DATA SECURITY MANAGEMENT

IMPROVING NETWORK-LEVEL SECURITY THROUGH REAL-TIME MONITORING AND INTRUSION DETECTION

Chris Hare

INSIDE

Today's Security Perimeter: How to Protect the Network; The Moat; The Threat of Attack; Unauthorized Computer Use; Financial Losses; Our Employees Are Against Us; Where Is the Critical Information? Future of Network Security; Secure Gateway Types; Security Layering; Security Goals; Types of Intrusion Monitoring and Detection Systems; Why Intrusion Monitoring and Detection? Implementation Examples; Monitoring at the Secure Gateway; Monitoring at the Remote Access Service Entry; Monitoring Within the Corporate Network; Monitoring the Extranet; Proactive and Reactive Monitoring; Computer Incident Response Team; Penetration and Compliance Testing; Types of Penetration Tests

INTRODUCTION

Corporations are seeking perimeter defenses without impeding business. They have to contend with a mix of employees and non-employees on the corporate network. They must be able to address issues in a short time period due to the small window of opportunity to detect inappropriate behavior.

PAYOFF IDEA

Since the Internet really gained popularity and companies have chosen to protect their corporate information from the threat of attack, network-based computing and information sharing have increased dramatically. The amount of data traffic carried on the internal network is growing at a phenomenal rate. It has been difficult to keep ahead of the attackers, as most computer and network security groups are limited to what they have available. Consequently, implementing a realtime monitoring and intrusion detection program can signal trouble areas that might otherwise go unnoticed.

TODAY'S SECURITY PERIMETER: HOW TO PROTECT THE NETWORK

Many companies protect their networks from unauthorized access by implementing a **security program** using perimeter protection devices, including the *screening router* and the *secure gateway*. A screening router is a network device that offers the standard network routing services, and incorporates filters or access

Auerbach Publications

control lists to limit the type of traffic that can pass through the router. A firewall or secure gateway is a computer that runs specialized software to limit the traffic that can pass through the gateway. (The term "secure gateway" is used here rather than the more generic term "firewall.")

While on the surface, they seem like they are doing the same thing, and in some respects they are, the router and the secure gateway operate at different levels. The screening router and the secure gateway both offer services that protect entry into the protected network. Their combined operation establishes the firewall as shown in Exhibit 1.

Establishing firewalls at the entry points to the corporate network creates a moat-like effect. That means that there is a "moat" around the corporate network that separates it from other external networks.

THE MOAT

While the moat provides good protection, it reduces the ability of the organization to respond quickly to changes in network design, traffic patterns, and connectivity requirements (see Exhibit 2). This lack of adaptability to new requirements has been evident throughout the deployment of the secure gateways within numerous organizations.

One of the major complaints surrounds the limited application access that is available to authorized business partner users on the external side of the firewall.

In some situations, this access has been limited not by the authorizations allowed to those users, but to the secure gateway itself. These same limitations have prevented the deployment of firewalls to protect specific network segments within the corporate network.

Many organizations are only connected to the internet and only have a need to protect themselves at that point of entry. However, many others connect to business partners and business partners, who are in turn connected to other networks. None of these points of entry can be ignored.

It is, in fact, highly recommended that today's organizations establish a centralized security team that is responsible for the operation of the various security devices. This places responsibility for the operation of that infrastructure into one group who must do the planning, implement, and take action to maintain it.

THE THREAT OF ATTACK

The threat of attack comes from two major directions: attacks based outside the corporate network and attacks based from within. The moat security model, which is working effectively at many organizations, addresses the "attack from without" scenario. Even then, it cannot reliably provide information on the number of attacks, types of attacks, and their point of origin.

EXHIBIT 1 — The Firewall Is Composed of Both the Screening Router and the Secure Gateway

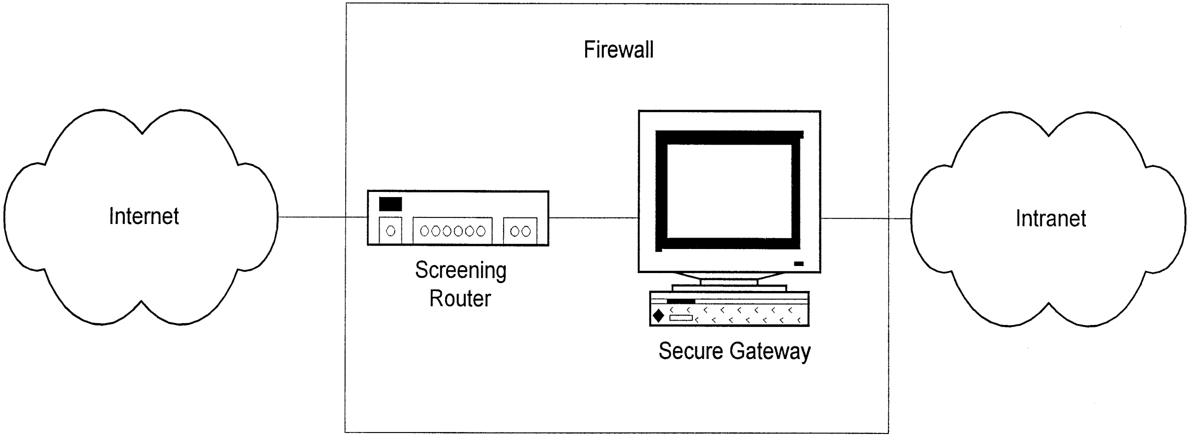

Auerbach Publications

EXHIBIT 2 — Establishing Firewalls at the Entry Points to the Corporate Network Creates a Moat-Like Effect

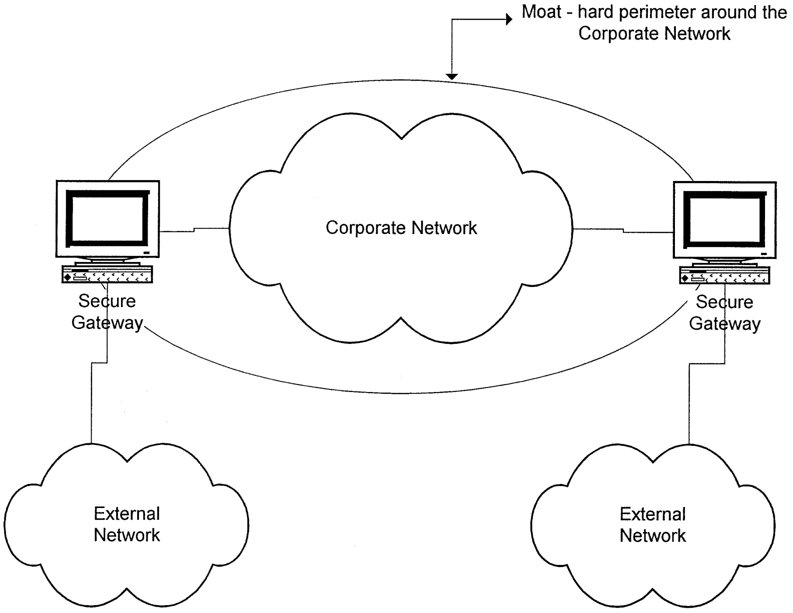

However, the moat cannot address the "attack from within" model, as the attack is occurring from within the walls. Consider the castle of medieval times. The moat was constructed to assist in warding off attacks from neighboring hostile forces. However, when fighting breaks out inside the castle walls, the moat offers no value.

The definition of an intrusion attempt is the potential possibility of a deliberate unauthorized attempt to:

- access information
- manipulate information
- render a system unreliable or unusable

However, an attack is a single unauthorized access attempt, or unauthorized use attempt, regardless of success.

UNAUTHORIZED COMPUTER USE

The problem is that the existing perimeter does not protect from an attack from within. The major security surveys continually report that the smallest percentage of loss comes from attacks that originate outside the

EXHIBIT 3 — Computer Security Institute 1999 Survey

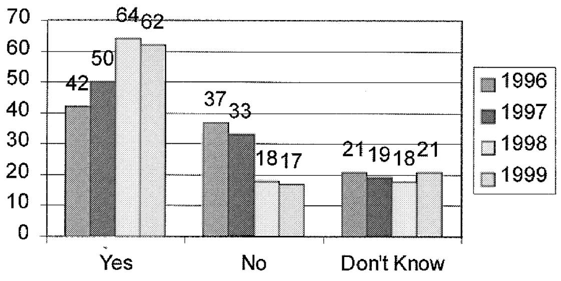

organization. This means that the employees are really the largest threat to the organization

The Computer Security Institute conducts an annual survey of its membership in conjunction with the FBI Computer Crime Unit. In the 1999 survey, the question was asked: "Has your organization experienced an incident involving the unauthorized use of a computer system?" (see Exhibit 3). As indicated, there was an overwhelming positive response, which had been climbing over the previous three years, but which saw a slight drop in an affirmative response. Many organizations could answer "Yes" to this question, but there is also a strong element of "I don't know." This element is because the only unauthorized use one is aware of is what is ultimately reported or found as a result of some other factor.

The cost of the information loss is staggering, as illustrated in the following information (also from the CSI Computer Security Survey). From that survey, it is evident that unauthorized insider access and theft of proprietary information has the highest reported cost. Given the potential value of the technical, R&D, marketing, and strategic business information that is available on the network, more and more companies need to focus additional attention to the protection of the data and securing the network.

FINANCIAL LOSSES

The financial impact to organizations continues to add up to staggering figures: a total of over $123 million as reported in the survey (see Exhibit 4). The survey identified that there has been a increase in the cost of unauthorized access by insiders, and the cost in other areas has also risen dramatically. The survey also identified that there continues to be an increase in the number of attacks driven from outside the reporting organizations. This is largely due to the increasing sophistication of the network attack tools and the number of attackers who are using them.

Intrusion detection and monitoring systems can assist in reducing the "I don't know" factor by providing a point where unauthorized or unde-

EXHIBIT 4 — Dollar Amount of Losses by Type

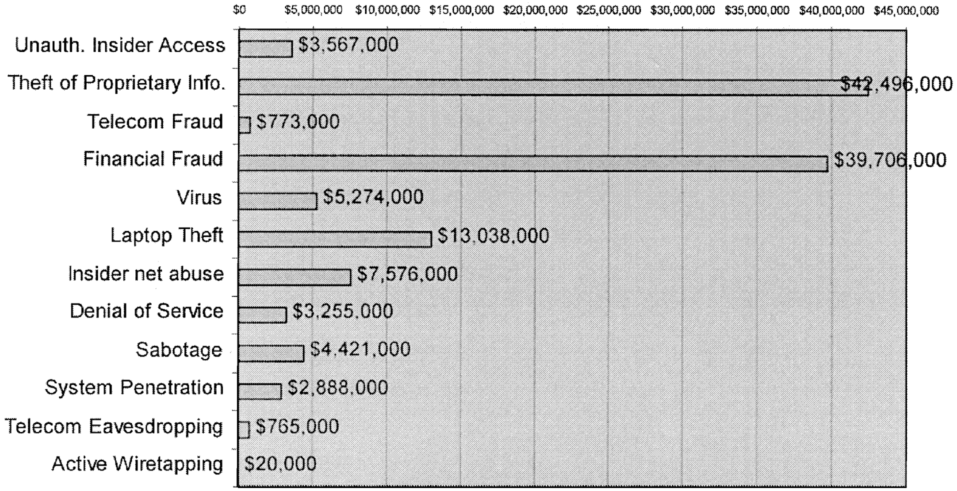

Auerbach Publications

sirable use can be viewed, and appropriate action taken either in realtime or after the fact.

OUR EMPLOYEES ARE AGAINST US

An often-quoted metric is that one of 700 employees is actively working against the company. This means that if an organization has 7000 employees, there are ten employees actively working against the organization's best interests. While this sounds like a small number of people, the nature of who they are in an organization will dictate what they have access to and can easily use against the company.

The most recent American Society for Industrial Security (ASIS at http://www.asisonline.org) "Trends in Intellectual Property Loss" survey report suggests that approximately 75 percent of technology losses occur from employees and those with a trusted relationship to the company (i.e., contractors and subcontractors). Computer intrusions involve approximately 87 percent of the insider issue.

While organizations typically have the perimeter secure, the corporate network is wide open, with all manner of information available to every one who has network access. This includes employees, contractors, suppliers, and customers! How does an organization know that their vital information is not being carried out of the network? The truth is that many do not know, and in many cases it is almost impossible to tell.

WHERE IS THE CRITICAL INFORMATION?

The other aspect to this is that many organizations do not know where their critical information is stored. This does not even mean where the source code or technical information is stored. That is important, but one's competitors will be building similar products. The critical information is the strategic business plan, bids for new contracts, and financial information. There are various systems in place to control access to various components, but there are problems with the security components in those systems.

Regardless, the strategic business plan will be scattered throughout the corporation on different desktops and laptops. What is the value of that information? Who has it? Where is it going? In the current environment, few organizations can adequately identify the information, let alone where it is stored within the network.

This situation is even worse in businesses such as government, military, or large corporations where they used to have dozens of filing cabinet to maintain a proper paper trail. Electronic mail has killed the chain of command and the proper establishment of a trail. Information is spread everywhere and important messages simply get deleted when employees leave the company.

The FBI has published a "Top Ten Technology List," which is still current according to the FBI's Awareness of National Security Issues and Response (FBI-ANSIR). This technology list includes:

- manufacturing processes and technologies
- information and communication technologies
- aeronautic and surface transportation systems
- energy and environmental-related technologies
- semiconductor materials and microelectronic circuits
- software engineering
- high-performance computing
- simulation modeling
- sensitive radar
- superconductivity

Many high-tech companies operate within these areas and, as such, are prone to increased incidents of attack and intelligence-gathering operations. Since the primary threat is from internal or authorized users, it becomes necessary to apply security measures within the perimeter.

THE FUTURE OF NETWORK SECURITY

However, the future of network security is changing. The secure gateway will be an integral part of that for a long time. However, implementation of the secure gateway is not the answer in some circumstances. Furthermore, users may be unwilling to accept the performance and convenience penalties created by the secure gateway.

SECURE GATEWAY TYPES

There are two major types of secure gateways — packet filters and application proxy systems — and companies choose one or the other for various reasons. This article does not seek to address the strengths or weaknesses of either approach, but to explain how they are different.

The packet-filter gateway operates at the network and transport levels, performing some basic checks on the header information contained in the packet. (See Exhibit 5.) This means that the packet examination and transfer happens very fast, but there is no logical break between the internal and external network.

The application proxy provides a clear break between the internal and external networks. This is because the packet must travel farther up the TCP/IP protocol stack and be handled by a proxy (see Exhibit 6). The application proxy receives the packet, and then establishes a connection to the remote destination on behalf of the user. This is how a proxy works. It provides a logical break between the two networks, and ensures that no packets from one network are automatically sent to the other network.

EXHIBIT 5 — The Packet-Filter Gateway Operates at the Network and Transport Levels, Performing Some Basic Checks on the Header Information Contained in the Packet

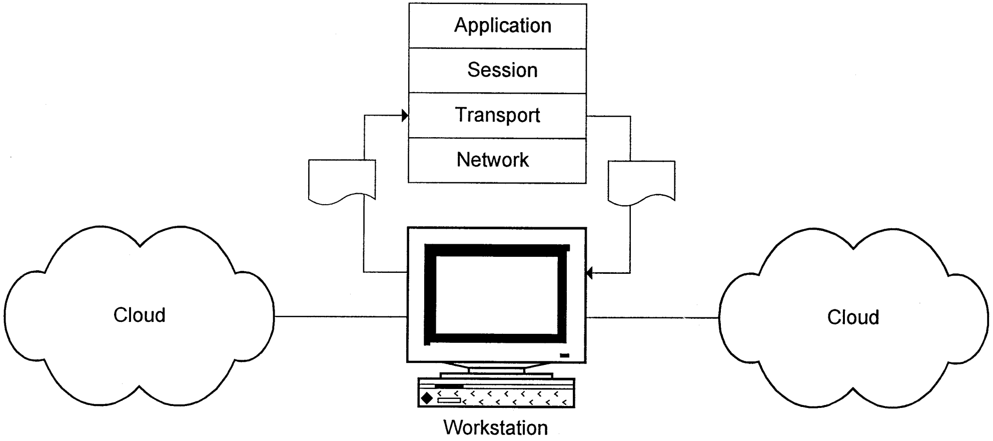

Auerbach Publications

EXHIBIT 6 — An Application Proxy Provides a Clear Break Between the Internal and External Network (This is because the packet must travel farther up the TCP/IP protocol stack and be handled by a proxy.)

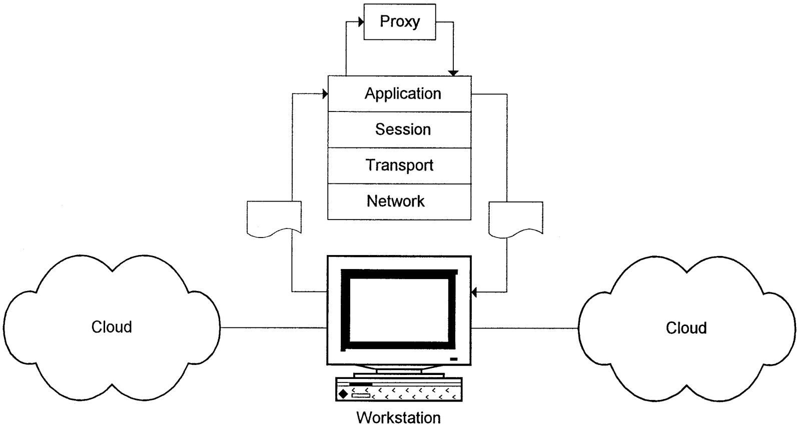

The downside is that there must a proxy on the secure gateway for each protocol. Most secure gateway vendors do not provide a toolkit to build application proxies. Consequently, companies are limited in what services can be offered until the appropriate proxy is developed by the vendor.

The third type of firewall that is beginning to gain attention is the adaptive proxy (see Exhibit 7). In this model, the gateway can operate as

EXHIBIT 7 — With an Adaptive Proxy, the Gateway Can Operate as Both an Application Proxy and a Packet Filter

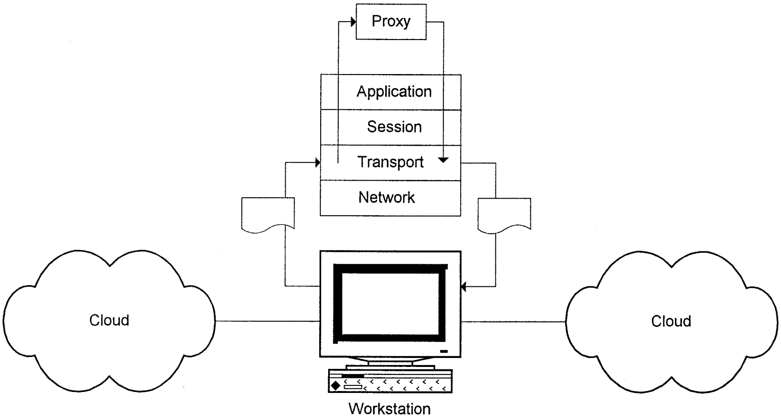

both an application proxy and a packet filter. When the gateway receives a connection, it behaves like an application proxy. The appropriate proxy checks the connection. As discussed earlier, this has an effect on the overhead associated with the gateway.

However, once the connection has been "approved" by the gateway, future packets will travel through the packet filter portion, thereby providing a greater level of performance throughput. There is currently only one vendor offering this technology, although it will expand to others in the future.

The adaptive proxy operates in a similar manner to stateful inspection systems, but it has a proxy component.

Whenever a firewall receives a SYN packet initiating a TCP connection, that SYN packet is reviewed against the firewall rule base. Just like a router, this SYN packet is compared to the rules in sequential order (starting with rule 0). If the packet goes through every rule without being accepted, the packet is denied. The connection is then dropped or rejected (RST is sent back to the remote host). However, if the packet is accepted, the session is then entered into the firewall's stateful connection table, which is located in kernel memory. Every packet that follows (that does not have a SYN) is then compared to the stateful inspection table. If the session is in the table and the packet is part of that session, then the packet is accepted. If the packet is not part of the session, then it is dropped. This improves system performance, as every single packet is not compared against the rule base; only SYN packets initiating a connection are compared to the rule base. All other TCP packets are compared to the state table in kernel memory (very fast).

This means that, to provide increased protection for the information within the corporate network, organizations must deploy security controls within the corporate network that consist of both secure gateways (where there is a good reason) and intrusion and network monitoring and detection. Intrusion detection systems are used in a variety of situations.

SECURITY LAYERING

Security is often layered to provide "defensive depths." This means that at each layer, there are security controls to ensure that authorized people have access, while still denying access to those who are not authorized (see Exhibit 8). As seen in this diagram, this layering can be visualized as a series of concentric circles, with the level of protection increasing to the center.

Layer 1, or the network perimeter, guards against unauthorized access to the network itself. This would include firewalls, remote access servers, etc. Layer 2 is the network. Some information is handled on the network without any thought. As such, layer 2 addresses the protection of the data as it moves across the network. This technology includes link encryptors, VPN, and IPsec. Layer 3 considers access to the server systems themselves. Many users do not need access to the server, but to an application

EXHIBIT 8 — Security Laying Provides Defensive Depths (This means that at each layer, there are security controls to ensure that authorized people have access, while still denying access to those who are not authorized.)

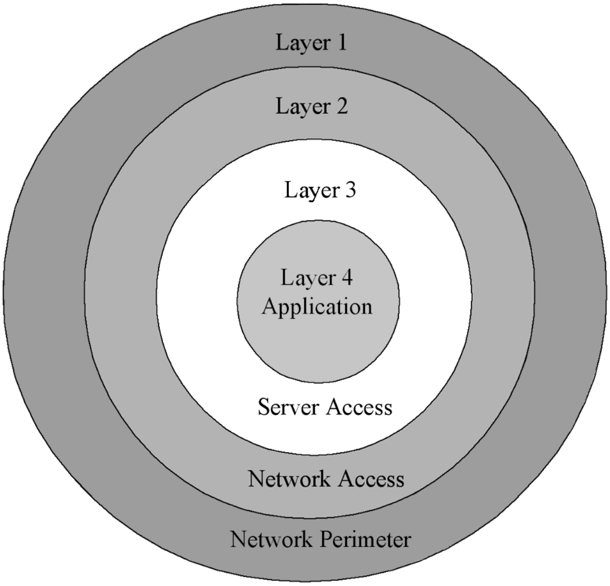

residing there. However, a user who has access to the server may have access to more information that is appropriate for that user. Consequently, layer 3 addresses access and controls on the server itself.

Finally, layer 4 considers the application-level security. Many security problems exist due to inconsistencies in how each application handles or does not handle security. This includes access and authorization for specific functions within that application.

There are occasions where organizations implement good technology in bad ways, which results in poor implementation. This generally leads to a false sense of security and lulls the organization into complacency.

Consequently, by linking each layer, it becomes possible to provide security that the user does not see in some cases, and will have to interact with at a minimal level with to provide access to the desired services. This corresponds to the goals of the three-year architecture vision.

SECURITY GOALS

Organizations place a great deal of trust in the administrators of computer systems to keep things running first, and then make sure that the need-

EXHIBIT 9 — Five Essential Steps in the Information Protection Arena: Protect, Detect, React, Assess, and Correct

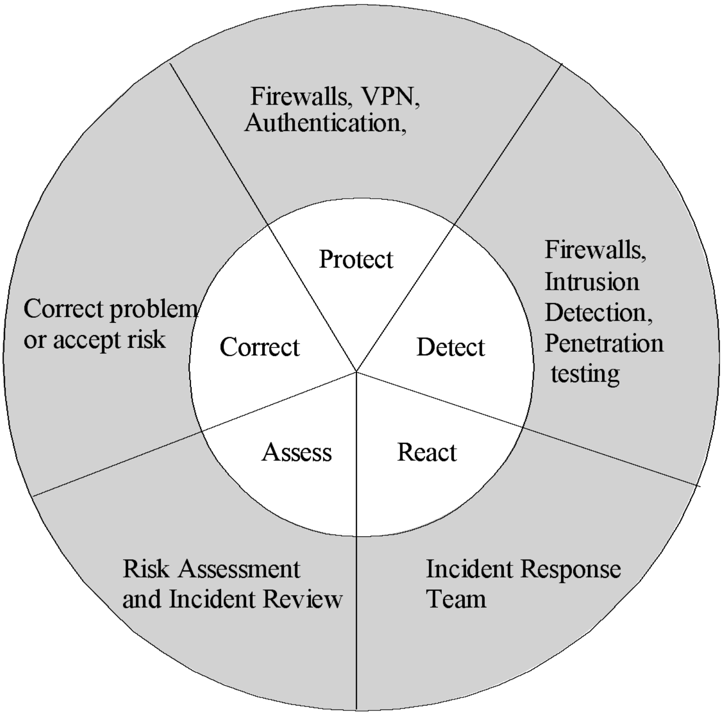

ed patches are applied whenever possible. It is very important that the security measures of any system be configured and maintained to prevent unauthorized access. The major threats to information itself are:

- disclosure, either accidental or intentional (confidentiality)
- modification (integrity)
- destruction (availability)

The goal of an information protection program is to maintain the confidentiality, integrity, and availability of information.

Exhibit 9 illustrates five essential steps in the information protection arena: protect, detect, react, assess, and correct.

Protection involves establishing appropriate policies procedures and technology implementations to allow for the protection of the corporation's information and technology assets.

Detection is the ability to determine when those assets have been, or are under attack from some source.

To be effective at maintaining the security goals of confidentiality, integrity, and availability, the corporation must be able to **react** to a detected intrusion or attack. This involves establishing a Computer Security Incident Response Team to review the alarm and act.

With the tactical response complete, the **assessment** phase reviews the incident and determines the factors that caused it. From there, a risk analysis is performed to determine:

- the risk of future occurrences
- what the available countermeasures are
- a cost/benefit analysis to determine if any of the available countermeasures should be implemented

The **correct** stage is where the countermeasures or other changes are implemented; or, if the level of risk is determined to be acceptable to the corporation, no action is taken.

Many of today's proactive organizations have the protection side operating well, as it relates to network protection. However, many have no systems in place to protect the internal data and network components.

Likewise, reaction mechanisms may be in place to address and investigate when an incident occurs. This is accomplished by establishing a Computer Incident Response Team to be used when an incident is detected in progress that requires the knowledge of a diverse group of computer and security specialists.

However, for many, their detection abilities are limited, which is the area that intrusion monitoring and detection is aimed at. By improving detection abilities, one can refine both protection strategies and technology, and how one reacts when an incident occurs.

Since today's computer systems must be able to keep information confidential, maintain integrity, and be available when needed, it is highly likely that any expectations of the system being able to completely prevent a security breach is unrealistic.

TYPES OF INTRUSION MONITORING AND DETECTION SYSTEMS

There are two major types of intrusion detection: host and network based. Host-based products are based on the computer system and look for intrusions into its own environment. These host-based systems are capable of examining their own configuration and reporting changes to that configuration or to critical files that may result in unauthorized access or modification. For example, a product such as tripwire can be considered a host-based intrusion detection system. Changes in the configuration of the system or its files are detected and reported by tripwire and then captured at the next report.

Network-based products are those that are not bound to looking at intrusions on a specific host. Rather, they are looking for specific activity on the network may be considered malicious. Network-based tools have the ability to find the attack in progress, while host-based tools can actually see the changes inside the system. In fact, it is recommend that one runs both types of systems.

There are essentially two types of intrusion detection "engines." These are statistical anomaly detection and pattern-matching detection engines. Statistical engines look at deviation from statistical measurements to detect intrusions and unusual behaviours. The baseline established for the statistical variables is determined by observing "normal" activity and behavior. This requires significant data collection over a period of time to establish this "normal" or expected behavior. Statistical anomaly systems are generally not run in real time due to the amount of statistical calculations required. Consequently, they are generally run against logs or other collected data.

Statistical anomaly systems offer some advantages. The well-understood realm of statistical analysis techniques is a major strength so long as the underlying assumptions in the data collection and analysis are valid. Statistical techniques also lend themselves better to analysis dealing with time.

However, the underlying assumptions about the data may not be valid, which causes false alarms and erroneous data reported. The tendency to link information from different variables to demonstrate trends may be statistically incorrect, leading to erroneous conclusions. The major challenge to this technique is establishing the baseline of what is considered expected behavior at the monitored site. This is easier if the users work within some predefined parameters. However, it is well-known that the more experienced users are, the less likely they will operate within those parameters.

One drawback to intrusion detection systems is false-positive alarms. A false positive occurs when the intrusion detection system causes an alarm when no real intrusion exists. This can occur when a pattern, or series of packets, occurs that resemble an attack pattern but are in fact legitimate traffic.

Worth noting is that some of the major issues with statistical engines involve establishing the baseline. For example, how does one know when a user has read too many files?

Pattern-matching systems are more appropriate to run in real- or near-real time. The concept is to look at the collected packets for a "signature," or activities that match a known vulnerability. For example, a port scan against a monitored system would cause an alarm due to the nature of packets being sent. Due to the nature of some of the signatures involved, there is some overlap between the pattern-matching and anomaly-detection systems.

The attack patterns provided by the vendors are compiled from CERT advisories, vendor testing, and practical experience. The challenge is for the vendor to create patterns that match on a more general class of intrusion, rather than being specific to a particular attack.

There are pros and cons to both types, but it is recommended that in the development of the tools, that both forms be run. This means collecting the packets and analyzing them in near-real time and collecting the log data from multiple sources to review it with an anomaly system as well.

In a pattern-matching system, the number and types of events that are monitored are constrained to only those items required to match a pattern. This means that if one is only interested in certain types of attacks, then one does not need to monitor for every event. As previously stated, the pattern matching engine can run faster due to the absence of the floating-point statistical calculations.

However, pattern-matching systems can suffer from scalability issues, depending on the size of the hardware and the number of patterns to match. Even worse is that most vendors do not provide an extensible language to allow the network security administrator to define his own patterns. This makes adding one's own attack signatures a complicated process.

For both systems, neither really has a "learning" model incorporated into it, and certainly none of the commercial intrusion detection systems have a learning component implemented in them.

WHY INTRUSION MONITORING AND DETECTION?

The incorporation of intrusion monitoring and detection systems provides the corporation with the ability to ensure that:

- *protected information is not accessed by unauthorized parties; and if it is, there is a clear audit record.* Organizations must identify the location of various types of information and know where the development of protected technologies takes place. With the installation of an intrusion detection system within the corporate network, one can offer protection to that information without the need for a secure gateway. The intrusion detection system can monitor for connection requests that are not permitted and take appropriate action to block the connection. This provides a clear audit record of the connection request and its origination point, as well as preventing the retrieval of the information. There is no impact to the authorized users.
- *the ability to monitor network traffic without impact to the network.* A secure gateway is intrusive: all of the packets must pass through it before they can be transmitted on the remote network. An intrusion monitoring system is passive: it "listens" on the network and takes appropriate action with the packets.
- *actively respond to attacks on systems.* Many implementations of intrusion monitoring systems have the ability to perform specific actions when an event takes place. Those actions range from

notification to a human to automatic reconfiguration of a device and blocking the connection at the network level.

- *information security organizations understand the attacks being made and can build systems and networks to resist those attacks.* As attacks are made against the organization, reviewing the information captured by the intrusion monitoring system can assist in the development of better tools, practices, and processes to improve the level of information security and decrease the risk of loss.
- *metrics reporting is provided.* As in any program, the ability to report on the operation of the program through good quality metrics is essential. Most organizations do not know if there has been a successful penetration into their network because they have no good detection methods to determine this.

IMPLEMENTATION EXAMPLES

As more and more organizations enter the electronic-business (E-biz) forum in full gear, the effective protection of those systems is essential to being able to establish trust with the customer base that will be using them. Monitoring of the activity around those systems will ensure that one responds to any new attacks in an appropriate fashion, and protects that area of the business — both from financial and image perspectives.

Implementing an intrusion monitoring and detection system enables monitoring at specific sites and locations within the network. For example, one should be immediately concerned with Internet access points and the extranets that house so many critical business services on the Internet.

Second, organizations should be working with information owners on the top-ten FBI list, on how to handle corporate strategic information. That venture would involve installing an intrusion monitoring system and identifying the information that people are not allowed to access, and then using that system to log the access attempts and block the network connections to that information.

The following examples are intended to identify some areas where an intrusion monitoring system could be installed and the benefits of each.

Monitoring at the Secure Gateway

In Exhibit 10, the intrusion monitoring system is configured to monitor the networks on both sides of the firewall. The intrusion monitoring system is unable to pass packets itself from one side to the other. This type of implementation uses a passive or nonintrusive mode of network data capture.

To illustrate this, first consider the firewall. The firewall must retransmit packets received on one network to the other network. This is intrusive as the packet is handled by the firewall while in transit. The intrusion monitoring system, on the other hand, does not actually handle the pack-

EXHIBIT 10 — An Intrusion Monitoring System is Configured to Monitor the Networks on Both Sides of the Fiirewall

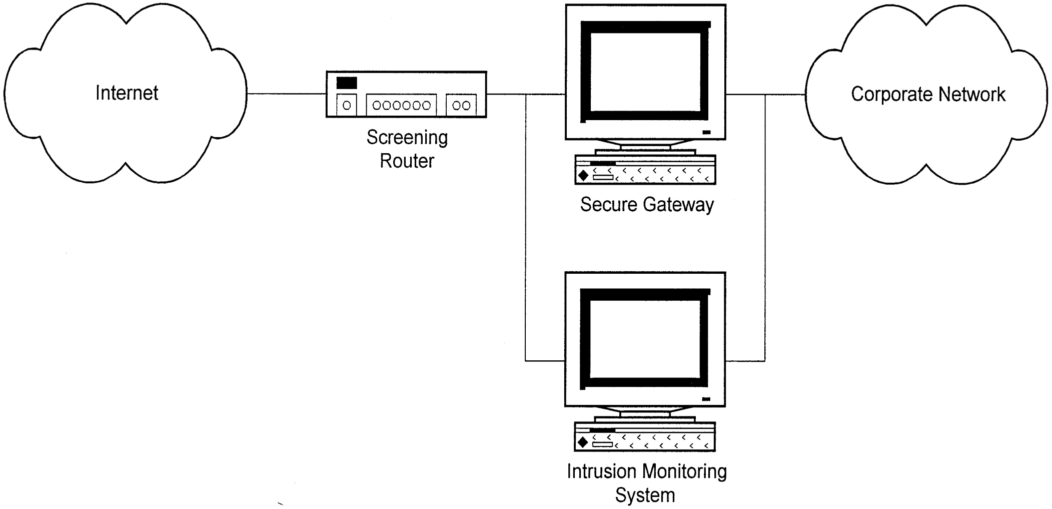

Auerbach Publications

et. It observes and examines the packet as it is transmitted on the network.

This example also lends itself to monitoring those situations where the traffic must be passed through the secure gateway using a local tunnel. As this provides essentially unrestricted access through the secure gateway, the intrusion monitoring system can offer additional support, and improved logging shows where the packet came from, and what it looked like on the other side of the gateway.

Using an intrusion monitoring system in this manner allows metrics collection to support the operation of the perimeter and demonstration that the firewall technology is actually blocking the traffic it was configured to block. In the event of unexpected traffic being passed through anyway, the information provided by the intrusion monitoring system can be used by the appropriate support groups to make the necessary corrections and, if necessary, collect information for law enforcement action.

Monitoring at the Remote Access Service Entry

A second example involves the insertion of an intrusion monitoring device between the RAS access points and their connection to the corporate network (see Exhibit 11). In this implementation, the intrusion monitoring system is installed at the remote access point. With the clear realization that most technical and intellectual property loss is through authorized inside access, it makes sense to monitor one's remote access points. It is possible to look for this type of behavior, active attacks against systems, and other misuse of the corporate computing and network services.

Monitoring Within the Corporate Network

As mentioned previously, there is no ability to monitor specific subnets within the corporate network where protected information is stored. Through the implementation of intrusion monitoring, it is possible to provide additional protection for that information without the requirement for a secure gateway.

Exhibit 12 reveals that the protected servers are on the same subnet as the intrusion monitoring system. When the corporate network user attempts to gain access to the protected servers, the intrusion monitoring server can log and, if configured, intercept the connection attempt. This also means that some guidelines on how to determine where to add an intrusion detection system within the corporate network are required. In many organizations, the corporate network is extensive and it may not be feasible to monitor them all.

Monitoring the Extranet

This will facilitate monitoring attacks against externally connected machines or, in the event that a proper extranet has been implemented, by

EXHIBIT 11 — The Intrusion Monitoring System is Installed at the Remote Access Point

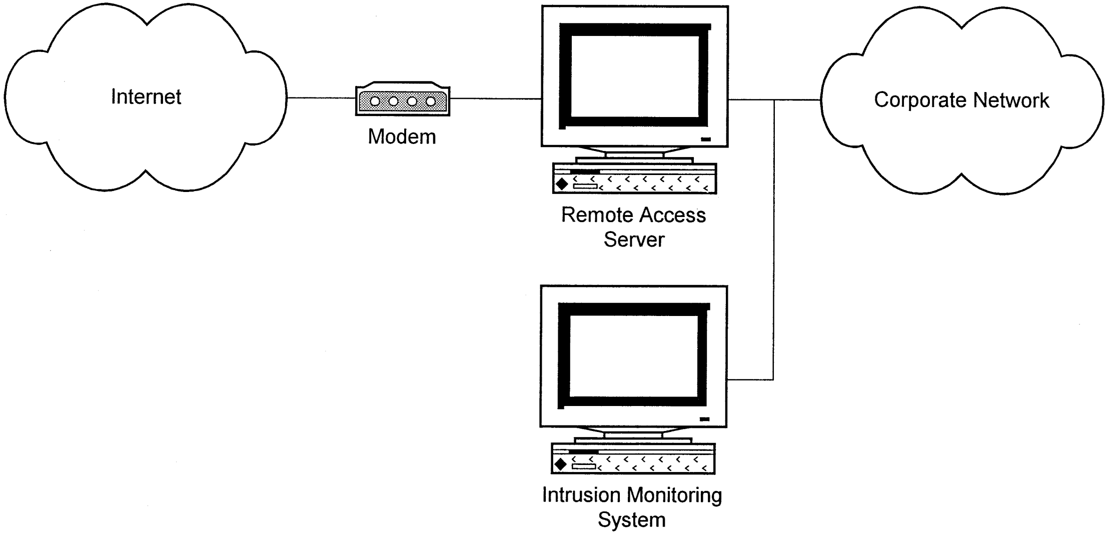

Auerbach Publications

EXHIBIT 12 — Protected Servers Are on the Same Subnet as the Intrusion Monitoring System

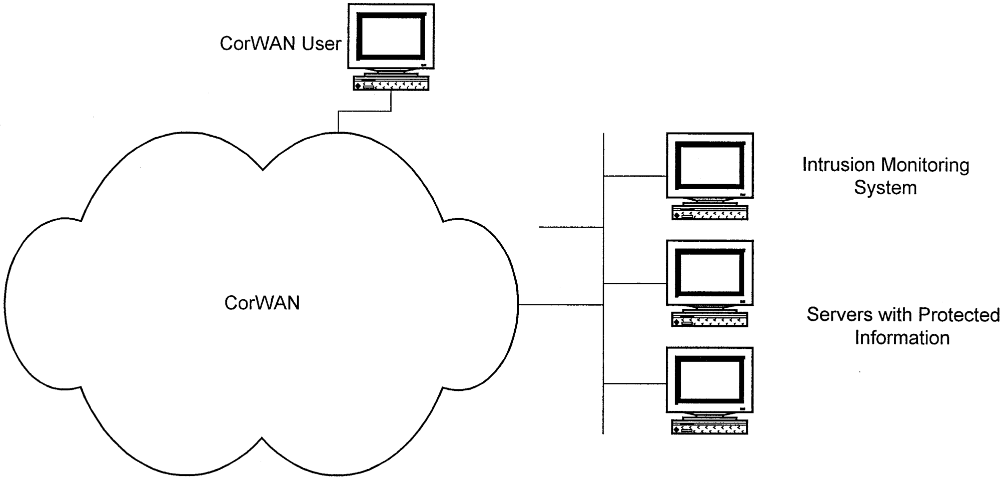
CorWAN User
CorWAN
Intrusion Monitoring System
Servers with Protected Information

EXHIBIT 13 — Two IDS Systems May be Required to Offer Detection Capabilities for Both the Extranet and the Firewall

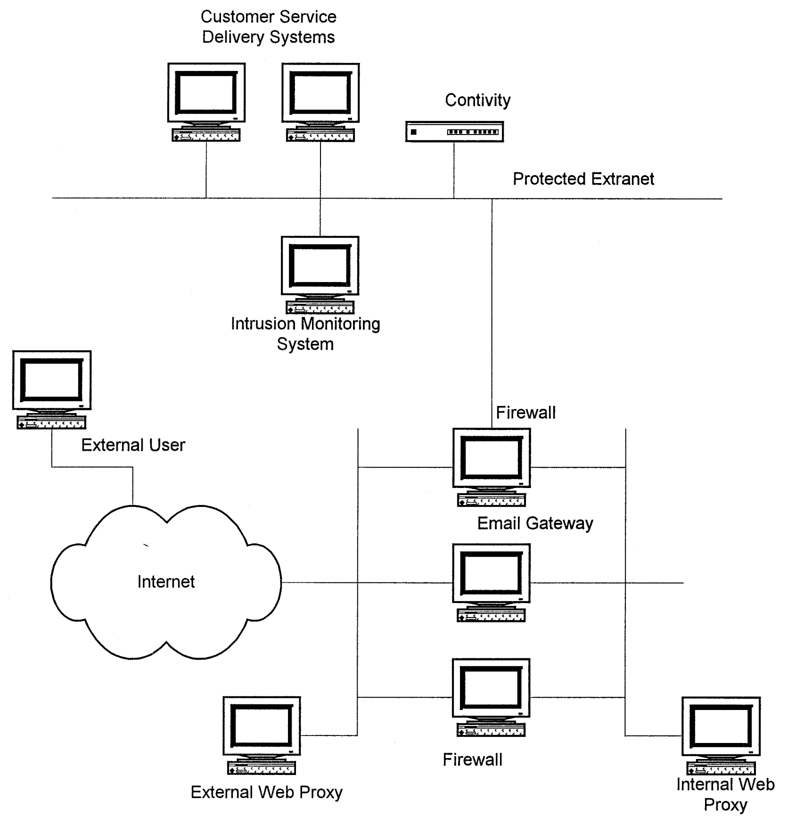

monitoring any attacks against the systems connected to the extranet. However, in this instance, two IDSs may be required to offer detection capabilities for both the extranet and the firewall, as illustrated in Exhibit 13.

In this illustration, all activity coming into the extranet is monitored. The extranet itself is also protected as it is not directly on the Internet, but in a private organizationally controlled network. This allows additional controls to be in operation to protect those systems.

SECURITY IS DIFFICULT TO QUANTIFY

Security is a business element that is often very difficult to quantify. This is because security is a loss prevention exercise. Until something is missing, most people do not bother with it. However, application of an intrusion monitoring system external to network access points can provide valuable information that includes metrics describing the state of the security perimeter.

Aside from the monitoring component, some intrusion detection systems offer the ability to block network sessions where they are deemed inappropriate or undesirable. These systems offer additional opportunities. Deployment of secure gateways can be problematic as the services that are available to users on the external network are reduced due to limitations at the secure gateway. Using the blocking technology, it may be possible to deploy an intrusion monitoring and detection system to monitor the traffic, but also block connection requests to protected information or sites.

PROACTIVE AND REACTIVE MONITORING

The situations illustrated in Exhibits 10 through 13 are proactive implementations of an intrusion detection system. The other implementation (not illustrated here) is reactive. A proactive approach calls for the installation and operation of the system in an ongoing mode, as well as ongoing maintenance to ensure that the intrusion monitor is processing information correctly. A reactive mode approach involves having an intrusion monitor system ready for installation, but not actually using it until some event occurs. The operation of an effective intrusion monitoring systems involves both of these elements.

However, there is the concept of realtime and interval-based intrusion detection. Real time implies that the monitoring agent is run on a continuous basis. Interval based means that the monitor is run as needed, or at intervals. Vulnerability scanning is also seen as a form of intrusion detection by exposing holes in an operating system configuration. This is interval-based monitoring, as it cannot be done all the time.

Information security organizations are often focused on the prevention aspect of network security. They operate systems that are intended to limit access to information and connectivity. This is a proactive activity that requires ongoing analysis and corrective action to ensure that the network is providing the services it should, and that is is properly protected.

COMPUTER INCIDENT RESPONSE TEAM

The benefits of the intrusion detection system (i.e., the ability to detect undesirable activities) will be lost without the ability to respond to it. This is done most effectively through the operation of a Computer Security Incident Response Team (or CIRT). Most CIRT teams are modeled after the Carnegie Mellon Computer Emergency Response Team.

The object of the CIRT is to accept alarms from intrusion detection and other sources. Its role is to review the incident and decide if it is a real incident or not.

The CIRT must include personnel from corporate and information security, internal audit, legal, and human resources departments. Other people may be called in as required, such as network engineering and application providers.

Normally, the alarm is provided to a small group of the CIRT to evaluate. If it is agreed that there is an incident, then the entire CIRT is activated. The operation of the CIRT becomes a full-time responsibility until the issue is resolved. There are a variety of potential responses and issues to be resolved in establishing a CIRT. These are well covered in other documents and will not be duplicated here.

The CIRT forms an integral part of the intrusion detection capability by evaluating and responding to the alarms raised by the intrusion detection systems. As such, the personnel involved must have time dedicated to this function; it cannot take a back seat to another project.

Once the tactical response is complete, the CIRT will closely evaluate the situation and make recommendations for review to prevent or reduce the risk of further occurrence. In the protection cycle, these recommendations are used to assess what further action is to be taken.

This being the case, a decision to implement intrusion detection is a decision to implement and support a CIRT. Intrusion detection cannot exist without the CIRT.

PENETRATION AND COMPLIANCE TESTING

The best method to test security implementation is to try it out. A penetration test simulates the various types of attacks — both internal and external, blind and informed — against the countermeasures of the network. Essentially, a penetration test attempts to gain access through available vulnerabilities.

Penetration testing is part of the detection strategy. While intrusion detection capabilities are required to monitor access and network status on an ongoing basis, penetration is an interval-based targeted approach to testing both the infrastructure, and the detection and reaction capabilities.

Penetration testing should be done as part of the network security strategy for several purposes:

- *To provide confidence or assurance of systems integrity:* Vulnerability scans often do not include attempts to exploit any vulnerability found, or any of the long list of known vulnerabilities. This is because many of the systems being tested currently are in production. A successful penetration test could seriously affect normal business operations. However, the integrity of the system can be effectively tested in a nonproduction role.
- *To verify the impact of the security program:* Penetration testing is used to determine if the security program is performing as it should.

There are a number of different products and services that work together to provide this infrastructure. Each can be evaluated on its own, but it is much more complicated to test them as a system.

- *To provide information that can be used in developing and prioritizing security program initiatives:* Any issues found during a penetration test can alter and affect the direction of the security program priorities. Should a major issue be found that requires correction, the security program goals may be altered to provide a timely resolution for the issue.
- *To proactively discover areas of the infrastructure that may be subject to intrusion or misuse:* People do not install an alarm system in their house and never test it. The same is true here. Ongoing evaluation allows for the identification of components in the infrastructure that may be less secure than desired, not operating as expected, or contain a flaw that can be exploited. Taking a proactive stance means that it becomes possible to find and correct problems before they are exploited.
- *To provide information that can be used in developing and prioritizing policy initiatives:* Policy is not cast in stone; it must be updated from time to time to reflect the changing needs of the business. Penetration tests can assist in the testing and development of policies. This is done using the information learned from the testing to evaluate whether one is compliant with the policies, and if not, which is correct — the implementation or the policy.
- *To assess compliance with standards and policies:* It is essential that the infrastructure, once in operation, be compliant with the relevant security policies and procedures. This verification is achieved through penetration testing, or what is also known as protection testing. Protection testing is the same as penetration testing but with a slightly different objective. While penetration testing attempts to find the vulnerabilities, protection testing proves that the infrastructure is working as expected.
- *To provide metrics that can be used to benchmark the security program:* The ability to demonstrate that the security infrastructure is operating as expected, and that improvement is visible, are important parts of the program. Metrics establish what has been *and* what is now. It is also possible from collected to metrics to make "educated guesses" about the future. By collecting metrics, one also gathers data that can be used to benchmark the operation of our infrastructure as compared to the companies.
- *For preimplementation assessments of systems or services:* It is important that appropriate evaluations are performed to ensure that the addition of new services to the infrastructure, or that is dependent on the infrastructure operating correctly, be certified to ensure that no vulnerabilities exist that could be exploited. When a new application

is developed that interconnects both internal and external systems, a penetration test against the application and its server is undertaken to verify that neither holds a vulnerability to be exploited. This also ascertains that if the external system is compromised, that the attacker cannot gain access to the corporate network resources.

Types of Penetration Tests

There are essentially three major types of penetration testing, each with their own tools and techniques:

- **Level 1 — Zero Knowledge Penetration Testing:** This attempts to penetrate the network from an external source without knowledge of its architecture. However, information that is obtained through publicly accessible information is not excluded.
- **Level 2 — Full Knowledge Penetration Testing:** This attempts to penetrate the network from an external source with full knowledge of the network architecture and software levels.
- **Level 3 — Internal Penetration Testing:** This attempts to compromise network security and hosts from inside one's network.

Penetration testing is interval based, meaning that it is done from time to time and against different target points. Penetration testing is not a realtime activity.

The process consists of collecting information about the network and executing the test. In a Level 1 test, the only information available is what is published through open source information. This includes network broadcasts, upstream Internet service providers, domain name servers, and public registration records. This helps simulate an attack from an unsophisticated intruder who may try various standard approaches. This approach primarily tests one's ability to detect and respond to an attack.

A Level 2 penetration test assumes full knowledge of the hardware and software used on the network. Such information may be available to meticulous and determined intruders using whatever means, including social engineering, to increase their understanding of your networks. This stage of the test assumes the worst-possible scenario and calls to light the maximum number of vulnerabilities.

A Level 3 penetration test, or acid test, is an attack from within the network. This is the best judge of the quality of the implementation of a company's security policy. A real attack from within a network can come from various sources, including disgruntled employees, accidental attacks, and brazen intruders who can socially engineer their way physically into a company.

Penetration testing should be considered very carefully in the implementation of an overall detection program, but it can lead to the negative side effects one is trying to prevent. Therefore, penetration testing should be used cautiously, but still be used to attempt to locate vulnerabilities and to assess the overall operation of the protection program.

SUMMARY

This article has presented several implementations of secure gateway and intrusion detection techniques, while focusing on the business impact of their implementation. It is essential that the security professional consider the use of both network and host-based intrusion detection devices, and balance their use with the potential for impact within the operating environment.

A key point worth remembering is that the implementation of technology is only part of the solution. There must be a well-thought-out strategy and a plan to achieve it.

Chris Hare, CISSP, ACE, works at Nortel.