by Donald Knuth. It consists of text between tags. A tag might tell the browser that the text enclosed is to be displayed in a specific point size, font, or color. Some tags define the enclosed text as a *URL*. Because HTML is sent to the browser as source code, it is possible to see how a Web page is written. This can lead to security breaches if a hacker can discern your host and network structure from the source code. Files can be downloaded from your host in HTTP mode as well as in FTP mode.

FTP is normally used to download files from a server to a client machine. Browsers as well as standalone programs can perform FTP transfers. Unless your site is required to provide FTP access, a firewall should be used to prohibit it. If the site is designed to provide FTP access, the system security must be carefully set to prevent unanticipated or unauthorized access. Inadequate server security could allow access to all directories as well as permit downloads of object code used to run the site. Even worse, it could allow information to be overwritten so that the next visitor will not get what you want to send but will get whatever the hacker uploaded.

## networks and routers

As mentioned previously, every host has at least one address. A network of hosts is defined by the addresses of the hosts on it and the *subnet mask* (see glossary). For example, a subnet of 255 hosts could have addresses from 192.244.38.0 through 192.244.38.254. Because the first three octets are the same, the hosts are de facto on the same subnet if the mask is set to 255.255.255.0. If any of these hosts need to talk to a host located at 193.220.6.72, they will have to communicate with another subnet. Communication between networks is done through *routers* or firewalls (see "Firewalls" subsection).

A host can be connected to more than one subnet. This is actually a security consideration because it is possible to jump from one subnet to another through this host without going through a router or firewall.

Routers can have static (nonchanging) or dynamic routes. A static route is usually defined and loaded by a network manager. Static routes are used to send traffic in a predefined way, for reasons the network designer defines. If static routes are not defined, traffic will flow as the routers determine, based on their configuration. Hackers can affect routing tables and send data to bogus servers instead of to legitimate servers.

## firewalls

A firewall is a special kind of router that can make routing decisions based on the contents of the data packets it sees. Firewalls are used to protect hosts and networks from unauthorized users. They can make routing decisions based on the hardware interface that receives the request, the source or destination IP address, the purpose of the request, and other criteria. Firewalls usually run on dedicated hardware; that is, no other processing should be done on the firewall machine. The less other work the firewall does, the more secure it tends to be.

Firewalls can be hardware boxes from vendors such as SonicWALL or Linksys, or software running on general-purpose operating systems such as Windows or UNIX® systems. Note that a firewall running on an insecure operating system is not secure. About 300 security changes need to be made on a Windows NT system before it is considered secure.

Firewalls can be used to separate networks into intranet, DMZ, extranet, and Internet zones. Requests can flow freely from the intranet out, but only specific requests can flow in other directions. A firewall can also do Network Address Translation (NAT). NAT is used so that the addresses presented to the Internet are not the same as the actual host addresses behind the firewall. This prevents hackers from determining the true structure of your intranet and minimizes hacking attacks of all but the firewall.

## proxy servers

A *proxy server* is like a firewall in that it also examines the packets it sees. Some proxy servers can look deeply at the content inside the packets they process to view the message that the packet is carrying. However, the purpose of a proxy server is to repeat or reroute specified requests. The network designer defines how it responds to those requests. Proxy servers can perform NAT on requests flowing from the intranet to the Internet as well as block or reroute access to some sites. They can reroute requests to a completely different site or allow some pages to be viewed by requests from specific hosts. A proxy server on a DMZ can be used to serve pages from an internal server without opening the internal server to the Internet.

## virtual private networks and tunnels

By building *tunnels* across the Internet to create a *virtual private network (VPN)*, a company can eliminate many leased lines and modem banks. Data flowing across the network is encrypted until it reaches the company's private LAN, where it is decrypted and routed like any other data. The three main kinds of VPNs are

- Intranet VPNs, which extend the intranet with a semipermanent secure WAN connection over a public network to a branch office

- Remote access VPNs, which provide remote access with a temporary secure connection to the enterprise over a public network
- Extranet VPNs, which extend the intranet to reach partners, customers, and suppliers over a public network for e-commerce and e-business

Businesses are attracted to the Internet's ubiquity and low cost as a WAN. The story is compelling:

- Roaming users can reach their home LAN from anywhere in the world.
- WAN management is simplified.
- CFOs win through reduced operations and capital costs.
- Service providers win with a new stream of revenue.
- According to IDC, "A corporation with 1,000 remote users can expect to save US\$1 million per year by moving away from leased lines and dial-up lines and adopting extranet access."

VPNs sound too good to be true, and there is a downside. First, if you choose the wrong vendor or tunnel scheme, your traffic runs the risk of being intercepted and decrypted. Even if it's not decrypted, the intercompany communication itself, and the amount of traffic involved, could provide valuable intelligence to a competitor. For example, suppose a lot of traffic suddenly began to occur from rootbeer.com to lemonlime.com, where no traffic had occurred before. Someone could infer from this pattern that something was happening between the two companies.

Traffic on a VPN is subject to the whims of the underlying network being used, whereas a leased-line (private) network is controlled 100 percent by the company that purchased the lines for their use. Assuming that a VPN is built over the Internet, if the Internet gets overloaded, packets could either take a long time to arrive or be dropped entirely. A VPN constructed in this manner is free, except for the cost of the equipment and software. But you also get what you pay for.

## summary

Internet security involves protecting the host, protecting data on the host, and protecting data flowing to the host. To maintain a secure site, you must authenticate users that access your hosts. If you are using a firewall or proxy server, you should evaluate the logs manually—or frequently if your system evaluates logs automatically—to look for attacks. For every new control, hackers are developing new threats. Your security infrastructure cannot be created, installed, and forgotten. It must be updated as your business changes and as threats change.

## glossary

### advanced encryption standard (AES)

A symmetric block cypher based on the Rijndael (pronounced Rhine-doll) algorithm that takes its name from its Belgian co-creators, Vincent Rijmen and Joan Daemon.

In October 2000, the National Institute of Standards and Technology (NIST) announced the end of a four-year search for a successor to the aging *Data Encryption Standard (DES)* used to protect nonclassified government information and systems. Like DES, AES is free for anyone to use (no license fees). Unlike DES, AES does not have a fixed key

size. This means that the level of protection (and the processing time required to gain that protection) can be linked to how valuable the data is that is being protected. AES requires 128-, 192-, and 256-bit keys.

### **certificate**

An electronic document that states a name or identifies the certification authority (CA), identifies the subscriber, contains the subscriber's public key, identifies the certificate's operational period, contains a certificate serial number, and is digitally signed by the CA.

A certificate is the electronic equivalent of a driver's license or passport. It contains information that uniquely identifies you. Once you identify yourself to the CA, the CA issues you a certificate encrypted under its private key. Because the certificate can be decrypted only under its public key, it is proved to have been issued by the CA.

A certificate is only as good as the CA behind it. If the CA does not require substantial proof of identity for the person or company asking for a certificate, the certificate is not worth much.

### **certification authority (CA)**

Also known as an *issuing authority*, a person or institution that issues and/or verifies certificates. Banks and security vendors may be CAs. Some companies operate solely as CAs, with no additional functions. VeriSign is a well-known CA.

The CA should maintain a secure database of the certificates it has issued and to whom. It does not necessarily generate the key pair. Depending on company policy, this is done by a security officer or, better yet, locally on the user's computer. The public key is then sent to the CA, which issues a certificate for it. The certificate verifies that it was generated by the person or agency indicated. To verify the certificate, the user sends it back to the issuing CA or another authorized CA with a request to verify.

To be taken seriously, a CA must be beyond reproach. If a CA loses its credibility, it loses its customer base. The CA must conduct checks and balances—audit trails, proper identification of requesters, trusted employees at the CA, and so on.

### **data encryption standard (DES)**

A popular symmetric encryption method developed in 1975 and standardized by ANSI in 1981 as ANSI X.3.92. DES uses a 56-bit key. It is the de facto standard for financial transactions. DES is treated like munitions, and its exportability is in flux. You should consult the proper authorities before discussing DES with foreign nationals or exporting hardware or software that uses DES.

### **decryption**

The process of decoding data that has been encrypted into a secret format. Decryption requires a secret key or password.

## **demilitarized zone (DMZ)**

The no-man's land in a network. A firewall is used to create a subnet of servers that are accessible from the Internet yet are protected from hacking. The firewall protects the servers from attacks or changes via the Internet while leaving them relatively unprotected from your intranet so they can be maintained easily. A proxy server on a DMZ can be used to forward specific pages from another server inside your intranet without leaving that system open to attack.

## **denial of service (DOS) attack**

A type of attack designed to bring down a network by flooding it with useless traffic. Many DOS attacks, such as the Ping of Death and TearDrop attacks, exploit limitations in the TCP/IP protocols. System administrators can install software fixes that limit the damage caused by any known DOS attacks. However, hackers are continually dreaming up new DOS attacks, as they do viruses.

Although the purpose of a DOS attack is not to steal information, it can result in lost revenues, profit, market share, and market capitalization, either by making your site appear to be down or by disrupting your business entirely. There have been rumors of adversaries attacking enemies' networks with DOS attacks.

## **Diffie-Hellman (DH) key exchange**

A key exchange protocol that allows participants to agree on a key over an insecure channel using properties of exponents. This was the first definition of a public-key infrastructure (PKI) system.

The DH key exchange is open to *man-in-the-middle attacks*. However, such attacks can be prevented by combining the DH key exchange with a signature algorithm and certificates.

## **encryption**

The translation of data into a secret code. Encryption is the most effective way to achieve data security. To read an encrypted file, you must have access to a secret key or password that enables you to decrypt it. Unencrypted data is called *plaintext*; encrypted data is referred to as *ciphertext*. There are two main types of encryption: asymmetric encryption (also called *public-key encryption*) and symmetric encryption.

## **encryption method**

The algorithm used to convert plaintext to ciphertext using a key. There are hundreds of encryption methods, and each has advantages and disadvantages. Some methods are considered standards. For example, DES is used for banking, *PGP* is used for personal e-mail, and *RSA* is used for Secure Sockets Layer (SSL).

It isn't the method used to encrypt a message that is the secret; rather, it's the key that is secret. It is not difficult to determine what the method is; once it is compromised, changing keys will not make your communications any more secure. Because the method is not secret, it will be exercised to its utmost, and any holes can be detected and fixed. It is better to know about a security hole and be able to fix it than to not

know and have it be cracked. Clients and servers should be able to negotiate which kind of encryption each implements so the strongest methods can be used.

## **firewall**

A system designed to prevent unauthorized access to or from a private network. Firewalls can be implemented in hardware, software, or a combination of the two. Firewalls are frequently used to prevent unauthorized Internet users from accessing private networks connected to the Internet, especially intranets. All messages entering or leaving the intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria.

There are several types of firewall techniques:

- Packet filter: Looks at each packet entering or leaving the network and accepts or rejects it based on user-defined rules. Packet filtering is fairly effective and transparent to users, but it is difficult to configure. In addition, it is susceptible to *IP spoofing*.
- Application gateway: Applies security mechanisms to specific applications, such as FTP and Telnet servers. This is very effective but can impose a performance degradation.
- Circuit-level gateway: Applies security mechanisms when a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) connection is established. Once the connection has been made, packets can flow between the hosts without further checking.
- Proxy server: Intercepts all messages entering and leaving the network. The proxy server effectively hides the true network addresses.

## **hypertext markup language (HTML)**

The authoring language used to create documents on the Web. HTML was originally created as a text formatting language, similar to T<sub>G</sub>AL or T<sub>F</sub>ORM, two languages used with NonStop systems. It was the implementation of HTML in a Web browser by Marc Andreessen around 1994 that led to the explosion of the World Wide Web. Since then, different browsers have expanded on HTML independently to gain market share.

Although there is a standard for HTML, neither Netscape nor Microsoft follows it completely. Some Web pages are written for one or the other browser, whereas many pages don't use any advanced features at all, so they can work with both browsers.

HTML is sent from the server to the browser in source code form, so knowledgeable Internet users can see how a Web page is written. This can lead to security breaches, because hackers can discern the host and network structure from the source code.

## **hypertext transfer protocol (HTTP)**

The underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what actions Web servers and browsers should take in response to various commands. For example, when you enter a URL in your browser, an HTTP command is sent to the Web server directing it to fetch and transmit the requested Web page.

HTTP is called a stateless protocol because each command is executed independently, without any knowledge of the commands that came before it (much like an HP Pathway server environment). This is the main reason it is difficult to implement websites that

react intelligently to user input. This shortcoming of HTTP is being addressed by new technologies—including Microsoft ActiveX, Java™, and JavaScript tools—and the use of cookies.

### **the internet corporation for assigned names and numbers (ICANN)**

A global nonprofit corporation formed to oversee a select range of Internet technical management functions currently conducted by the U.S. government, or by its contractors and volunteers. ICANN is gradually taking over responsibility for coordinating the assignment of protocol parameters, the management of the domain name and root server systems, and the allocation of IP address space.

### **InterNIC**

A concept for an integrated network information center developed by several companies, including Network Solutions, in cooperation with the U.S. government.

### **IP spoofing**

A technique used to gain unauthorized access to computers, whereby the intruder sends messages to a computer with an IP address indicating that the message is coming from a trusted port. To engage in IP spoofing, a hacker must first use a variety of techniques to find an IP address of a trusted port and then modify the packet headers so it appears the packets are coming from that port.

IP spoofing is one of the hardest attacks to defend against, because the server perceives that the request is coming from a valid address. One way to minimize spoofing from outside of the intranet is to use a proxy server firewall; such a device can hide valid intranet addresses from the Internet.

Spoofing from the intranet is much harder to detect and prevent. Extensive routing tables and internal firewalls can be used to separate a secure intranet (for example, the accounting department) from an insecure intranet. (See also *IPSec*.)

### **IPSec, or secure internet protocol**

A set of protocols being developed by the Internet Engineering Task Force (IETF) to support secure exchange of packets at the IP layer. Once it is complete, IPSec is expected to be deployed widely to implement virtual private networks (VPNs).

IPSec supports two encryption modes: transport and tunnel. Transport mode encrypts only the data portion (payload) of each packet but leaves the header untouched. The more secure tunnel mode encrypts both the header and the payload. On the receiving side, an IPSec-compliant device decrypts each packet. IPSec may also be a solution to man-in-the-middle attacks.

### **key pair**

A pair of related keys used in public-key encryption. The key pair consists of a public key and a private key. Because there are two keys and only the public key is ever transmitted, it is highly unlikely that the private key will be compromised if a hacker taps the transmission line.

Even though the private key is never transmitted, some implementations could keep the private key in the clear either on a disk or in memory where it can be discovered. One of the benefits of using HP products (SignMaster, PayMaster/PCI, or WebSafe2/PCI Internet security processors) is that the private key is never seen outside of the box.

### **man-in-the-middle attack**

A method of attacking networks that includes interception, insertion, deletion, and modification of messages; reflecting messages back at the sender; replaying old messages; and redirecting messages.

This type of attack is hard to guard against and can be used to reroute packets to a bogus server (pretending to be the name server), steal information, alter funds transfers, or make it appear that a server is down. In classic man-in-the-middle attacks, hackers intercept the first message that supplies a public key and supplant it with their own. They can then alter subsequent messages before re-encrypting them under the recipient's key and sending them on. IPSec has some features that can minimize man-in-the-middle attacks.

Future versions of TCP/IP also have features that limit man-in-the-middle attacks. Such features include serial numbered packets and information that can trace a packet back to the host that sent it. These involve a tradeoff between security and privacy, however.

### **octet**

Eight contiguous bits, also called a byte. An octet is usually displayed in decimal form instead of a byte's octal (base 8) or hexadecimal (base 16) display. A 32-bit TCP/IP address is almost always displayed as four decimal octets, each of which can vary from 0 to 255.

### **PGP, or "pretty good privacy"**

A technique developed by Philip Zimmerman for encrypting messages. PGP is one of the most common ways to protect messages on the Internet because it is effective, easy to use, and free. PGP is based on the public-key method, which uses two keys: One is a public key that you disseminate to anyone from whom you want to receive a message; the other is a private key that you use to decrypt messages that you receive.

The PGP encryption package is available at no charge from numerous sources for almost any operating system. The official repository is at the Massachusetts Institute of Technology.

PGP is such an effective encryption tool that the U.S. government actually brought a lawsuit against Zimmerman for putting it in the public domain and hence making it available to enemies of the United States. After a public outcry, the U.S. lawsuit was dropped, but it is still illegal to use PGP in many countries.

### **proxy server**

A server that sits between a client application, such as a Web browser, and a real server. It intercepts all requests to the real server to see if it can fulfill the requests itself. If not, it forwards the request to the real server.

A proxy server has two main purposes. It can dramatically improve performance for groups of users, because it saves the results of all requests for a specific amount of time. Consequently, it doesn't need to contact the real server for the second and subsequent requests. It can also be used to filter requests. For example, a company might use a proxy server to prevent its employees from accessing some websites.

In addition, proxy servers have an incidental function. They can be used to serve pages from inside a company network through a firewall without compromising the security of the network.

### **public-key encryption (PKE)**

An asymmetrical encryption method that uses two keys—a public key and a private key. Plaintext encrypted using one key can be decrypted only by using the other.

Because one of the keys is public, you do not need to exchange an encryption key with the person to whom you send the encrypted message. This means that you do not need to worry that a third party can see the key you've chosen and encrypt subsequent messages.

Note that if you encrypt using your private key (which proves your identity), anyone can decrypt the message with your public key. If you encrypt a message under the recipient's public key, only the recipient can decrypt the message (proving his or her identity). If you want to ensure the identity of the sender and the recipient, you can encrypt under your private key and the recipient's public key. The recipient will then decrypt under his or her private key and your public key.

PKE is a processor-intensive method of encryption. Rather than encrypting everything using PKE, you can send a *session key* encrypted using PKE and send subsequent messages under the session key. See also *Secure Sockets Layer (SSL)*.

### **public-key infrastructure (PKI)**

The infrastructure that supports PKE. This includes PKE encryption methods, CAs, and certificates. Without all of the pieces in place, PKE can only be used between two people who already trust each other and who have exchanged public keys in a manner that ensures each other's identity. PKI can authenticate the identities of the two public key users.

### **router**

A device that connects any number of subnets. Routers use headers and a forwarding table to determine where packets go. They also use Internet Control Message Protocol (ICMP) to communicate with each other and configure the best route between any two hosts.

Routers can act as basic firewalls by looking at the packets that pass through them. Data can be hijacked by hackers attacking routing tables.

## **RSA**

A PKI cryptosystem named for its inventors (Rivest, Shamir, and Adleman). The RSA cryptosystem is the most popular form of public-key cryptography. RSA Security is one of the leading PKI vendors.

## **secure electronic transaction (SET)**

A standard that has been endorsed by virtually all of the major players in the electronic commerce arena, including Microsoft, Netscape, Visa, and MasterCard. By employing digital signatures, SET enables merchants to verify the identity of buyers. And it protects buyers by providing a mechanism for their credit card number to be transferred directly to the credit card issuer for verification and billing without the merchant being able to see the number.

SET was overtaken by SSL for Web transactions. Although SET is a secure end-to-end protection scheme for all transactions related to e-business, it was a licensed architecture. In addition, no one wanted to pay the license fees or install the special software required.

SET is in the process of dying out and will probably be replaced by the Visa Open Platform for smart card transactions.

## **secure hypertext transfer protocol (S-HTTP)**

An extension to the HTTP protocol to support sending data securely over the Web. S-HTTP was developed by Enterprise Integration Technologies (EIT), which was acquired by Verifone, Inc., in 1995.

Not all Web browsers and servers support S-HTTP. SSL, another technology for transmitting secure communications over the Internet, is more prevalent. However, SSL and S-HTTP have different designs and goals, so it is possible to use the two protocols together. Whereas SSL is designed to establish a secure connection between two computers, S-HTTP is designed to send individual messages securely.

## **secure/multipurpose internet mail extension (S/MIME)**

A version of the MIME protocol that supports encryption of messages. S/MIME is based on RSA Security's PKE technology. MIME is a specification for formatting non-ASCII messages so they can be sent over the Internet. S/MIME is expected to be widely implemented, which will make it possible for people to send secure e-mail messages to one another, even if they are using different e-mail clients.

## **secure sockets layer (SSL)**

A communications protocol that uses PKI to create a session key for use between a client (browser) and server. The key is created by the client and encrypted under the server's public key. The server then decrypts it using its private key, and they both use the key during the transaction. Once it is enabled, SSL is protocol independent; therefore, HTTP, FTP, and Telnet can be used securely. Major browsers implement this protocol transparently to the user.

SSL was created by Netscape and is supported by both major browsers. It is currently at version 3; each version fixes holes found in the previous versions. When SSL is invoked, the client and server can negotiate the level they are both capable of using.

By convention, Web pages that require an SSL connection start with *https*: instead of *http*:. (See also *S-HTTP*.) SSL can also verify the client using certificates. Currently, this function is not widely implemented, which means that only the server is verified.

### **session key**

A key for symmetric cryptosystems that is used for the duration of one message or communication session—for example, the messages needed to place and pay for an order on an e-commerce site. See also *Secure Sockets Layer (SSL)*.

PKE is a processor-intensive method of encryption. Symmetric encryption using a session key is much less so. The session key is used in both directions, to encrypt and decrypt (hence symmetric).

Once you know the key, it is easy to intercept all subsequent messages. Therefore, session keys should be used only for a short period of time.

### **smart card**

A small electronic device about the size of a credit card that contains electronic memory and sometimes an embedded integrated circuit (IC). Smart cards containing an IC are sometimes called integrated circuit cards (ICCs). Smart cards are used for a variety of purposes, including storing a patient's medical records, storing digital cash, and generating network IDs (similar to a token). To use a smart card, either to pull information from it or add data to it, you need a smart card reader, a small device into which you insert the smart card.

Authentication involves three factors: who you are, what you have, and what you know. Unlike a password, which is based only on what you know, a smart card is more secure because it is based on two of these three factors. Although you can reproduce a credit card, including its magnetic stripe, and steal or guess the PIN, a smart card is much harder to replicate.

### **sniffer**

A program and/or device that monitors data traveling over a network. Sniffers can be used both for legitimate network management functions and for stealing information from a network. On TCP/IP networks, they're often called *packet sniffers*.

Unauthorized sniffers can be extremely dangerous to a network's security because they are virtually impossible to detect and can be inserted almost anywhere on the network. This makes them a favorite weapon in the hacker's arsenal.

### **subnet**

A portion of a network that shares a common address component. On TCP/IP networks, subnets are defined as all devices whose IP addresses have the same prefix. For example, all devices with IP addresses that start with "100.100.100" would be part of

the same subnet. Dividing a network into subnets is useful both for security and for performance. IP networks are divided using a subnet mask.

Subnets can be used to separate hosts by traffic, function, security domains, and other criteria. Firewalls and routers can be used to link subnets in a secure manner.

### **subnet mask**

A mask used to determine which subnet an IP address belongs to. An IP address has two components, the network address and the host address.

Consider the IP address 150.215.017.009, for example, which consists of four octets. Assuming there are only 255 hosts on this subnet, the first three octets (150.215.017) represent the network address. The last octet, which can vary from 0 to 255, identifies a particular host on this network. The subnet mask is the network address (in this case, three octets) plus the bits reserved for identifying the subnet (in this case, one octet). By convention, the bits for the network address portion are all set to 1, although it would also work if the bits were set exactly as in the network address and the bits for the host portion were all set to 0. For this example, the subnet mask would be 255.255.255.0.

It's called a mask because it can be used to identify the subnet to which an IP address belongs by performing a binary logical AND operation on the mask and the IP address. The important fact is that the subnet mask is used so routers know when to send packets to another subnet instead of keeping them on the same subnet.

### **transmission control protocol/internet protocol (TCP/IP)**

The suite of communications protocols used to connect hosts on the Internet. Each host must have a unique address if other hosts are to reach it.

Because the number of addresses is not infinite, some address ranges have been declared as nonroutable. These addresses cannot appear on the Internet but must only be used in private networks. Proxy servers can be used to hide the nonroutable addresses of these hosts.

### **Triple-DES**

The process by which plaintext is run through the DES encryption method three times using an additional key: encrypt under key 1, decrypt under key 2, and encrypt again under key 1. When it was introduced, DES, with its 56-bit key, was considered secure, but the increase in speed of affordable computing power has rendered it much less secure against brute-force attacks (trying all key combinations to crack a message). Triple-DES has a 112-bit key, which gives it significant protection against such attacks.

### **tunneling**

A technology that enables one network to send its data via another network's connections. Tunneling works by encapsulating a network protocol within packets carried by the second network. For example, Point-to-Point Tunneling Protocol (PPTP), Layer 2 Tunneling Protocol (L2TP), or IPSec technology enables organizations to use the Internet to transmit data across a virtual private network. It does this by embedding its own network protocol within the TCP/IP packets carried by the Internet. Tunnels can be built between two networks, a network and a client, or two clients.

Companies can let employees who are working at home connect via the employee's Internet service provider to a company tunnel server. This eliminates the need for expensive leased lines or modem banks. Packets from the employee's computer are encrypted and sent over the Internet to the company's tunnel server, where they are decrypted and put on the corporate LAN. Similarly, a company can securely exchange information with its suppliers or customers using tunnels.

There are secure and insecure tunneling methods. And vendor equipment isn't always compatible. Before installing tunnel servers and software, companies should plan well.

### **uniform resource locator (URL)**

The global address of documents and other resources on the Web. The first part of the address indicates what protocol to use, and the second part specifies the IP address or the domain name where the resource is located. For example, the two URLs below point to two different files at the domain "pcwebopedia.com." The first specifies an executable file that should be fetched using FTP; the second specifies a Web page that should be fetched using HTTP:

- `ftp://www.pcwebopedia.com/stuff.exe`
- `http://www.pcwebopedia.com/index.html`

Allowing FTP access to a website (HTTP) could enable hackers to determine your system and network structure and break into it more easily.

### **virtual private network (VPN)**

A temporary, secure connection over a public network that provides privacy, integrity, authentication, and controlled access between two points. Numerous systems enable you to create networks using the Internet as the medium for transporting data. These systems use encryption and other security mechanisms to ensure that only authorized users can access the network and that the data cannot be intercepted.

Traffic on a VPN is subject to the whims of the underlying network being used, whereas a leased-line network (which is private) is entirely controlled by the company that purchased the lines for its use. Assuming that a VPN is built over the Internet, if the Internet gets overloaded, packets could either take a long time to arrive or they could be dropped entirely. A VPN constructed in this manner is free, except for the cost of the equipment and software. But you also get what you pay for.

For more information, go to [www.hp.com/go/nonstop](http://www.hp.com/go/nonstop).

September 2002, first published May 2001. Java is a U.S. trademark of Sun Microsystems, Inc. Microsoft, Windows, and Windows NT are U.S. registered trademarks of Microsoft Corporation. Netscape is a U.S. trademark of Netscape Communications Corporation. UNIX is a registered trademark of The Open Group. All other product names mentioned herein may be trademarks of their respective companies. HP shall not be liable for technical or editorial errors or omissions contained herein. The information is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

02-0403

©2002 Hewlett-Packard Company

