Visit techguide.com

The Technology Guide Series

How to Build Secure LANs with IPSec

This Guide has been sponsored by

techguide.com

Table of Contents

About the Editor

Jerry Ryan is a principal at ATG and the Editor-in-Chief of techguide.com. He is the author of numerous technology papers on various aspects of networking. Mr. Ryan has developed and taught many courses in network analysis and design for carriers, government agencies and private industry. He has provided consulting support in the area of WAN and LAN network design, negotiation with carriers for contract pricing and services, technology acquisition, customized software development for network administration, billing and auditing of telecommunication expenses, project management, and RFP generation. Mr. Ryan has been a member of the Networld+Interop Program Committee and the ComNet steering Committee. He holds a B.S. degree in electrical engineering.

---

209 West Central Street, Suite 301, Natick, MA 01760
Tel: (508) 651-1155, Fax: (508) 651-1171
E-mail: info@techguide.com, Web site: http://www.techguide.com

Abstract

Traditional network security focuses on establishing a perimeter to keep outsiders at bay and on limiting access through password protection, smart cards, or biometrics. Emerging Virtual Private Networks (VPNs) focus on secure site interconnection and remote access over the Wide Area Network (WAN). As e-business becomes imperative, organizations are exerting considerable effort and investing sizable budgets to secure back-office systems they have made accessible to customers and business partners. But a recent FBI study confirms what security experts have acknowledged for years: most security breaches occur within the corporate network, over a Local Area Network (LAN).

Protecting the privacy and integrity of data transported over LANs is paramount to securing a corporation's most valuable asset: its intellectual property. Standards-based technologies like Internet Protocol Security (IPSec) can be used to create trusted, secure LAN workgroups. This Technology Guide examines the risks associated with LAN-based Intranets and Extranets and the enabling technologies that can be employed to mitigate these risks. It describes how to build secure LANs using an integrated network-level approach that offers an optimum balance between security and performance. Finally, this Guide considers deployment strategies for rolling out secure LAN workgroups in a manageable, cost-effective manner.

Introduction

Over the past decade, Internet growth and its impact on business communication has been astounding. According to Forrester Research, business-to-business e-commerce already exceeds $100 billion, with sales projected to reach $1 trillion by the year 2002. Businesses rushing to “get connected” place their most valuable asset—information—at risk. By conducting on-line transactions with employees, customers, and business partners, companies make vast quantities of mission-critical and highly sensitive data accessible through corporate intranets, extranets—even over the public Internet.

As e-business grows, the increasing value placed on intellectual capital has spawned a new kind of criminal: the cyber-thief. Some cyber-thieves are outsiders, leveraging the Internet to access public servers and the data they contain. But the threat doesn’t stop at the edge of the corporate network.

The shared media technologies that power today’s corporate LANs make “insider theft” trivial to accomplish and difficult to detect. The same qualities that make LANs ubiquitous—simple, plug-and-play, shared media access—also make LANs the most vulnerable segment of the corporate network.

Furthermore, today’s corporate network is no longer an island, protected by isolation. LANs are morphing into “virtual LANs” (VLANs): geographically distributed workgroups that are related logically rather than physically. Increasingly, these virtual networks extend beyond the corporate premise, providing remote intranet access for teleworkers and travelers, even connecting business partners, customers, and suppliers through extranets. Intranets and extranets improve business efficiency and profitability, but also add risk. More information is exposed to a larger number of people, and corporate liability grows as every extranet member becomes responsible for safeguarding shared data.

Fortunately, an emerging standards-based technology called IPSec (Internet Protocol Security) can be used to protect IP packets against unauthorized disclosure and modification. IPSec is a security enhancement to the standard Internet Protocol (IP)—the network layer protocol embedded within virtually every network device today. IPSec is versatile: it can protect both WAN traffic exchanged between routers and LAN traffic exchanged between desktop and server. An integrated network-level implementation that combines embedded operating system support with hardware encryption can provide efficient end-to-end LAN security, significantly reducing the threat of cybertheft.

Rolling out IPSec—or any new technology—on an enterprise scale requires planning. Software and hardware upgrades must be installed in a coordinated fashion. Security policies must be defined to reflect business goals. Policies must be distributed and secure connectivity must be verified. With the proper tools, gradual deployment can be an affordable, manageable, and scalable strategy.

This Guide will show what secure LANs are, why they are important, provide strategies for how to implement them using IPSec, and illustrate practical deployment models.

Growing Security Risk: LAN-Based Intranets and Extranets

According to the Forrester Research Group, US business-to-consumer and business-to-business sales more than doubled between 1998 and 1999. An IDC study indicates that 97% of large corporations are connected to the Internet today, a third conducting sales over the Internet. The rise of both e-commerce and e-business has changed the structure of the corporate network. The need to communicate seamlessly with suppliers, customers, and business partners requires us to incorporate forms of public access into previously private networks.

How Extranets, Intranets Increase Risk

For example, business partners may create a joint venture to collaborate on new product design and, in doing so, exchange trade secrets. Whether the multi-enterprise team is co-located at one site or distributed across an extranet, the highly sensitive data of both partners now become accessible to anyone on the LAN. How can both partners be assured against cyber-theft of information through unauthorized eavesdropping on the affected LAN segments?

Extranets that connect customers with their suppliers can be just as vulnerable. Typically, these business-to-business extranets permit authenticated access to selected back-office servers and databases—for example, the parts inventory and production-planning systems operated by a manufacturer. But reaching these back-office systems may require “punching a hole” through the corporate network firewall and other perimeter defense measures. Extranet tunnels provide high-speed, unfettered access to trusted partners—and to cyber-thieves able to find and exploit this back door to the corporate LAN.

Even a business that grants no “public” access is not immune to cyber-theft. Private LAN access may be provided to contract and temporary employees. Just as password authentication prevents unauthorized access to internal payroll and benefits systems, so too must steps be taken to avoid disclosure of private data sent over the corporate LAN. And insiders who eventually become outsiders—including former permanent employees—leave the enterprise carrying privileged information: system names, logins, passwords, customer lists, trade secrets. Limiting insider information access based on "need to know" is a common sense measure to guard against future theft.

Cyber-Crime Begins At Home

These examples illustrate situations in which data exchanged within the corporate network becomes vulnerable to insider theft. In fact, a 1999 FBI/CSI study on computer crime showed that 55% of companies reported enterprise network security breaches within the LAN, due to unauthorized access by insiders. A whopping 51% of the respondents in this study acknowledged financial loss due to security breaches. And theft of proprietary information was the most costly type of loss, averaging $142K per incident, reaching $25M in one case. The American Society for Industrial Security estimates that intellectual property theft is costing US companies $24B each year. And the rate of insider theft is growing, up 10% between 1998 and 1999.

Figure1: The Insider Threat[1]

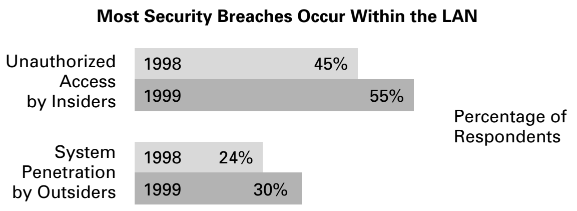

Source: FBI/CSI Computer Crime & Security Survey, 1999

Insider cyber-theft has been around—and ignored—for as long as corporate networks have existed. Local area networks, based on broadcast technologies that operate over shared media, are inherently insecure. It can be difficult to safeguard every bit of information, so instead we convince ourselves that prohibiting outsider access is sufficient. We trust insiders to enforce corporate policies preventing disclosure and improper use of proprietary information.

When insider theft inevitably occurs, companies are reluctant to disclose it. Enterprises keep mum to avoid public embarrassment and additional cost due to liability or copycat repetition. Even the damage from an insider attack may be greater than any an external attacker could inflict due to the sheer volume and sensitivity of systems and data accessible to an insider on the corporate LAN.

Securing Today's Corporate Network

"A company must protect its data at the source" is a security norm. An emerging corollary to this axiom is: "When transporting sensitive data, no shared link should be considered trustworthy." The common assumption that LANs and VLANs are trusted networks is invalid, even when the LAN is surrounded by traditional security measures.

Enterprises typically employ premise security measures to prevent physical theft and damage. Locked doors and keypad or card access protect against unauthorized access to data centers and individual offices. Asset management and monitoring tools can detect physical theft of computer assets (e.g., processors, storage devices, integrated peripherals).

But physically secure devices remain vulnerable to network threats. Traditional firewall security builds a network perimeter by blocking outside access to inside resources. Through packet inspection, filtering, and/or proxy services, firewalls implement policies that selectively permit access by authorized networks, hosts, and users. All other traffic originating outside the firewall is denied access to internal servers and desktops.

For years, companies have employed private dial pools and leased or switched lines to connect roaming users and remote offices. But once the

[1] Source: Intel White Paper "IP Security: Building Blocks for the Trusted Virtual Network" Figure A

remote user or router answered a password challenge, authenticated private WAN links were usually thought of as trusted. Today, many enterprises are leveraging the public Internet to reduce the cost of communication. Emerging VPNs are protecting packets that leave the LAN; while in transit across a public WAN. Stronger authentication methods are being combined with standards-based message integrity and confidentiality protocols to keep WAN traffic exchanged between access routers and firewalls secure.

These measures, while essential, still leave the most sensitive part of the corporate network—the local area network—unprotected. The vast majority of business communication occurs between desktops and servers, or between desktops, within a LAN-based workgroup. As we've seen, LANs also present the greatest threat of insider cyber-theft. The same standards-based protocols used to secure WAN traffic can also be used to safeguard LAN traffic. IPSec provides flexible options that meet the unique performance and security needs of both environments.

Figure 2: Securing Intranets & Extranets at All Levels

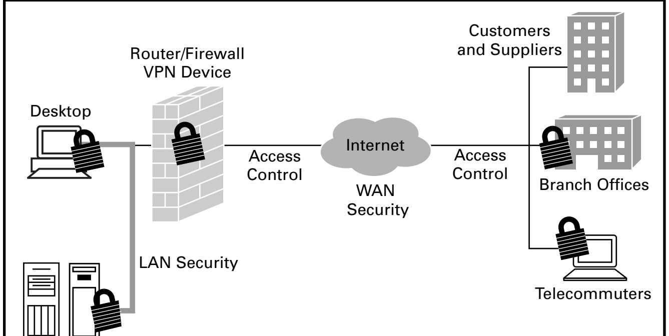

Source: Intel Presentation "Making Money by Adding Security to Your Customers' Networks"

Indeed, consistent enterprise-wide deployment of solutions that implement a centrally managed policy is essential to ensure end-to-end security in today's corporate network. An affordable, scalable, manageable, interoperable combination of physical security, access control, perimeter defense, and WAN and LAN security measures is the enterprise's best defense against cyber-theft.

Integrated IPSec as an Enabling Solution

Over the years, improvements have been necessary to adapt the Internet protocol (IP) to accommodate the phenomenal growth of the public Internet and changes in network topology: roaming users, wireless networks, virtual private networks. Recognizing the increase in e-business activity and consequent heightened concern over network security, the Internet Engineering Task Force (IETF) defined IPSec: a set of security extensions designed to enable access control, data confidentiality, and protection against message modification and replay.

The design of IPSec was influenced by early B2B extranet adopters, most notably the Automotive Exchange Network (ANX). The ANX is a large supply-chain network that serves the automotive industry. Through a process of stepwise-refinement, IPSec standards and early products have been implemented and tested over several years, now reaching a level of maturity appropriate for large scale enterprise use.

Today, IPSec has been integrated into Microsoft Windows 2000 and most analysts predict widespread de facto deployment of this network layer security solution.

What is IPSec?

IPSec consists of two Internet protocol extensions—the Authentication Header (AH) and the Encapsulating Security Payload (ESP)—that support creation of secure networks. Access is controlled by a companion key management protocol called the Internet Key Exchange (IKE). Together, IPSec and IKE ensure that authorized parties may exchange private IP packets securely over a public network. Packet content is kept confidential by applying encryption and protected against modification through digital signing. IPSec can be used to securely "tunnel" packets to routers or firewalls over a WAN, or to securely "transport" packets end-to-end between desktops and servers.

IPSec works transparently: the end user need not be aware that packets are being intercepted and transformed by IPSec. This transparency ensures that IPSec-based security policies are not circumvented accidentally or intentionally. Because it operates at the network layer, IPSec is perfectly positioned to enforce corporate network security. As IPSec matures, embedded operating system support, policy-based management tools, hardware encryption, and other advances will promote seamless integration of security into the network infrastructure.

Message Integrity

IPSec provides message integrity services that allow a packet recipient to be certain of the sender's identity. This prevents session hijacking: for example, preventing a thief from using captured packets to masquerade as a legitimate client and steal information from or perform bogus transactions with an enterprise server.

Message integrity services also allow a packet recipient to be certain that a packet has not been modified in transit and that packets are not being replayed. This prevents a thief from sniffing legitimate packets on the LAN, modifying their content, then resending them either immediately or at a later time—for example, repeating an accounting system transaction authorizing payment to the thief's own bank account rather than the legitimate payee.

Both AH and ESP use digital signatures to provide message integrity services. They take part of the IP packet, or the entire IP packet, and generate a message hash, then encrypt the hash with a secret key known only to the sender and receiver. This is analogous to applying a tamper-proof identity seal to the back of a sealed envelope. Upon receipt, if the seal has not been broken, the reader knows the envelope's contents are authentic. AH signs the entire IP packet, including the source address, while ESP leaves the source address unprotected.

The hashed message authentication code (HMAC) algorithms typically used to create this digital signature are MD5 and SHA-1. Choosing the algorithm, providing the secret key to sender and receiver, and deciding whether to use AH or ESP are matters dictated by a company's security policy. Sender and receiver must have complementary security policies for communication to be successful.

Figure 3: How IPSec Protects Packets

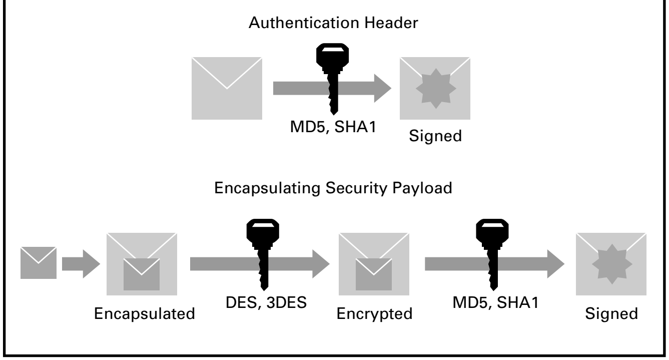

Data Confidentiality

In addition, IPSec ESP provides data confidentiality, assuring the packet sender and receiver that others cannot see the packet content. ESP does so by scrambling part, or the entire original packet, with a symmetric encryption algorithm like DES or 3DES. Again, only the sender and receiver hold the secret key used to encrypt and decrypt the packet.

Data confidentiality prevents sniffing of LAN packets that might otherwise reveal trade secrets, payroll data, or benefits information. A LAN segment used for both extranet and intranet traffic can employ different keys for each workgroup or pair of communication partners, ensuring that only those with a legitimate "need to know" can read each packet. Any thief intercepting encrypted packets sees only garbled data, not the sensitive information conveyed by the packet.

The more sensitive the information and the longer the communication exchange, the more vital it becomes to employ a strong encryption algorithm and key. But stronger encryption requires greater processing power. Balancing risk against cost is the role of enterprise security policy: for example, an enterprise might use 3DES to encrypt financial transactions but choose the faster, weaker DES to encrypt all other private data.

Tunnel over WAN, Transport over LAN

When packets flow across a WAN, an outsider might deduce something by viewing source and destination addresses contained in the IP header. For example, company A is exchanging high-volume traffic with company B; a thief might use this knowledge to predict a pending joint venture and illegally profit from use of insider information. Addresses can also expose details of the corporate network's internal topology, facilitating denial-of-service (DOS) attacks whereby a hacker floods an enterprise server with requests that block access by legitimate users.

To keep addresses private, IPSec can be used in tunnel mode. The entire private IP packet—header and payload—is hidden inside a public IP packet "envelope". Tunnel mode is typically employed by security gateways: edge devices like routers and firewalls that relay packets on another system's behalf.

But, inside a LAN, the threat of traffic analysis and denial-of-service attacks is minimal. To reduce processing overhead and packet length without sacrificing security, the original header can be used on packets exchanged between hosts. In transport mode, ESP hides only the private packet's payload. Transport mode IPSec can be used to efficiently protect data end-to-end between clients and servers, peers in a workgroup, and extranet partners.

Figure 4: Tunnel vs. Transport Mode IPSec

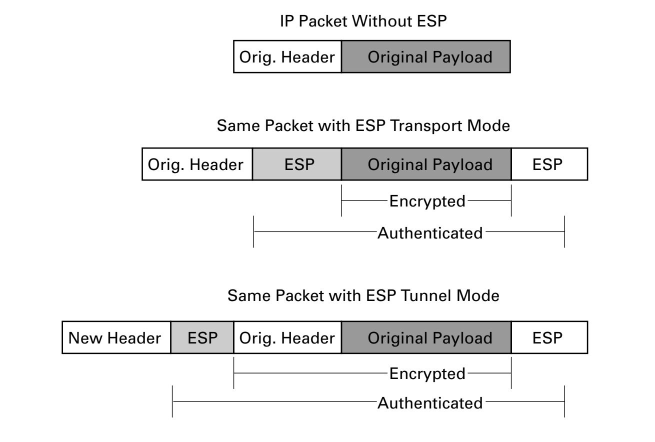

Transport and tunnel mode can be used in conjunction to secure the total enterprise network by applying each where appropriate: tunnel mode to WAN security, transport mode to LAN security.

Automating IPSec Setup with IKE

As we've seen, establishing IPSec communication requires agreement on protocols, algorithms, and secret keys—these details must be securely negotiated between partners. Furthermore, communication partners must exchange credentials so that access can be controlled based on authenticated identity. These two steps are automated by IKE:

1. A security policy database determines the type of protection required by an IP packet, based on information contained within the packet: the source and destination address, the type of application protocol it carries, etc.
2. When an IPSec-enabled device wants to send a packet, it consults security policy. If it is not already communicating with the packet's destination, it proposes a "security association" by sending an IKE message to the destination system (transport mode) or the security gateway responsible for routing traffic to the destination (tunnel mode).
3. The receiver consults its own policy, authenticates the sender's identity, and verifies the sender has permission to access the destination. The receiver returns its own credentials so the IKE sender can do the same. Credentials can be something simple, like a numeric password known as a preshared key. For environments where stronger authentication is desired, digital certificates permit identity verification with a trusted third party.
4. Once authenticated, the IKE sender and receiver negotiate message integrity and data confidentiality measures to protect future IKE exchanges between them. They use this secure control channel to negotiate parameters and share secret keys used to authenticate and encrypt IPSec packets.
5. Because keys can be cracked if thieves are given enough time to do so, this setup process must be repeated at specified intervals. These IKE and IPSec security association "lifetimes" are determined by security policies that balance risk against setup overhead.

By automating endpoint authentication, parameter negotiation, and secret key exchange, IKE enables scalable IPSec deployment and enforcement of an organization's security policy.

How to Build Secure LANs Using IPSec-Enabled Network Adapters

To provide efficient end-to-end security, IT managers will deploy security measures at several levels throughout the corporate network. As we've seen, security is just as important for LANs as it is for WANs; LANs are not more trustworthy simply because they are hidden from outside access by a firewall. Deploying transport mode IPSec on desktops and servers is necessary to provide adequate protection against insider theft.

Taking the First Step: Software Support

New operating systems like Windows 2000 provide embedded support for IPSec, simplifying desktop and server security deployment. Over time, embedded OS support will provide 100% coverage for every device participating in the enterprise network. IPSec will become as ubiquitous as TCP/IP is today.

Until then, other measures are needed. Add-on software can be used to add IPSec support to desktops and servers running legacy operating systems such as Windows NT 4.0 and soon Windows 98[LAP1]. Companies require standards-based IPSec solutions that enable every affected device to interoperate smoothly when implementing enterprise-wide security policies.

Increasing Efficiency with Encryption Co-Processors

Software support is the first step. However, IT managers must always strive for operational efficiency. In an ideal world, adding security measures would not impact network or system performance. In reality, some additional overhead is inevitable; the trick is to minimize it. Encryption is processor-intensive: a software-only IPSec solution performing strong 3DES encryption can triple CPU load and significantly reduce packet throughput, as shown in the following chart.

Enterprises need an efficient IPSec solution that allows desktops and servers to focus on the task at hand and not security. Fortunately, encryption can be offloaded to hardware, including network encryption co-processors and desktop and server network adapter cards.

IPSec-enabled network adapters balance security and performance by allowing the operating system and hardware to share the workload. The operating system supports policy configuration and enforcement, while the co-processor chip performs the processor-intensive work: encryption/decryption, hashed message authentication, and TCP checksums. This integrated approach frees the CPU to support business tasks at optimal speed without sacrificing network performance.

Hardware support increases IPSec efficiency, whether encryption is performed on a firewall tunnel endpoint or a server transport endpoint. As a rule of thumb, devices that support many users and high-volume transactions benefit the most from hardware encryption.

Tools to Assist with Enterprise-Scale Deployment

Once an enterprise decides to implement LAN security with IPSec-enabled operating systems and network adapters, the next step is to plan and carry out deployment. In very small companies, ad hoc installation and update may suffice. However, in larger enterprises, IT managers require tools that enable scalable, automated deployment and produce a manageable network environment. This challenge can be broken down into three phases: installation, on-going management, and security policy implementation.

Efficient Installation and Management

Automated management tools for distributed LANs can simplify large-scale deployment of operating system software, updates, and configuration changes. Some products make it possible to boot new operating systems over the LAN. Others reduce the impact of on-going update by enabling off-hours software installation. For example, a PC workgroup can be remotely upgraded—even remotely powered on—to install new Windows NT service packs or drivers over the weekend, avoiding disruption of normal business operation.

Installing new hardware requires hands-on effort, but this process can also be streamlined. Look for systems and network adapters that support hot install so that servers need not be powered down during hardware upgrade. Eliminate the complexity of finding the right adapter for each device by installing “plug and play” cards that employ a uniform driver that supports all adapters in the product line (desktop, server and mobile).

To reduce the cost of on-going asset management, configuration, and maintenance, use systems and adapters that offer built-in support for Internet-standard management interfaces like SNMP and the DMTF’s Desktop Management Interface (DMI) and

Common Information Model (CIM). Enterprises that employ a distributed management solution like Tivoli TME or Intel LANDesk should seek out "managed" adapters that include compatible agents.

Centralized Security Policy Implementation

IPSec-enabled operating systems and network adapters provide the infrastructure for implementing enterprise network security. However, this is like having a well-equipped workbench—the carpenter must still create a design and implement that vision. With IPSec, companies must first define a suitable network security policy and then use tools to automate its implementation.

For example, the Windows 2000 Microsoft Management Console (MMC) security policy manager provides the ability to centrally-configure IPSec policies and make them available to homogeneous servers and clients. MMC snap-ins reduce the cost of security configuration and analysis of Windows 2000 networks by automating policy distribution to a large number of PCs.

The final piece of the puzzle is the creation, secure distribution, verification and revocation of credentials used for authentication. Preshared keys are manageable on a small scale. Public Key Infrastructure (PKI) products like Entrust/PKI support strong third-party authentication based on digital certificates. PKI is widely recognized as an important enabler for scalable deployment of IPSec in large, distributed enterprise networks.

Secure LAN Deployment

Enterprises should add IP security to local area networks using a gradual three-phase strategy that combines start-up simplicity with the experience developed over time.

Many benefits can be achieved simply by allowing LAN resources at highest risk to communicate securely, without disrupting "business as usual" elsewhere. One generic policy that permits optional encryption can be deployed without security expertise, on-going maintenance, or synchronized rollout. This simple, non-disruptive first step lets an organization get its feet wet with minimum risk or investment.

As a company becomes more comfortable with IPSec, it can gradually increase the level of security for resources at highest risk. Custom policies can be used to create secure workgroups—for example, require encryption for all traffic exchanged within the finance department and prohibit outsider access to finance servers. Fine-tuning can focus on meeting business goals, applying custom policies when and where they are most needed. Over time, as customization increases, the organization will grow its security expertise and infrastructure.

Eventually, a company may decide to tackle more advanced security measures, adding strong authentication using a public key infrastructure. Companies do not need to reach this phase to benefit from secure LAN deployment. However, those that do can achieve very strong security, maximizing their return-on-investment in IPSec.

This Technology Guide describes this three-phase strategy in general terms that can be applied to many different enterprises, from small, single-site businesses to large, multi-national corporations. Of course, implementation details differ for each enterprise, influenced by network size, use, and a variety of business factors. Companies can use this generic strategy as a guide to develop their own secure LAN implementation plans.

Phase One: A Simple, Generic Policy

When IPSec capabilities are added to a desktop or server, security policy dictates how that system may communicate with other systems on the LAN; which systems require secure, IPSec-encrypted communication; and which systems are permitted "clear text" (non-encrypted) communication. The easiest way to get started is to implement one generic security policy on every IPSec-capable system.

Of course, IPSec capabilities will be added to desktops and servers over time. Systems that host proprietary information or perform sensitive transactions may be given priority during rollout. When IPSec has been added to some systems but not others, care must be taken to avoid service disruption. This is easily accomplished by making IPSec optional during this start-up phase.

Simple, gradual secure LAN deployment can begin with one generic secure initiator policy. A system with this policy will exchange secure, encrypted transactions with any other IPSec-capable system as its preferred mode of operation. However, this policy treats security as discretionary. If a secure association cannot be established with another system, the IPSec-capable system will fall back to a regular "clear text" connection. This policy facilitates gradual rollout by providing concurrent support for both non-secure and secure transactions on the same desktop or server. There is no need to synchronize rollout; systems can be upgraded independently. There is no need to create or maintain complex security policies; one simple "canned" policy, applied uniformly on every upgraded system, does the job.

Figure 5: Phase One - Simple, Generic Policy

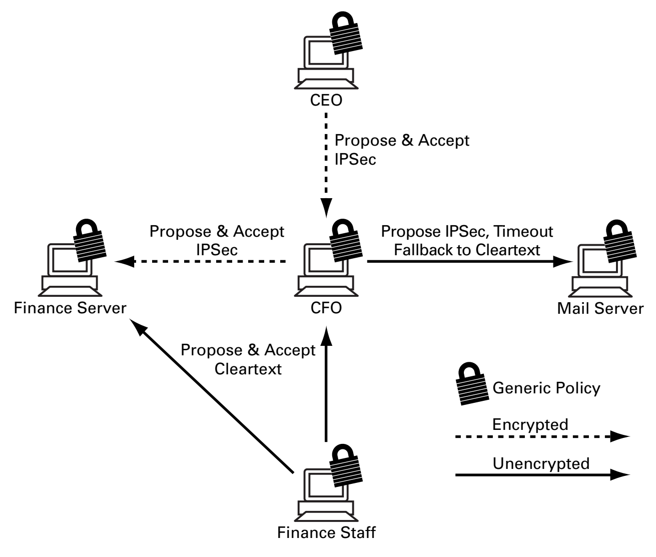

Consider the example shown in Figure 5. Executive desktops and a finance server are given priority during secure LAN rollout, and are the first to be upgraded to support IPSec. After upgrade, whenever the CEO and CFO communicate with each other, they always use IPSec-encrypted security associations. The CFO now also connects securely to the finance server. But, of course, the CFO still needs to communicate with desktops and servers that have not yet been upgraded. To ensure a smooth transition, the CFO will accept clear text connections initiated by others. The CFO will also fall back to clear text when connecting to public servers (e.g., the company's mail server). This "initiate IPSec, fall back to clear text" behavior applies to every IPSec-capable desktop or server because all share the same generic policy.

During this initial phase, many secure LAN benefits are realized: traffic at highest risk is protected from modification and disclosure. In this example, communication between the CEO, CFO, and finance server can no longer be "sniffed" by others on the LAN. But all systems using the same simple policy share credentials and are treated as peers. Access control is limited to operating system passwords. Because the server can be accessed by anyone on the LAN, traditional file system security is still required to restrict file access to authorized staff members.

Phase Two: Custom Workgroups

Over time, IPSec capabilities will be rolled out to many—perhaps every—desktop and server in the enterprise. IT and network administrators will gain experience with IPSec, developing procedures that streamline system upgrades. They will also become more comfortable with IPSec technology and more familiar with choices that can be made within security policies. At this point, an organization may begin to experiment with distinct security policies for individual workgroups.

As before, this phase also starts by identifying the most sensitive and vulnerable data, the applications and servers at highest risk, and the desktops or users involved in these transactions. But this phase goes a step further: high-risk systems are compartmentalized into logical workgroups.

Workgroups are based on organizational affinity and assigned group credentials. Different security policies may be defined for communication within each workgroup and for communication between workgroups. Inter-workgroup communication may require stronger security than internal transactions between trusted members. Workgroup policies control can be prioritized so that strong security can be applied to some connections, no security to others.

Figure 6: Phase Two - Custom Workgroup Policies

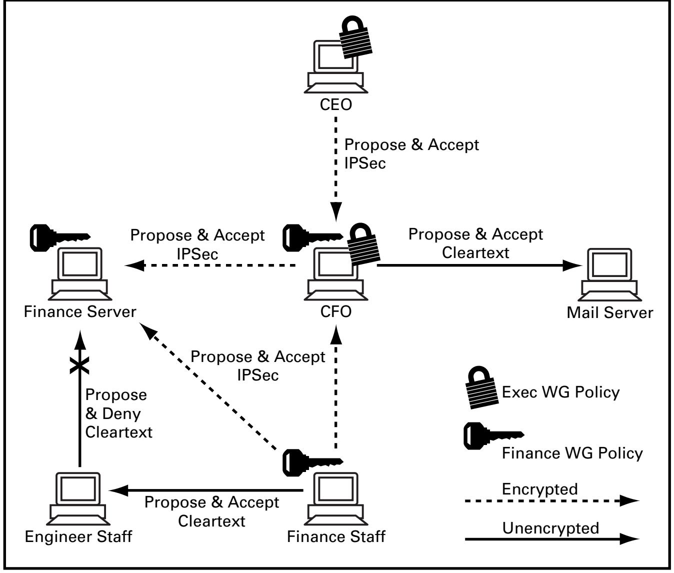

Workgroup policies can control access by individual users or the entire workgroup. Consider the example shown in Figure 6. After rolling out one generic policy to the executive staff and the entire finance department, this company has decided to strengthen LAN security by creating two workgroups. The executive team, including the CEO and CFO, adopt a custom "Executive Workgroup" security policy. Every finance department member, including the CFO, adopts a custom "Finance Workgroup" security policy. These two custom policies differ from the simple, generic policy deployed in phase one: each custom policy has a group identity used to prevent outsider access. For example, perhaps Executive Workgroup members share one secret key, while Finance Workgroup members share another secret key. By using custom policies, the CEO can differentiate between the CFO and other finance department members.

In addition, custom policies can use different connection initiator, responder, and fallback rules. With the generic policy, the finance server was always willing to fall back to a clear text connection. With a custom policy, server access can be restricted to members of the finance department by eliminating this fall back rule. With the generic policy, systems always propose an IPSec-encrypted connection. With a custom policy, members of the finance department can propose clear text connections to outsiders, including the public mail server. Customized policies let a company fine-tune the behavior of workgroups. Workgroups at highest risk for cyber-theft can be "locked down", while others can use either secure or non-secure connections, depending upon the identity of the far-end system.

When a system adopts more than one policy, filters are used to determine which policy applies to each IP packet. For example, the Finance Workgroup policy might be applied using the department's LAN subnet as a filter. But the more tightly policies are bound to network addresses, the more likely they are to require update over time. When possible, it is better to filter on other values, such as organizational unit. Designing good, easily maintained filters is an important part of custom policy deployment.

A few custom policies can be maintained easily. Workgroups should be small enough to capitalize on trust between members, but large enough for the policy database to remain manageable. A small set of commonly used policies, augmented by granular exception policies that reflect business needs, can simplify on-going maintenance.

Phase Three: Adding Trust Infrastructure

In any IPSec security association, credentials identify communication partners (i.e., IPSec endpoints). The way in which partners identify themselves has a direct impact on access control strength and deployment effort. For simplicity, enterprises can start with shared credentials, adding stronger authentication over time.

During initial deployment, generic policies use a common preshared key. This approach is quite simple to implement. But anyone who has the common preshared key has the same IPSec access. All IPSec-capable systems on the LAN are peers. Former employees leaving the company may know the common preshared key.

During phase two, custom policies target systems at greatest risk and increase the level of security provided for them. Companies may continue to use preshared key authentication at this more granular level. But, during this phase, most companies begin using a management system for central configuration and automated "push" or "pull" policy deployment to desktops and servers. Organizations in phase two gradually grow their security expertise and build management infrastructure.

Eventually, some companies enter a third, advanced phase that replaces preshared keys with credentials enabling strong authentication. In homogenous Windows 2000 environments, Kerberos tickets might be used. In all other environments, including mixed Windows networks; a Public Key Infrastructure should be used. IPSec partners can be authenticated with "raw" public keys, but most enterprises will prefer X.509 digital certificates verified by a trusted Certification Authority (CA). PKI combines strong access control with scalable administration through the use of CA hierarchies, cross-certification, and delegation.

Continuing the example introduced in Figures 5 and 6, X.509 digital certificates might be issued to uniquely identify each member of the executive team. Doing so allows the CEO to differentiate between requests made by the CFO and the CTO, for example. When the CFO leaves the company or changes jobs, his digital certificate can be revoked. Managing the creation, secure distribution, and revocation of digital certificates on a per-employee basis requires a trust infrastructure. Companies enter this phase when they have gained security expertise through previous phase experience and are ready to adopt the tools required for large-scale deployment.

Network and System Infrastructure Considerations

Each deployment model and phase involves upgrading desktops and servers throughout the enterprise LAN. Network and system considerations that may influence deployment include the following:

- Wherever possible, upgrade affected resources with Windows 2000 and IPSec-enabled adapters to form secure LAN workgroups.
- Software-only encryption is not recommended on servers due to high performance demands. It is generally better to upgrade the operating system and adapters before implementing IPSec security on a server.
- For existing workgroups, target Windows 2000 early adopters first, following the company's strategy for phased introduction of new operating systems.
- Deploy IPSec-enabled operating systems and adapters in all new workgroups, especially those where security risk is highest (e.g., Extranets).
- Build in IPSec support from the ground up: make every new adapter an IPSec-enabled adapter, even if the desktop or server is not (yet) running Windows 2000. Be ready to enable security when operating system support is deployed.

Summary

Security is emerging as one of the pivotal issues in today's enterprise network. Experience shows that because local area networks are such ready targets for insider theft, they pose the greatest security risk. Intellectual capital on intranets and extranets must be safeguarded.

End-to-end network security built upon IPSec with hardware-based encryption offers an optimum balance between security and performance. IPSec-enabled operating systems and network adapters provide the building blocks needed to create secure LAN workgroups. Companies must develop policies that reflect business needs—then implement them in a phased deployment strategy that secures data at highest risk first.

Automated installation, configuration, and monitoring tools can reduce the upgrade cost and ongoing administration. By starting with a simple deployment model, companies can grow their security expertise and build up the necessary infrastructure. In large enterprise networks, policy-based management systems and PKI enable scalable secure LAN administration.

LAN security will soon become a check-off item, an expected part of every company's network infrastructure. Enterprises that start phased IPSec deployment today will reach this target faster and gain a competitive edge in today's rapidly expanding e-business environment.

This Technology Guide is one in an ongoing series of over 100 solutions-focused Guides. These Guides assist IT professionals in making informed business decisions about specific aspects of technology development and strategic deployment.

In a non-biased, easy-to-understand style and tone, The Technology Guide Series® offers a broad array of titles, each presenting objective information and practical guidance. Our editorial writing team has many years of experience in IT and communications technologies, and is highly conversant in today's emerging technologies.

The Technology Guide Series and techguide.com are supported by a consortium of leading technology providers. The Sponsor has lent its support to produce and publish this Guide.

This Guide, as well as the entire Technology Guide Series, is made available to view and print at no charge by visiting techguide.com.

Over 100 Technology Guides in the following categories:

Software Applications
Network Management
Enterprise Solutions
Network Technology
Telecommunications
Convergence/CTI
Internet
Security

produced and published by

techguide.com