Information Security Management Handbook
Edited by Harold F. Tipton and Micki Krause
Boca Raton: CRC Press LLC, 2003

Firewalls, Ten Percent of the Solution: A Security Architecture Primer

Chris Hare, CISSP, CISA

A solid security infrastructure consists of many components that, through proper application, can reduce the risk of information loss to the enterprise. This chapter examines the components of an information security architecture and why all the technology is required in today's enterprise.

A principal responsibility of the management team in any organization is the protection of enterprise assets. First and foremost, the organization must commit to securing and protecting its intellectual property. This intellectual property provides the organization's competitive advantage. When an enterprise loses that competitive advantage, it loses its reason for being an enterprise.

Second, management must make decisions about what its intellectual property is, who it wants to protect this property from, and why. These decisions form the basis for a series of security policies to fulfill the organization's information protection needs.

However, writing the policies is only part of the solution. In addition to developing the technical capability of implementing these policies, the organization must remain committed to these policies, and include regular security audits and other enforcement components into its operating plan. This is similar to installing a smoke alarm: if you do not check the batteries, how will you know it will work when you need it?

There are many reasons why a corporation should be interested in developing a security architecture. These include:

- Telecommunications fraud
- Internet hacking
- Viruses and malicious code
- War dialing and modem hacking
- Need for enhanced communications
- Globalization
- Cyber-terrorism
- Corporate espionage
- E-commerce and transaction-based Web sites

Telecommunications fraud and Internet and modem hacking are still at the top of the list for external methods of attacking an organization. Sources of attack are becoming more sophisticated and know no geographical limits. Consequently, global attacks are more predominant due to the increased growth in Internet connectivity and usage.

With business growth has come the need for enhanced communications. No longer is remote dial-up sufficient. Employees want and need high-speed Internet access and other forms of services to get their jobs done, including videoconferencing, multimedia services, and voice conferencing. Complicating the problem is that many corporate networks span the globe, and provide a highly feature-rich, highly connected environment for both their employees and for hackers.

The changes in network requirements and services has meant that corporations are more dependent on technologies that are easily intercepted, such as e-mail, audio conferencing, videoconferencing, cellular phones, remote access, and telecommuting. Employees want to access their e-mail and corporate resources through wireless devices, including their computers, cell phones, and personal digital assistants such as the PalmPilot and Research in Motion (RIM) BlackBerry.

With the Information Age, more and more of the corporation's knowledge and intellectual capital are being stored electronically. Information technology is even reported as an asset on the corporation's financial statements. Without the established and developed intellectual capital, which is often the distinguishing factor between competitors, the competitive advantage may be lost.

Unfortunately, the legal mechanisms are having difficulty dealing with this transnational problem, which affects the effectiveness and value of the legislation — expertise of law enforcement, investigators, and prosecutors alike. This legal ineffectiveness means that companies must be more diligent at protecting themselves because these legal deficiencies limit effective protection.

Add to this legal problem the often limited training and education investment made to maintain corporate security and investigative personnel in the

legal and information technology areas. Frequently, the ability of the hacker far surpasses the ability of the investigator.

Considering the knowledge and operational advantages that a technology infrastructure provides, the answer is this: the corporation requires a security infrastructure because the business needs one.

Over the past 15 years, industry has experienced significant changes in the business environment. Organizations of all sizes are establishing and building new markets. Globalization has meant expanding corporate and public networks and computing facilities to support marketing, sales, and support staff. In addition to the geographical and time barriers, enterprises are continually faced with cultural, legal, language, and ethical issues never before considered.

In this time frame, we have also seen a drive toward electronic exchange of information with suppliers and customers, with E-commerce and transaction-based Web sites being the growth leader in this area.

This very competitive environment has forced the enterprise to seek efficiencies to drive down product costs. The result of this activity has been to outsource non-core activities, legacy systems, consolidation of workforces, and a reduction in non-essential programs.

The mobile user community reflects the desire to get closer to our customer for improved responsiveness (e.g., automated salesforce). In addition, legislation and the high cost of real estate have played a role in providing employees with the ability to work from home.

The result of these trends is that information is no longer controlled within the confines of the data center, thereby making it easier to get access to, and less likely that this access would be noticed.

WHERE ARE THE RISKS?

The fact is that firewalls provide the perimeter security needed by today's organizations. However, left on their own, they provide little more than false assurance that the enterprise is protected. Indeed, many organizations believe the existence of a firewall at their perimeter is sufficient protection. It is not!

The number of risks in today's environment grows daily. There have been recent documented instances in which members of some of these areas, such as outsourced consultants, have demonstrated that they are more at risk than some organizations are prepared to handle. For example, *Information Week* has reported cases where outsourced consultants have injected viruses into the corporate network. A few of the many risks in today's environment include:

- Inter-enterprise networking with business partners and customers
- Outsourcing
- Development partners
- Globalization
- Open systems
- Access to business information
- Research and development activities
- Industrial and economic espionage
- Labor unrest
- Hacking
- Malicious code
- Inadvertent release or destruction of information
- Fraud

These are but a few of the risks to the enterprise the security architecture must contend with. Once the organization recognizes that the risk comes from both internal and external sources, the corporation can exert its forces into the development of technologies to protect its intellectual property.

As one legitimate user community after another have been added to the network, it is necessary to identify who can see what and provide a method of doing it. Most enterprises have taken measures to address many of the external exposures, such as hacking and inadvertent leaks, but the internal exposures, such as industrial or economic espionage, are far more complex to deal with. If a competitor really wants to obtain valuable information, it is easier and far more effective to plant someone in the organization or engage a business partner who knows where the information can be found.

Consider this: the U.S. FBI estimates that one out of every 700 employees is actively working against the company.

ESTABLISHING THE SECURITY ARCHITECTURE

The architecture of the security infrastructure must be aligned with the enterprise security policy. If there is no security policy, there can be no security infrastructure. As security professionals, we can lead the best technologists to build the best and most secure infrastructure; however, if it fails to meet the business goals and objectives, we have failed. We are, after all, here to serve the interests of the enterprise — not the other way around.

The security architecture and resulting technology implementation must, at the very least, meet the following objectives:

- It must not impede the flow of authorized information or adversely affect user productivity.
- It must protect information at the point of entry into the enterprise.

- It must protect the information throughout its useful life.
- It must enforce common processes and practices throughout the enterprise.
- It must be modular to allow new technologies to replace existing ones with as little impact as possible.

Enterprises and their employees often see security as a business impediment. Consequently, they are circumvented in due course. For security measures to work effectively, they must be built into operating procedures and practices in such a way that they do not represent an "extra effort." From personal experience, this author has seen people spend up to ten times the effort and expense to avoid implementing security.

The moment the security infrastructure and technology are seen, *or perceived*, to impact information flow, system functionality, or efficiencies, they will be questioned and there will be those who will seek ways to avoid the process in the interest of saving time or effort. Consequently, the infrastructure must be effective, yet virtually transparent to the user.

Once data has entered the system, it must be assumed that it may be input to one or more processes. It is becoming impractical to control the use of all data elements at the system layer; therefore, any data that is considered sensitive, or can only be "seen" by a particular user community, must be appropriately protected at the point of entry to the network or system and, most importantly, wherever it is subsequently transferred. This involves the integration of security controls at all levels of the environment: the network, the system, the database, and the application.

A centralized security administration system facilitates numerous benefits, both in terms of efficiency and consistency. Perhaps the most significant advantage is knowing who has access to what and if, for whatever reason, access privileges are to be withdrawn, that can be accomplished for all systems expeditiously.

Quite clearly, it is not economically feasible to rewrite existing applications or replace existing systems. Therefore, an important aspect of the security architecture must be the ability to accommodate the existing infrastructure. Along the same lines of thinking, the size of existing systems and the population using them precludes a one-time deployment plan. A modular approach is an operational necessity.

The infrastructure resulting from the architecture must also provide specific services and meet additional objectives, including:

- Access controls
- Authorization
- Information classification
- Data integrity

Achieving these goals is not only desirable, it is possible with the technology that exists today. It is highly desirable to have one global user authentication and authorization system or process, a single encryption tool, and digital signature methodology that can be used consistently across the enterprise for all applications. Authenticating the user does not necessarily address the authorization criteria; it may prove that you are who you say you are but does not dictate what information can be accessed and what can be done with it.

Given the inter-enterprise electronic information exchange trend, one can no longer be certain that the data entering the corporate systems is properly protected and stored at the points of creation. Data that is submitted from unsecured areas represents a number of problems, primarily related to integrity, the potential for information to be modified (e.g., the possibility of the terminal device being "spoofed," collecting data, modifying it, and retransmitting it as if from the original device), and confidentiality (e.g., "shoulder surfing").

Unfortunately, one cannot ignore the impact of government in our infrastructure. In some way or another, domestic and foreign policies regarding what one can and cannot use do have an effect. Consider one of the major issues today being the use of encryption. The United States limits the export of encryption to a key length, whereas other governments (e.g., France) have strict rules regarding the use of encryption and when they require a copy of the encryption key.

In addition, governments also impose import and export restrictions on corporations to control the movement of technology to and from foreign countries. These import/export regulations are often difficult to deal with due to the generalities in the language the government uses, but they cannot be ignored. Doing so may result in the corporation not being able to trade with some countries, or lose its ability to operate.

AN INFRASTRUCTURE MODEL

The security infrastructure must be concerned with all aspects of the information, and the technology used to create and access it. This includes:

- Physical security for the enterprise and security devices
- Monitoring tools
- Public network connectivity
- Perimeter access controls
- Enterprise WAN and LAN
- Operating systems
- Applications
- Databases
- Data

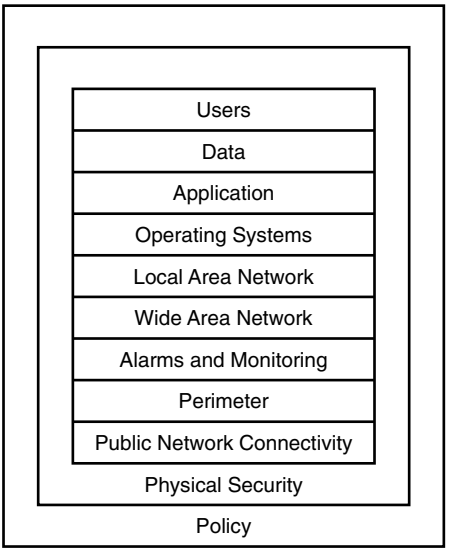

Exhibit 39-1. The infrastructure model.

This also does not discount the need for proper policies and an awareness program as discussed earlier. The protection objects listed above, if viewed in a reverse order (see Exhibit 39-1), provides an outside in view to protecting the data.

What this model also does is incorporate the elements of physical security and awareness, including user training, which are often overlooked. Without the user community understanding what is expected from them in the security model, it will be difficult — if not impossible — to maintain.

The remainder of this chapter focuses on the technology components and how to bring them together in a sample architecture model.

ESTABLISHING THE PERIMETER

The 1980s brought the development of the microcomputer, and despite its cost, many enterprises that were mainframe oriented could now push the work throughout the enterprise on these lower-cost devices. Decentralization of the computing infrastructure brought several benefits and, consequently, several challenges.

As connectivity to the Internet increased, a new security model was developed. This consisted of a "moat," where the installation of a firewall provided protection against unauthorized access. Many organizations then, as today, took the approach that information contained within the network was available for any authorized employee to access. However,

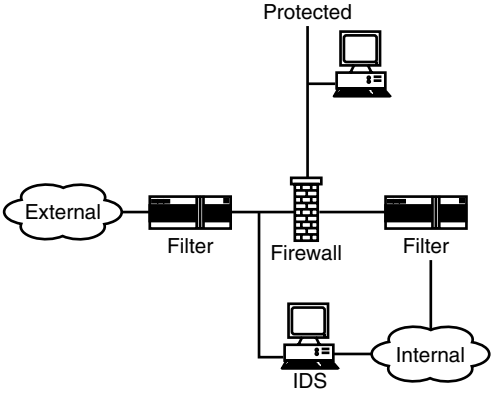

Exhibit 39-2. Perimeter access point.

this open approach meant that the enterprise was dependent upon other technology such as network encryption devices to protect the information and infrastructure.

The consequence many organizations have witnessed with this model is that few internal applications and services made any attempt to operate in a secure fashion. As the number of external organizations connected to the enterprise network increases, the likelihood of the loss of intellectual property also increases.

With the knowledge that the corporate network and intellectual property were at risk, it was evident that a new infrastructure was required to address the external access and internal information security requirements.

Security professionals around the globe have embarked on new technology and combinations. Consequently, it is not uncommon for the network perimeter to include:

- Screening or filter routers
- Firewalls
- Protected external networks
- Intrusion detection systems

When assembled, the perimeter access point resembles the diagram in Exhibit 39-2.

The role of the screening or filter router between the external network and the firewall is to limit the types of traffic allowed through, thereby reducing the quantity of network traffic visible to the firewall. This establishes the first line of defense. The firewall can then respond more effectively to the traffic that is allowed through by the filter router. This first filter router

performs the ingress traffic filtering, meaning it limits the traffic inbound to your network based on the filter rules.

Traditionally, enterprises have placed their external systems such as Web and FTP servers outside their firewall, which is typically known as the DMZ (demilitarized zone). However, placing the systems in this manner exposes them to attack from the external network. An improved approach is to add additional networks to the firewall for these external systems. Doing so creates a protected network, commonly known as a service network or screened subnet.

The filters on the external filter router should be written to allow external connections to systems in the protected network, but only on the allowed service ports. For example, if there is a Web server in the protected network, the filter router can be designed to send all external connection requests to the Web server to only the Web server. This prevents any connections into the internal network due to an error on the firewall.

> *Note:* The overuse of filters on routers can impact the overall performance of the device, increasing the time it takes to move a packet from one network to another. For example, adding a single rule: <any IP address> to <any IP address> adds ten percent to the processing load on the router CPU. Consequently, router filter rules, while recommended, must be carefully engineered to not impede network performance.

The firewall is used to create the screened or protected subnet. A screened subnet allows traffic from the external network into the screened subnet, but not directly into the corporate network. Additionally, firewall rules are also used to further limit the types of traffic allowed into the screened subnet, or into the internal network.

Should a system in the protected network require access into the internal network, the firewall provides the rules to do so, and limits the protocols or services available into the internal network.

The second filter router between the firewall and the internal network is used to limit outbound traffic to the external network. This is particularly important to prevent network auto-discovery systems such as HP Openview from trying to use its auto-discovery features to map the entire Internet. This filter router can also allow other traffic that the enterprise does not want sent out to the Internet to be blocked. This is egress filtering, or using the router to limit the traffic types being sent to the external network. Some enterprises combine both filters on one router, which is acceptable depending on the ultimate architecture implemented.

The final component is an intrusion detection system (IDS) to identify connection attempts or other unauthorized events and information. Additionally, content filtering systems can be used to scan for undesirable content in various protocols such as Web and e-mail. Many vendors offer solutions

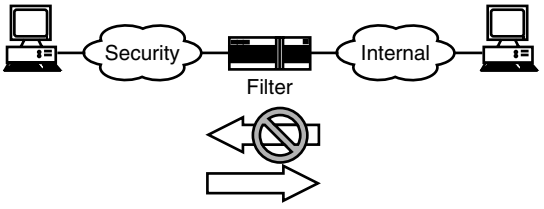

Exhibit 39-3. Local area network with security domains.

for both, including those that can prevent the distribution of specific types of attachments in e-mail messages. E-mail attachment scanning should also be implemented in the enterprise to prevent the distribution of attachments such as malicious code within the enterprise.

THE NETWORK LAYER

The network layer addresses connectivity between one user, or system, and another for the purposes of information exchange. In this context, information may be in the form of data, image, or sound and may be transmitted using copper, fiber, or wireless technologies. This layer will include specific measures to address intra- and inter-enterprise information containment controls, the use of private or public services, protocols, etc.

Almost all enterprises will have some level of connectivity with a public data network, be it the Internet or other value-added networks. The security professional must not forget to examine all network access points and connectivity with the external network points and determine what level of protection is needed. At the very least, a screening router must be used. However, in some cases, external legislation determines what network access control devices are used and where they must be located.

The enterprise wide area network (WAN) is used to provide communications between offices and enterprise sites. Few enterprises actually maintain the WAN using a leased line approach due to the sheer cost of the service and associated management. Typically, WAN services are utilized through public ATM or Frame Relay networks. While these are operated and managed by the public telecommunications providers, the connectivity is private due to the nature of the ATM and Frame Relay services.

Finally, the local area network (LAN) used within each office provides network connectivity to each desktop and workstation within the enterprise. Each office or LAN can be used to segregate users and departments through security domains (see Exhibit 39-3).

In this case, the security professional works with the network engineering teams to provide the best location for firewalls and other network

access devices such as additional filter routers. Utilizing this approach can prevent sensitive traffic from traveling throughout the network and only be visible to the users who require it. Additionally, if the information in the security domain requires it, network and host-based IDSs should be used to track and investigate events in this domain.

Finally, the security professional should recommend the use of a switched network if a shared media such as coaxial or twisted-pair media is used. Traditional shared media networks allow any system on the network to see all network traffic. This makes it very easy for a sniffer to be placed on the network and packets collected, including password and sensitive application data. Use of a switched network makes it much more difficult, although not impossible.

Other controls should be used in the design of the LAN. If the enterprise is using DHCP, any person who connects to the LAN and obtains an IP address can gain access to the enterprise network. For large enterprises, it is unrealistic to attempt to implement MAC-level controls due to the size of the network. However, public areas such as lobbies and conference rooms should be set up in one of the following manners:

- No live network jacks
- DHCP on a separate subnet and security domain
- Filtered traffic

The intent of these controls is to prevent a computer in a conference room from being able to participate fully on the network, and only offer limited services. In this context, security domains can be configured to specifically prevent access to other parts of the network or specific systems based on the source IP address.

Other LAN-based controls for network analysis and reporting, such as Nicksun Probe and NetVCR, provide network diagnostics, investigation, and forensics information. However, on large, busy networks, these provide an additional challenge, that being the disk space to store the information for later analysis.

Each of the foregoing layers provides the capability to monitor activities within that layer. Monitoring systems will be capable of collecting information from one or more layers, which will trigger alarm mechanisms when certain undesirable operational or security criteria are met. The alarm and monitoring tools layer will include such things as event logging, system usage, exception reporting, and clock synchronization.

PHYSICAL SECURITY

Physical security pertains to all practices, procedures, and measures relating to the operating environment, the movement of people, equipment

or goods, building access, wiring, system hardware, etc. Physical security elements are used to ensure that the corporate assets are not subjected to unwarranted security risks. Items addressed at this layer include secure areas, security of equipment off-premises, movement of equipment, and secure disposal of equipment.

The physical security of the network access control devices, including the

- Firewall
- IDS
- Filter routers
- Hubs
- Switches
- Cabling
- Security systems

is paramount to ensuring the ongoing protection of the network and enterprise data. Should these systems not be adequately protected, a device could be installed and no one would notice. Physical security controls for these devices should include locked cabinets and cable conduits, to name only two.

SYSTEM CONTROLS

Beyond the network are the systems and applications that users use on a daily basis to fulfill enterprise business objectives. The protection of the operating system, the application proper, and the data are just as important as the network.

Fundamentally, information security is in the hands of the users. Regardless of the measures that may be implemented, carelessness on the part of individuals involved in the preparation, consolidation, processing, recording, or movement of information can compromise any or all security measures. This layer then looks at the human-related processes, procedures, and knowledge related to developing a secure environment, such as user training, information security training and awareness, and security policies and procedures.

Access to the environment must be controlled through a coordinated access control program, as discussed later in this chapter. Access control provides the control mechanisms to limit access to systems, applications, data, or services to authorized people or systems. It includes, for example, identification of the user, their authorization, and security practices and procedures. Examples of items that would be included in access control systems include identification and authentication methods, privilege management, and user registration. One could argue that privilege management is part of authorization; however, it should be closely coupled to the authentication system.

The operating system controls provide the functionality for applications to be executed and management of system peripheral units, including connectivity to network facilities. A heterogeneous computing environment cannot be considered homogeneous from a security perspective because each manufacturer has addressed the various security issues in a different manner. However, within your architecture, the security professional should establish consistent operating system baselines and configurations to maintain the overall environment.

Just as the security professional will likely install a network-based intrusion detection system, so too should host-based systems be considered for the enterprise's critical systems and data. Adding the host-based element provides the security professional with the ability to monitor for specific events on the system itself that may not be monitored by or captured through a network-based intrusion detection system.

The data aspect of the architecture addresses the measures taken to ensure data origination authenticity, integrity, availability, non-repudiation, and confidentiality. This layer will address such things as database management, data movement and storage, backup and recovery, and encryption. Depending on the applications in use, a lot of data is moved between applications. These data transfers, or interfaces, must be developed appropriately to ensure that there is little possibility for data compromise or loss while in transit.

The application and services layer addresses the controls required to ensure the proper management of information processing, including inputs and outputs, and the provision of published information exchange services.

ESTABLISHING THE PROGRAM

The security architecture must not only include the elements discussed so far, but also extend into all areas to provide an infrastructure providing protection from the perimeter to the data. This is accomplished by linking security application and components in a tightly integrated structure to implement a security control infrastructure (see Exhibit 39-4).

The security control infrastructure includes security tools and processes that sit between the application and the network. The security control infrastructure augments or, ideally, replaces some of the control features in the applications — mostly user authentication. This means that the application does not maintain its own view of authentication, but relies on the security control infrastructure to perform the authentication. The result is that the user can authenticate once, and let the security control infrastructure take over. This allows for the eventual implementation of a single sign-on capability.

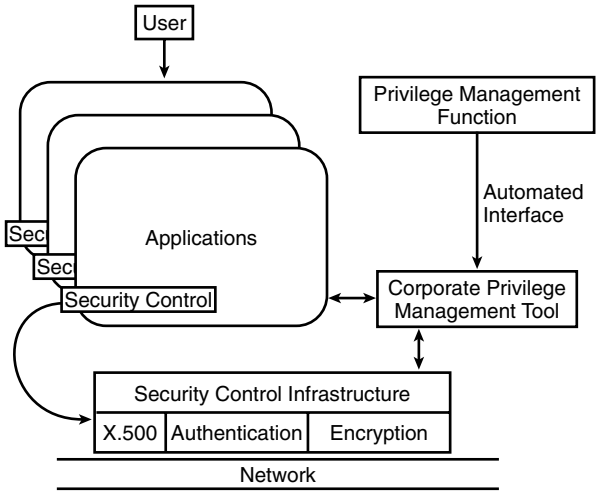

Exhibit 39-4. Security control infrastructure.

A centralized tool for the management of individual user and process privileges is required to enable the security control infrastructure to achieve this goal. The centralized user management services interact with the control infrastructure to determine what the user is allowed to do. Control infrastructure and other services within it depend on the existence of an enterprisewide privilege database containing the access and application rights for every user.

The result is a security infrastructure that has the ability to deliver encryption, strong authentication, and a corporate directory with the ability to add single sign-on and advanced privilege management in the future.

THE CORPORATE DIRECTORY

The corporate directory, which is a component of the security control infrastructure, contains elements such as:

- Employee number, name, department, and other contact information
- Organizational information such as the employee's manager and reporting structure
- Systems assigned to the employee
- User account data
- E-mail addresses
- Authorized application access
- Application privileges
- Authentication information, including method, passwords, and access history
- Encryption keys

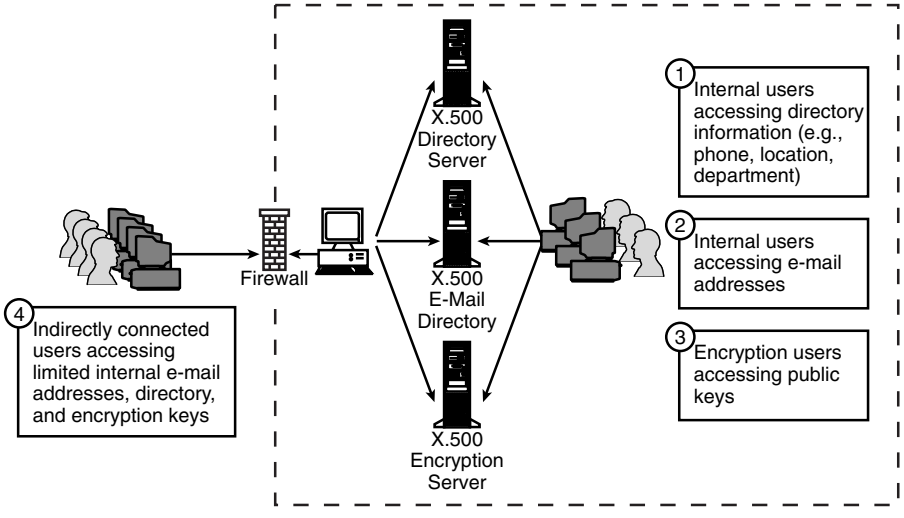

Exhibit 39-5. Authentication information for network, system, and application access.

All of this information is managed through the enterprise user and privilege management system to provide authentication information for network, system, and application access on a per-user basis (see Exhibit 39-5).

With the wide array of directory products available today, most enterprises will not have to develop their own technology, but are best served using X.500 directory services as they provide Lightweight Directory Access Protocol (LDAP) services that can be used by many of today's operating systems, including Windows 2000.

The enterprise directory can be used to provide the necessary details for environments that cannot access the directory directly, such as NIS and non-LDAP-ready Kerberos implementations.

Using the enterprise privilege management applications, a new user can be added in a few minutes, with all the necessary services configured. New applications and services can be added at any time. Should an employee no longer require access to specific applications or application privileges, the same tool can be used to remove them from the enterprise directory, and subsequently the application itself.

A major challenge for many enterprises is removing user access when that user's employment ends. The enterprise directory removes this problem because the information can be removed or invalidated within the directory, thereby preventing the possibility of the employee's access

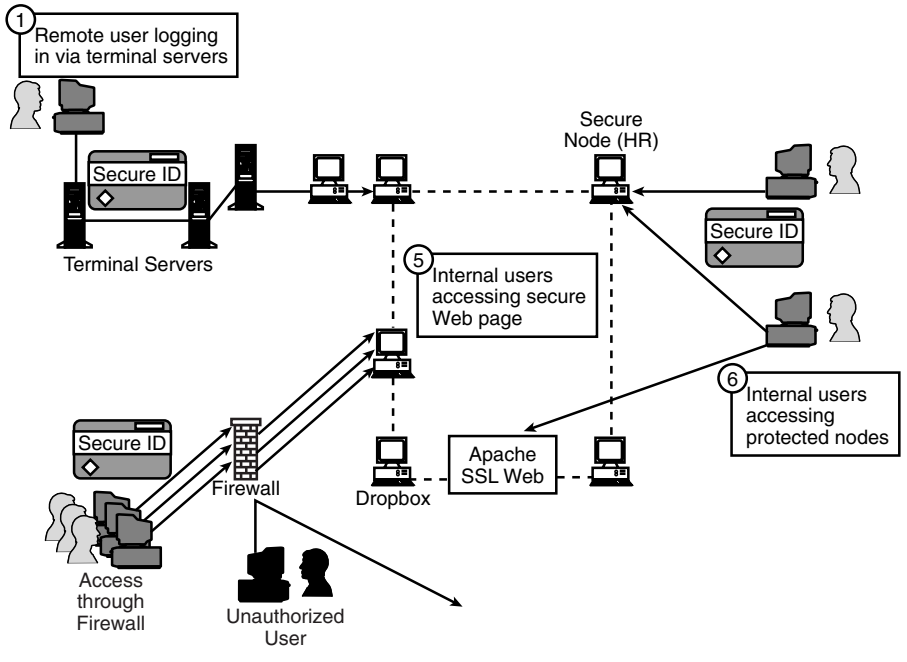

Exhibit 39-6. Authentication systems.

remaining active and exposing the company beyond the user's final day of work.

AUTHENTICATION SYSTEMS

There are many different identification and authentication systems available, including passwords, secure tokens, biometrics, and Kerberos to name a few (see Exhibit 39-6). The enterprise must ultimately decide what authentication method makes sense for its own business needs, and may require multiple systems for different information types within the enterprise. However, the common thread is that in today's environment, the simple password is just not good enough anymore.

When a user authenticates to a system or application, his credentials are validated against the enterprise directory, which then makes the decision to allow or deny the user's access request. The directory can also provide authorization information to the requesting application, thereby limiting the access rights for that user. Using this methodology, the exact authentication method is irrelevant and could be changed at any time. For example,

using a password today could be replaced with a secure token, biometrics, or Kerberos at any time, and multiple authentication technologies can easily coexist within the enterprise.

However, one must bear in mind that user authentication is only one aspect. A second aspect concerns authentication of the information. This is achieved through the use of a digital signature, which provides the authentication and integrity of the original message.

It is important to remember: no authentication method is perfect. As security professionals, we can only work to establish even greater levels of trust to the authenticating users.

ENCRYPTION SERVICES

Encryption is currently the only way to ensure the confidentiality of electronic information. In today's business environment, the protection of enterprise and strategic information has become a necessity. Consequently, the infrastructure requirements include encryption and digital signatures (see Exhibit 39-7).

Encryption of files before sending them over the Internet is essential, given the amount of business and intellectual property stolen over the Internet each year. The infrastructure must provide for key management, as well as the ability to handle keys of varying size. For example, global companies may require key management abilities for multiple key sizes.

Encryption of enterprise information may be required within applications. However, without a common application-based encryption method, this is difficult to achieve. Through the use of virtual private network (VPN) technologies, however, one can construct a VPN within the enterprise network for the protection of specific information, regardless of the underlying network technologies. Virtual private networking is also a critical service when sessions are carried over insecure networks such as the Internet.

In addition, the mobile user community must be able to protect the integrity and confidentiality of its data in the event a computer is stolen. This level of protection is accomplished with more than encryption, such as disk and system locking tools.

CUSTOMER AND BUSINESS PARTNER ACCESS

The use of the security infrastructure allows for the creation of secure environments for information exchange. One such example is the customer access network (see Exhibit 39-8) or those entry points where non-enterprise employees such as customers and suppliers can access the enterprise network and specific resources. In our global community, the

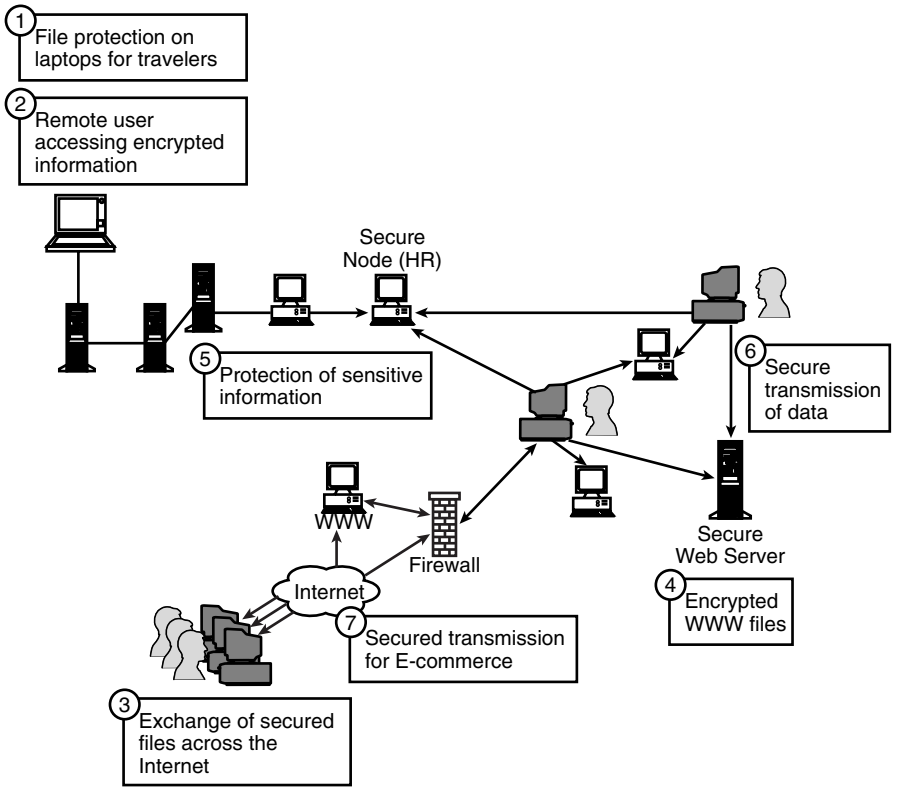

Exhibit 39-7. Encryption services.

number of networks being connected every day continues to grow. However, connecting one's corporate network to "theirs" also exposes one to all of the other networks "they" are connected to. Through the deployment of customer access networks, the ability to provide connectivity with security is achieved.

The customer access network is connected to the customer network and to one's corporate network, configured to prevent access between connected partners, and includes a firewall between it and the corporate network. In fact, the customer may also want a firewall between its network and the access point.

With VPN technologies, the customer access network may not be extremely complicated, but does result in a VPN endpoint and specific rules within the VPN device for restricting the protocol types and destinations that the customer is permitted to access.

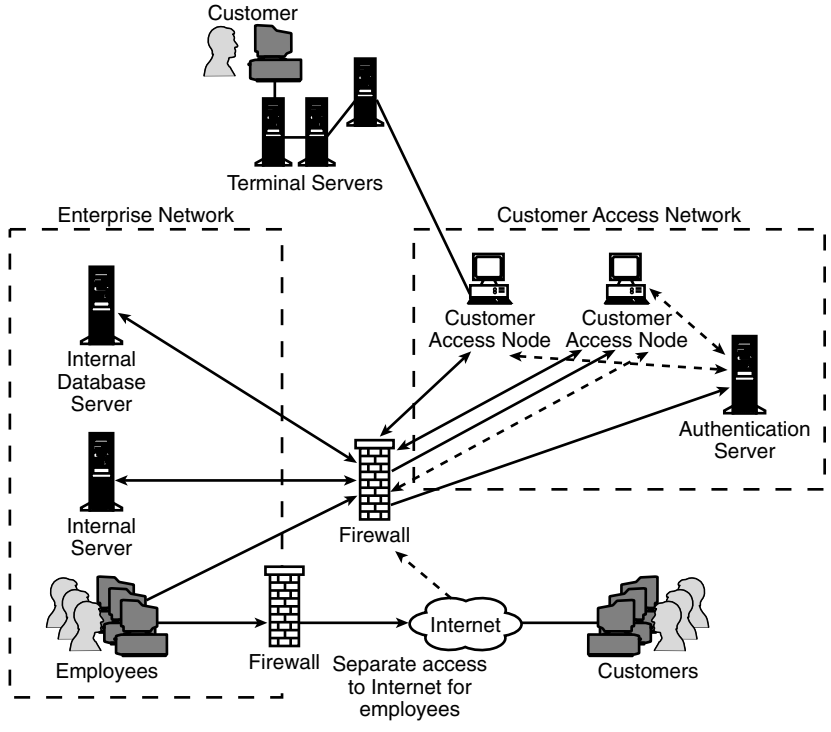

Exhibit 39-8. Customer access network.

The rules associated with the individual customer should be stored in the enterprise directory to allow easy setup and removal of the VPN access rules and keys. The real purpose behind the customer access network is not only to build a bridge between the two networks, but also to build a **secure** bridge.

CONCLUSION

This chatper focused on the technologies and concepts behind a security infrastructure. There are other elements that ideally should be part of the security infrastructure, including:

- Desktop and server anti-virus solutions
- Web and e-mail content filtering
- Anti-spam devices

At the same time, however, one's infrastructure must be designed at the conceptual level using the business processes and needs, and not be driven by the available technology. The adage that "the business must

drive the technology" is especially true. Many security and IT professionals forget that their jobs are dependent upon the viability and success of the enterprise — they exist to serve the enterprise, and not the other way around!

Many infrastructure designers are seduced by the latest and greatest technology. This can have dire consequences for the enterprise due to unreliable code or hardware. Additionally, one nevers knows when one has something that works because one is constantly changing it. To make matters worse, because the users will not know what the "flavor of the week" is, they will simply refuse to use it.

Through the development of a security infrastructure that is global in basis and supported by the management structure, the following benefits are realized:

- The ability to encourage developers to include security in the early stages of their new products or business processes
- The risk and costs associated with new ventures or business partners are reduced an order of magnitude from reactive processes
- Centralized planning and operations with an infrastructure responsive to meeting business needs
- Allow business application developers to deliver stronger controls over stored intellectual capital
- The risks associated with loss of confidentiality are minimized
- A strengthening of security capabilities within the installed backbone applications (e.g., e-mail, servers, WWW)
- The privacy and integrity associated with the corporation's intellectual capital are increased
- The risks and costs associated with security failures are reduced

In short, we have created a security infrastructure that protects the enterprise assets, is manageable, and is a business enabler.

Above all this, the infrastructure must allow the network users, developers, and administrators to contribute to the corporation's security by allowing them to "do the right thing."

ABOUT THE AUTHOR

Chris Hare, CISSP, CISA, is an information security and control consultant with Nortel Networks in Dallas, Texas. A frequent speaker and author, his experience includes from application design, quality assurance, systems administration and engineering, network analysis, and security consulting, operations, and architecture.