

Security

Key strategies for safeguarding IT resources from interior and exterior threats

TABLE OF CONTENTS

- 2** A Multilayered Security Approach
- 2** Interior Network Threats
- 3** Data Leakage Prevention
- 4** Secure Remote Access
- 5** Threats to Remote Access
- 6** Endpoint Security
- 6** Threats to Endpoints
- 8** Manufacturer Options

Executive Summary

Systems will never be totally secure, as hackers keep coming up with new uses for old methods of attack. Over the past several years, there has also been an erosion of network perimeters as organizations strive to enable roving users and facilitate interoperation with partners.

In addition, regulatory requirements have upped the stakes for failure. A security lapse might now imply fines, operations sanctions or other penalties if the victim organization fails in its security responsibilities.

In response to these developments, many organizations have changed tactics. Instead of devoting their efforts to protecting their computers, the focus has shifted to protecting the critical data residing on them.

This might seem like a superficial change. However, changes in application architecture, host-based data leakage prevention tools and encryption technologies allow organizations to protect at least some information — even under conditions where an attacker has compromised portions of the target computer.

In the long view, all of this is good for security. Threats do not remain static, and so security measures against those threats must constantly evolve as well. This white paper provides a brief update to help your organization freshen up its security strategies and keep you secure against the latest developments in threats.

.....

A Multilayered Security Approach

An informed, up-to-date security approach includes multiple layers of defense. This approach has evolved as a response to developments in network design. Many organizations now have multiple Internet uplinks, making it impractical to deploy a traditional firewall sandwich (utilizing two different brands of firewalls at a perimeter location) defense.

Even in a hub-and-spoke network where remote sites only use the Internet for virtual private network (VPN) connectivity with the central site, each remote site necessarily has an Internet connection. The more remote sites there are, the higher the probability that one of them will experience a security incident.

In order to secure these remote sites, organizations need to apply a multilayered security strategy. The following measures can reduce these threats significantly:

- Apply access control lists on remote site firewalls to restrict outbound traffic to only necessary services — typically those provided by the server virtual local area network (VLAN) at the main office.
- Apply access control lists to the main site firewall, restricting inbound traffic from the remote sites to only the application protocols served up by the central servers.
- Implement a configuration management system to keep the many remote site firewalls, routers and switch configurations in sync.
- Turn on the intrusion detection feature sets in the remote site firewalls.
- Implement a security event management system to collect and analyze log and alert information from infrastructure devices across the network.

While there are other types of threats out there, gateway security cannot be completely forgotten. Intrusion prevention systems (IPS) continue to be indispensable in helping secure the perimeter.

An IPS is a solution that looks more closely at the content coming into the network than a firewall can. It can also look for known methods of attack and can be fine-tuned to look for attacks targeting the platforms and applications used by your organization.

Security as a Process

Security is an evolving, never-ending process. This process, or lifecycle, consists of four phases: design, implementation, testing and monitoring.

- **DESIGN:** A system's security begins with a definition of its functional requirements and the development of a design able to meet them. Adding security features onto a design as an afterthought is ineffective and costly.
- **IMPLEMENTATION:** Even the best design may not turn out as planned when executed, so it's critical to pay attention to security concerns as a design is translated into a working system.
- **TESTING:** Once a system is ready to be deployed, it's crucial to verify that its security features function properly and don't expose the system to unnecessary risk. Outside expertise is needed to make sure that the testing covers all the areas it should address.
- **MONITORING:** It's important to remember that accounting is a key component of a system's security capabilities. Nearly everything in the IT environment has the capacity to keep some sort of log or send alert messages.

Interior Network Threats

Sensitive data may be at risk in foreign environments where we already know enough to be wary, but the problem is present in subtle forms within our own networks as well.

Nearly every organization has some pool of highly confidential data: Social Security numbers, health records, student grades/transcripts, tax and property records and so on. Access to these materials is often carefully restricted so that only a small group of people can view a database or folder of files.

The organization feels safe because these users only print to a printer in their area, with physical access to that area well controlled. Yet

print jobs are rarely protected as they traverse the network on the way to the printer.

Anyone physically in a position to capture this traffic (and software to capture network packets is freely available) would be able to reconstruct the print job and produce their own copy of the document. How easy would it be to intercept that traffic and extract the pertinent information?

Detecting and Preventing Intrusions

Since the inside-out approach to gateway-network security has blurred a bit, organizations really need to be aware of threats inside their network. Security has to go beyond simply watching for suspicious behavior. The user's identity needs to be factored in, too.

For example, internal network security systems should raise a flag if someone from an organization's HR department (without authorization) tries to access financial records. This is because, based on the user's identity, they would not typically be given access to financial records.

Organizations need to set policies that specify who can go where inside the network. There are security solutions available that allow organizations to set access policies based on user identity, device security state and location information that is session-specific.

Intrusion detection and prevention (IDP) solutions are also invaluable for internal security. IDP software and hardware solutions protect the network from a wide range of attacks, as well as provide information on rogue servers. They can also raise alerts on types and versions of applications and operating systems that may have unknowingly been added to the network.

Protect Your Network Chokepoints

To protect the core of your organization's data, place intrusion prevention devices (IPS) not at just the Internet connection, but at the chokepoints where hackers can grab big chunks of data.

Add an IPS unit at these chokepoints:

- The aggregation point of your wireless network
- The frame relay or WAN connection
- In front of your database server farm

Data Leakage Prevention

For most organizations, data is their most valuable resource, so they need to keep it secure and confidential. This can be difficult, as a variety of staffers often end up accessing private data, so it tends to accumulate in unexpected places and is handled unsafely or is disclosed inadvertently.

Every time a file or document is stored or moved, it can be intercepted or simply left out for later abuse. Data leakage prevention focuses on restricting the flow of private information across organizational boundaries. Organizations need to set policies that govern how information is handled.

For example, a policy may prohibit sending e-mails that contain Social Security numbers or financial information over instant messaging systems that are not logged. Ideally, organizations would enforce such policies by technical means, but often this strategy is not completely feasible.

Removable and Portable Media

Any time information is put into digital form it can be lost or stolen. Storage media are portable and easy to misplace. It's also hard to know what happens to media once a user is finished with them. If a USB drive is used to bring a file to another site, what happens to the file once it has been copied? Most often, it still resides on the USB drive.

If the USB drive is lost, the contents go with it — likewise with CDs, DVDs and floppies. When copying the contents onto hard drives, staff rarely take the final step of destroying the disks or wiping them clean.

To help address the potential for data leakage, many organizations have deployed smartphones with "remote-kill" functionality or full disk encryption on their notebooks. This can certainly help, but an approach of this type has limitations.

For example, law enforcement officers now place phones that they confiscate in bags lined with wire mesh, which prevents remote kill and ensures the recovery of information residing in those devices. If a phone can't receive instructions to wipe its memory, it is vulnerable.

Likewise, whole disk encryption protects only the copy of a device's data written to the hard drive. If a device is captured while powered on, a great deal of sensitive information may reside in memory in decrypted form.

E-Mail

E-mail is an even worse problem for data leakage than portable media. When replies to messages are chained together, the e-mail thread forwarder seldom checks to see whether the whole thread is suitable for the recipient. Files of all sorts get attached to e-mail, and these can contain sensitive information.

In fact, the sender of a file may not even be aware of its full content. For instance, spreadsheets can contain hidden rows; word-processing documents may contain historical data about changes and edits; presentations can contain embedded objects and metadata that give clues about confidential data.

Finally, many users remain unaware of the types of information they should avoid exchanging via e-mail. What is considered acceptable e-mail usage within an organization may not be safe for transmission across the Internet.

Once e-mail is sent, it can't reliably be recovered or erased. It can linger indefinitely in inboxes over which you have no control, and others can forward it onto recipients you have not authorized.

Other Transmissions and Synthesis

As instant messaging has become an increasingly popular, it too has presented security liabilities. Many organizations have deployed technologies that inspect outbound e-mail for suspicious content, but instant messaging can bypass these controls and thereby provide another avenue for exporting the organization's internal secrets.

Blogs and social networks are also capable of leaking data. Most organizations have a policy regarding staff participation in these online activities. But even if users carefully refrain from posting any private data, an attentive observer can certainly draw inferences that amount to a disclosure of sensitive data.

For example, web logs contain the IP addresses of visitors. Even if a discussion in an online forum never mentions the name of a participant's organization, the source IP address may provide a dead give-away.

Whenever an organization has serious confidentiality obligations, as many government and educational organizations do, it's important to keep in mind that comments and photos posted to the Internet will be archived.

The amount of information cataloged by search engines will only increase over time. As a result, a comment that contains no specific harmful or incriminating information today might have a completely different meaning when placed in the context of a long series of posts over time.

This risk, sometimes called "information synthesis," is not a new one by any means. It's receiving new attention now because of the volume of online content that is cataloged and searchable.

Secure Remote Access

Today, staff members are often expected to keep up with tasks while on the road or at home, and notebooks have become standard-issue in many organizations. This means that people will be accessing their organization's networks remotely. Remote access must be as secure as possible. There are a number of key steps to ensuring this security.

Varieties of Remote Access

Remote access encompasses any access to private network resources from beyond the organization's physical perimeter. This characterization of remote access includes some modes that aren't usually taking into consideration in security planning.

Remote access includes the following scenarios:

- The typical understanding of remote access: Users log in via a VPN to do work on internal systems.
- Vendors and other partners may have access to support or administer systems on the organization's network.
- On-the-road users make use of lightweight remote access solutions such as web e-mail or various support tools.
- Wireless networking signals, whether from access points or from endpoints, seep out of buildings. Unauthorized users may receive those signals and connect to systems on internal networks.

Though this list touches on the major modes of remote access, each of these might have many variations. VPN, for example, can encompass Internet protocol security (IPsec), Secure Sockets Layer (SSL), point-to-point tunneling protocol (PPTP), secure shell (SSH) and other protocols.

Unified Threat Management

In its bid to stay ahead of new attack strategies, has your organization acquired a mountain of security boxes and services that aren't integrated and all use different management systems? Then consider a unified threat management (UTM) solution, a single appliance that implements a firewall's intrusion protection plus a range of other security features, including virus protection, content filtering, and spam, phishing and spyware blockers.

Threats to Remote Access

Having identified the basic types of remote access, it's time to start thinking about what might go wrong. Threats to remote connectivity can target the establishment of the connection, information in transit, the endpoints themselves or the ability of either end to transmit or receive.

Denial of Service

Denial of service refers to an attempt to make remote connectivity usage impossible. Many ways exist to accomplish this goal:

- **RESOURCE CONSUMPTION:** An attacker can overwhelm the systems that provide connectivity either with traffic floods that crowd out legitimate transactions or with requests for services from the endpoints.
- **CRASHING:** An attacker can send specially crafted network traffic designed to disrupt the operations of either endpoint, typically by crashing the system or rendering some portion of it inoperative.
- **LOCKOUT:** With a list of users, an attacker can supply intentionally incorrect passwords for each account until all users are locked out.

The unavailability of remote access functionality can have a serious impact on an organization's security, especially if the organization relies on remote access for any critical tasks.

Traffic Analysis

Even though a remote access solution may protect the contents of communication from eavesdropping, traffic interception can still prove a problem. An eavesdropper can learn a great deal from encrypted traffic, despite the fact that the actual contents of packets may not be accessible.

The eavesdropper can see what the endpoints of the conversation are. For some organizations, this is can be dangerous. The contents of any given interaction might be innocuous, but the mere fact that communication is taking place can expose a person or project to risk.

Even if the existence of encrypted traffic doesn't pose a problem, an observer can often make powerful inferences about patterns in operations by looking for deviations from normal patterns.

Attacks on Authentication Credentials

The possibility of unauthorized users accessing sensitive resources remotely poses a serious risk. By supplying the right credentials, an attacker can join the organization's VPN, or establish a dial-up connection, or get onto the wireless network and proceed from there.

An organization can handle authentication for remote access in several ways. Standard user name/password login, two-factor tokens and digital certificates are the most common, but others exist.

With each of these approaches, users prove their identity by providing some combination of secrets (or information based on secrets) that only a legitimate user should have. Several possible vectors of attack exist:

- **BRUTE FORCE:** With a legitimate user name, an attacker can systematically begin supplying guesses as to the secret information, with the goal of eventually trying all possible alternatives.
- **SHALLOW DICTIONARY:** If an attacker can access or generate a list of user names, it's easy to supply a list of simple guesses. A typical attack might involve supplying the following guesses for each account: blank password, password is "password" and password is the same as the user's name.

Even with a secure underlying means of transporting data, a remote access solution can easily be compromised if authentication mechanisms are weak.

Man in the Middle

Organizations must also be concerned with the threat of an attacker impersonating the remote access system itself. In a man-in-the-middle attack on a remote access system, the adversary convinces the remote user that an imposter system is, in fact, the legitimate source of remote-access connectivity.

When the user connects, the attacker simply takes the same credentials and passes them along to the actual remote access on-ramp, impersonating the user's endpoint system. This is a particularly insidious attack: The interceptor can inspect and alter all traffic on its way between the user and the organization's internal network.

Endpoint Security

While the servers that store and process information may be secure in the data center, the vast majority of our interactions with computers take place on network endpoints: workstations, notebooks, PDAs and the like. A successful security strategy will ensure that endpoints are secure.

Varieties of Endpoints

The list of endpoint varieties is longer than you might expect:

- Workstations, terminals and notebooks owned by the organization and used locally are the most obvious endpoints and represent the main avenue for users to interact with network resources.
- Workstations and notebooks can also be used remotely, and some might not belong to the organization.
- Though often overlooked, network-connected printers and fax machines are endpoints as well. Such devices are often overlooked as a path for sensitive information to enter and exit the network.
- Smartphones, PDAs, music players and other multifunction devices are endpoints. They often have large storage capacities and can be used to store confidential files or e-mail.

Each of these device types have different security limitations. As a result, varying constraints exist around what network and system administrators can do with them.

Protecting Inside the Network

Protecting along the entry points of your network is important, but don't forget internal network layers of protection. Some additional tips to secure your network internally include:

- Aggregating wireless links through a single gateway, and having intrusion protection.
- Placing web application firewalls in front of your web servers.
- Setting up database auditing devices or firewalls in front of your database servers.
- Adding encryption to your databases.
- Making sure your VPN connection is not bridged directly to your network, and having a device that examines all the packets that have been decrypted.

Threats to Endpoints

Because we tend to focus attention on centralized services or massive repositories of critical data, network endpoints are often neglected security-wise, but they need protection. Attacks on network endpoints can put an organization's private information at risk the same as attacks on the central resources.

Unauthorized Acquisition

The loss of hardware is problematic, but probably not as big an issue as the loss of the information that resides in it. Nearly all endpoint devices store some valuable information, whether proprietary data, internal communications or cached passwords.

Most printers and multifunction copiers contain hard drives where images of recently processed documents might persist, perhaps along with network credentials. It's possible for someone in physical possession of one of these devices to extract sensitive data from them.

Many organizations consider and address the threat of equipment theft, but they should also bring a security mindset to how they approach systems repair and expired asset disposal. Each of these situations represents an opportunity for someone to gain access to stored data.

The most common method for protecting against unauthorized possession of your physical devices is to encrypt the data on them. Properly managed cryptosystems can protect private information from unauthorized access in the event that the physical device or media falls into unauthorized hands.

Some strategies involve encrypting all data on a device (such as whole disk encryption), while others selectively encrypt only sensitive materials. The difficulty with the latter approach is that it's not always easy to ensure that all sensitive information receives protection.

Subversion

Encryption is suitable protection for information stored on powered-off systems, but information is unencrypted when in use. As a result, if an attacker gains control of a system while it is up and running, a great deal of information that would otherwise be encrypted is accessible.

For example, spyware typically has access to whatever materials the logged-in user can see; bots running as background services might be able to access anything that the operating system can. If a hacker finds a means to execute code on a network endpoint, it becomes possible to read files, log keystrokes and capture screenshots of any user's activity.

Also, the compromised system might contain stored passwords that would give the hacker access to other systems: help-desk accounts, cached domain login credentials, service account passwords, passwords to websites, VPN authentication materials and so on.

These possibilities make it critical to protect network endpoints against attack. Host-based intrusion prevention, antivirus/malware products, local firewalls, diligent system administration and user awareness training are all key components in a program to protect network endpoints.

Eavesdropping

Network endpoints rely on other network resources in order to be productive. This means that sensitive information (such as the data that the endpoint handles, as well as authentication credentials) necessarily traverses the network.

So network traffic always remains vulnerable to eavesdropping. Passively listening to network traffic can, in fact, reveal a great deal of private data, system administration practices and even passwords for services.

Remote network endpoints have always dealt with the problem of eavesdropping: VPNs and cryptographically protected application protocols such as SSH or hypertext transfer protocol over Secure Sockets Layer (HTTPS) make it possible to conduct private transactions on public networks without major information leakage.

However, these measures also create a false sense of security because users don't realize all the avenues for disclosure of their passwords. Suppose a user fires up a notebook, connects to the Internet, checks web-based e-mail, hops on the organization's VPN and signs into an instant messaging service. An eavesdropper can capture all of this traffic.

In this example, the VPN logon may be secure, but the other two may not be. Moreover, the instant messaging logon might happen automatically, without the user even thinking about it.

If the user employs the same password for more than one of these systems, the eavesdropper could read the user's e-mail. The potential may also exist for the eavesdropper to access resources over the VPN. This is because, in addition to the instant messaging credentials, the attacker has likely seen the addresses of the VPN concentrator and the web mail server.

Impersonation

When a network endpoint sends a user name and password, how does it know that the mail server to which it sends is the intended recipient and not an imposter? All too often, the endpoint simply requests the address of a remote system from a domain name server (DNS) and takes the resulting response automatically.

We're familiar with the requirement that remote users need to prove their identities before they can access private systems. In fact, nearly everything requires a password, a certificate, a two-factor token or a thumbprint. On the other hand, a great many application protocols don't require that the remote system prove its identity to the user.

HTTPS, for instance, uses SSL certificates to prove that remote web servers are what they appear to be. Unfortunately, incorporating such authentication measures is often impossible or requires additional layers of configuration.

If an attacker controls DNS, all kinds of mischief can ensue. Moreover, a great deal of network interaction takes place behind the scenes. In particular, antivirus software, software update facilities, instant messaging applications and many other applications pull information from the network, sometimes without any means of verifying the authenticity of the source.

The problem does not end merely with network name lookups; it extends to the network itself. Users are accustomed to selecting the wireless networks to which they'll connect by name — or allowing their notebooks to do this automatically. What assurance do users have that the access point broadcasting a particular ESSID (network name) truly belongs to the network they have in mind?

By cryptographically protecting the network, users can breathe easier. Open wireless networks, on the other hand, can be easily impersonated.

Manufacturer Options

Organizations have many options for finding the right security solutions for their needs. Here is a breakdown of some of the key gateway and network security solutions and their manufacturers:

FIREWALLS are a good frontline defense, preventing unwanted traffic from entering the network. They can be either a software or a hardware appliance. Leading manufacturers of firewalls are SonicWALL, WatchGuard, Cisco, Netgear and Barracuda.

CONTENT, E-MAIL AND SPAM FILTERING are additional gateway security tools. Prominent manufacturers of filtering solutions include Jupiter Networks, SonicWALL, McAfee and Websense.

INTRUSION PREVENTION SYSTEMS examine incoming content on a deeper level than a firewall can. Reliable IPS manufacturers include McAfee, TippingPoint, Trend Micro and 3Com.

DATA ENCRYPTION SOFTWARE keeps your data safe regardless of whether the device it is on is safe or not. Top manufacturers for encryption solutions are McAfee and HP.

Here is a breakdown of some of the key remote access solutions and their manufacturers:

VIRTUAL PRIVATE NETWORKS allow for secure remote access onto the network. They can take the form of either software or hardware solutions. Leading VPN manufacturers include SonicWALL, Netgear, Cisco, Barracuda, WatchGuard, Juniper Networks, Check Point and D-Link.

NETWORK ACCESS CONTROL solutions verify the user attempting to access the network, as well as make sure that the user device is security compliant. NAC solutions are typically software based, the most prominent manufacturers being Symantec, Cisco, McAfee and TippingPoint.

KEY FOBS are token devices that enforce two-factor authentication for remote users. RSA is the go-to manufacturer of key fob solutions.

Put It in Writing

Security starts with your organization's users. Educating them about being mindful of security can go a long way toward helping secure your organization. You should craft a security policy that provides general security guidelines for the entire organization.

Guidelines should include the appropriate use of equipment and the Internet and the ramifications if the rules are not followed. A security policy should also include procedures for creating and changing passwords, high-level statements about authentication requirements to get on systems and how information should be stored and destroyed.